

Pilot HandBook



For Flight Simulator

Contents

Introduction

Readership, Topics, and Objectives

How to Use this Book

1 Energy Awareness and Energy Management

1.1 Total Energy Cannot Change

1.2 Energy Conversion

1.2.1 Converting Speed to Altitude and Back

1.2.2 Energy Per Unit Mass

1.2.3 Converting Fuel to Altitude

1.2.4 Power versus Energy

1.2.5 Drag and the Power Curve --- Introduction

1.2.6 Rates of Energy Conversion

1.3 Effect of Controls on Energy

1.3.1 Power Budget --- Using the Engine

1.3.2 The Effects of the Throttle

1.3.3 The Effects of the Yoke

1.3.4 Sizes of Energy Reservoirs

1.4 Energy Management Strategy

1.5 Summary: Energy Management

2 Angle of Attack Awareness and Angle of Attack Management

2.1 The Importance of Angle of Attack

2.2 Definition of Angle of Attack

2.3 Trim for Angle of Attack!

2.4 Three Contributions to Angle of Attack

2.5 Perceiving Pitch Angle

2.6 Making Changes in Angle of Attack

2.7 Fly with a Light Touch

2.8 Trim Won't Solve All The World's Problems

2.9 Pitch Attitude versus Angle of Attack

2.10 Power plus Attitude does not equal Performance

2.11 Estimating the Relative Wind

2.12 Airspeed Is Related to Angle of Attack

2.12.1 Airspeed versus Coefficient of Lift

2.12.2 Coefficient of Lift versus Angle of Attack

2.12.3 Correcting for Reduced Density

2.12.4 Correcting for Reduced Lift Requirements

2.12.5 Correcting for Increased Lift Requirements

2.12.6 Compute with Calibrated not Indicated Airspeed

2.12.7 Correcting for Slip

2.12.8 Drag and Lift-to-Drag Ratio

2.13 Not Everything Depends on Angle of Attack

2.13.1 Explicit Airspeed Limits

2.13.2 Maneuvering Speed

2.13.3 Overview of Limits and Performance Numbers

2.14 Absolute versus Geometric Angle of Attack

2.15 Summary

3 Airfoils and Airflow

3.1 Flow Patterns Near a Wing

3.2 Pressure Patterns Near a Wing

3.3 Stream Line Curvature

3.4 Bernoulli's Principle

3.4.1 Magnitude

3.4.2 Altimeters; Static versus Stagnation Pressure

3.4.3 Compressibility

3.5 Stall Warning Devices

- 3.6 Air Is A Fluid, Not A Bunch of Bullets
- 3.7 Other Fallacies
- 3.8 Inverted Flight, Cambered vs. Symmetric Airfoils
- 3.9 Thin Wings
- 3.10 Circulation
 - 3.10.1 Visualizing the circulation
 - 3.10.2 How Much Circulation? The Kutta Condition
 - 3.10.3 How Much Lift? The Kutta-Zhukovsky Theorem
 - 3.10.4 Quantifying the Circulation
- 3.11 Mechanically-Induced Circulation
- 3.12 Lift Requires Circulation & Vortices
 - 3.12.1 Vortices
 - 3.12.2 Wake Turbulence
 - 3.12.3 Induced Drag
 - 3.12.4 Soft-Field Takeoff
- 3.13 Frost on the Wings
- 3.14 Consistent (Not Cumulative) Laws of Physics
- 3.15 Momentum in the Air
- 3.16 Summary: How a Wing Produces Lift

4 Lift, Thrust, Weight, and Drag

- 4.1 Definitions
- 4.2 Balance of Forces
- 4.3 Forces During a Turn
- 4.4 Types of Drag
- 4.5 Coefficients, Forces, and Power
- 4.6 Induced vs. Parasite Drag

5 Vertical Damping, Roll Damping, and Stalls

- 5.1 Introduction and Overview
- 5.2 Vertical Damping
 - 5.2.1 Origins of Vertical Damping
 - 5.2.2 Loss of Vertical Damping
- 5.3 The Stall
 - 5.3.1 Definition of Stall
 - 5.3.2 Flying Beyond the Stall?
- 5.4 Roll Damping
 - 5.4.1 Origins of Roll Damping
 - 5.4.2 Loss of Roll Damping
 - 5.4.3 Schemes to Increase Roll Damping
- 5.5 The Effect of Flaps
 - 5.5.1 Effect on Stalling Speed
 - 5.5.2 Effect on Incidence
 - 5.5.3 Effect on Washout
 - 5.5.4 Effect on Drag
 - 5.5.5 Effect on Trim
 - 5.5.6 Effect on Top Speed
- 5.6 Summary

6 Angle of Attack Stability, Trim, and Spiral Dives

- 6.1 The Basic Stability Principle
 - 6.1.1 Center of Mass Too Far Aft
 - 6.1.2 Center of Mass in the Middle
 - 6.1.3 Center of Mass, Lift, and Area
 - 6.1.4 Pitch-Wise Equilibrium
 - 6.1.5 Canards Operate on the Same Principle
 - 6.1.6 Beyond Decalage
 - 6.1.7 Springs and Bobweights
 - 6.1.8 Pitch Damping
 - 6.1.9 Center of Mass Too Far Forward
 - 6.1.10 Other Failure Modes

- 6.1.11 Practical Considerations
- 6.1.12 Phugoid Oscillations
- 6.2 Spiral Dive
 - 6.2.1 Which Way Is Up?
 - 6.2.2 Overview
 - 6.2.3 General Discussion
 - 6.2.4 Recovering From a Spiral Dive
 - 6.2.5 Try It Yourself
- 6.3 Summary

7 More About Energy and Power

- 7.1 Introduction
- 7.2 Making Changes in Airspeed
 - 7.2.1 Front Side of the Power Curve
 - 7.2.2 Back Side of the Power Curve
 - 7.2.3 Right versus Wrong Procedures
- 7.3 You Can Get Away With A Lot During Cruise
- 7.4 Let "George" Do It
- 7.5 Max Performance using the Power Curve
 - 7.5.1 Best Rate of Climb
 - 7.5.2 Zero Power Available
 - 7.5.3 Best Angle of Climb
 - 7.5.4 Power Depends on Altitude via True Airspeed
 - 7.5.5 Other Power and Altitude Effects
 - 7.5.6 Best Glide: Wind, Downdrafts, etc.
 - 7.5.7 Weight Effects
- 7.6 Variations in the Power Curve
 - 7.6.1 Power Curve Depends on Aspect Ratio
 - 7.6.2 Sketching the Curve
 - 7.6.3 Some Theory
 - 7.6.4 Power Requirements versus Speed
 - 7.6.5 Power Requirements versus Altitude
- 7.7 Energy Management Stunts
 - 7.7.1 High-Speed Steep Descent
 - 7.7.2 Low-Speed Steep Descent
 - 7.7.3 Skimming in Ground Effect
- 7.8 Summary

8 Yaw-Wise Torque Budget

- 8.1 Overview
- 8.2 Yaw Stability
- 8.3 Yaw Damping
- 8.4 Helical Propwash
- 8.5 P-Factor
 - 8.5.1 Blade Speed
 - 8.5.2 Blade Angle
 - 8.5.3 Initial Takeoff Roll
 - 8.5.4 Observing P-Factor
- 8.6 Gyroscopic Precession
- 8.7 Canted Engine
- 8.8 Rudder Usage During Rolls
 - 8.8.1 Analysis of a Roll
 - 8.8.2 Designers' Tricks
 - 8.8.3 Transitory Adverse Yaw
 - 8.8.4 Steady Adverse Yaw -- Twisted Lift
 - 8.8.5 Yaw-Wise Inertia
 - 8.8.6 Amount of Rudder Required
 - 8.8.7 Summary: Coordinated Turning Procedures
- 8.9 Long-Tail Slip
- 8.10 Boat Turn
- 8.11 Weathervaning During Taxi
- 8.12 Asymmetric Thrust

8.13 Yaw-Wise Torque Budget --- Summary

9 Roll-Wise Torque Budget

- 9.1 Dihedral
- 9.2 Other Forms of Slip-Roll Coupling
- 9.3 Roll-Wise Stability
- 9.4 Differential Wingtip Speed; Overbanking
- 9.5 Rolling Moment due to Propeller Drag
- 9.6 Engine Inertia
- 9.7 Climbing and Descending Turns
- 9.8 Roll-Wise Torque Budget --- Summary

10 Equilibrium, Stability, and Damping

- 10.1 Equilibrium
- 10.2 Stability
- 10.3 Damping
- 10.4 Relationship of Stability and Damping
- 10.5 Oleo-Pneumatic Struts
- 10.6 Oscillations
 - 10.6.1 Analysis of Dutch Roll
 - 10.6.2 How to Fight Oscillations

11 Slips, Skids, and Snap Rolls

- 11.1 A Lesson on Snap Rolls
- 11.2 Intentional Slips
- 11.3 Skids
- 11.4 Anticipate Correct Rudder Usage
- 11.5 Perceiving Slip, Perceiving Coordination
 - 11.5.1 Looking Out the Side
 - 11.5.2 Looking Out the Front
 - 11.5.3 Using the Inclinator Ball
 - 11.5.4 Using the Seat of Your Pants
 - 11.5.5 Intentional Slips
 - 11.5.6 Slip Angle versus Bank Angle
- 11.6 Summary

12 Landing

- 12.1 Planning the Approach
 - 12.1.1 Other Planning Issues
 - 12.1.2 Traffic
 - 12.1.3 Obstacle Clearance
- 12.2 Judging Left or Right
- 12.3 Judging High or Low; Rule of Thumb
- 12.4 Judging Pitch Attitude and Angle of Attack
 - 12.4.1 Use Outside References and Trim
 - 12.4.2 Observe and Control More Than One Thing
 - 12.4.3 Correct for Wind
- 12.5 Other Perceptions
- 12.6 Basic "Normal" Landing
 - 12.6.1 Short Final
 - 12.6.2 Flare
 - 12.6.3 Timing the Flare
 - 12.6.4 Touchdown and Rollout
- 12.7 High-Performance Landing
 - 12.7.1 Choose the Right Runway
 - 12.7.2 Use the Right Configuration
 - 12.7.3 Touch Down at the Right Point
 - 12.7.4 Touch Down at a Low Speed
 - 12.7.5 Use the Brakes
 - 12.7.6 Summary: High-Performance Landing

- 12.8 Soft-field Landing
- 12.9 Crosswind Landing
 - 12.9.1 Basics
 - 12.9.2 Heading Control
 - 12.9.3 Drift Control
 - 12.9.4 Flare and Touchdown
 - 12.9.5 737 Scheme
- 12.10 Going Around
- 12.11 Learning to Land the Airplane
 - 12.11.1 Maneuver by Reference to the Edge
 - 12.11.2 Hesitation Takeoff
 - 12.11.3 Practice Maneuvering at Altitude
 - 12.11.4 Practice Flaring and Stalling at Altitude
 - 12.11.5 Practice Flying in the Runway Environment
 - 12.11.6 Learn Soft-Field Procedure First
 - 12.11.7 Nose-High Rollout
 - 12.11.8 Recovering from an Evil Zoom
 - 12.11.9 Salvaging an Imperfect Flare
- 12.12 Fly with a Light Touch
- 12.13 Critique Your Own Landings

13 Takeoff

- 13.1 Simplest Takeoff
- 13.2 Normal Takeoff
- 13.3 Obstructed-Field Takeoff
- 13.4 Soft-Field Takeoff
- 13.5 Crosswind Technique
- 13.6 Multi-Engine Takeoff
- 13.7 Planning and Decisionmaking
 - 13.7.1 Monitoring Takeoff Performance (wrong)
 - 13.7.2 Monitoring Takeoff Performance (right)
 - 13.7.3 Causes of Diminished Power
 - 13.7.4 Plan & Practice Rejected Takeoffs
 - 13.7.5 After Liftoff: Departure Climb
- 13.8 Other Elements of the Takeoff
- 13.9 Summary

14 Cross-Country Flying

- 14.1 Pilotage
 - 14.1.1 Airports Make Good Landmarks
 - 14.1.2 One-Dimensional Landmarks
 - 14.1.3 Choose Distinctive Landmarks
 - 14.1.4 Doglegs
 - 14.1.5 Reality-Based Navigation
- 14.2 Dead Reckoning
 - 14.2.1 Course
 - 14.2.2 Distance, Time, and Airspeed
 - 14.2.3 Crosswind Correction
 - 14.2.4 The Wind Triangle
 - 14.2.5 Discussion
- 14.3 Navigating by Instruments
 - 14.3.1 Don't Be a Gauge Junkie
 - 14.3.2 Navigation Systems (Brief Survey)
 - 14.3.3 Intended Heading
- 14.4 VOR Techniques
 - 14.4.1 Off-Course Distance
 - 14.4.2 Approaching the Station
 - 14.4.3 Progress Along the Course
 - 14.4.4 Twisted VORs
- 14.5 Combined Techniques
- 14.6 Staying Un-Lost
- 14.7 Getting Un-Lost

- 14.7.1 Basics
- 14.7.2 When in Doubt, Climb
- 14.7.3 GPS or LORAN
- 14.7.4 VOR Cross Radials or VOR/DME
- 14.7.5 Ask ATC
- 14.8 Flight Planning

15 Emergency Procedures

- 15.1 Engine Out Procedures
 - 15.1.1 Emergency Checklist
 - 15.1.2 Configuring for Glide
 - 15.1.3 Return to Airport?
 - 15.1.4 Power-Off Glide Perception and Planning
- 15.2 Preventing Emergencies
 - 15.2.1 Safety Margins
 - 15.2.2 Fuel Management
- 15.3 Dealing with Emergencies

16 Flight Maneuvers

- 16.1 Fundamentals
- 16.2 Seeing and Avoiding Other Traffic
- 16.3 Speeding Up and Slowing Down
- 16.4 Phugoids
- 16.5 Turns
- 16.6 Coordination Exercises
- 16.7 Constant-Heading Slips
- 16.8 Crabbing Along a Road
- 16.9 Slipping Along a Road
- 16.10 Familiarization Exercises; Configuration Changes
- 16.11 Transitioning to Fast and Complex Aircraft
- 16.12 Turns around a Point
- 16.13 Eights Around Pylons
- 16.14 Chandelles
- 16.15 Lazy Eights
- 16.16 Eights on Pylons
 - 16.16.1 Turns on a Pylon
 - 16.16.2 Eights on Pylons
- 16.17 Changing Headwinds and Tailwinds
 - 16.17.1 Steady Wind
 - 16.17.2 Albatross Effect: Winds that Vary with Altitude
 - 16.17.3 Windshear on Approach and Departure
 - 16.17.4 Turning Downwind; Energy Budget
 - 16.17.5 Summary: Changing Headwinds and Tailwinds
- 16.18 Remarks: Ground Reference Maneuvers
 - 16.18.1 Accounting for the Wind
 - 16.18.2 Entry Strategy
 - 16.18.3 Visual Reference
 - 16.18.4 Checklist
- 16.19 Slow Flight
 - 16.19.1 Airspeed and Altitude
 - 16.19.2 Yaw and Roll
 - 16.19.3 Procedures and Perceptions
- 16.20 Stall Practice
 - 16.20.1 Preliminaries
 - 16.20.2 Provoking a Distinct Stall
 - 16.20.3 Stall Recovery
 - 16.20.4 Power-On Stalls
 - 16.20.5 Accelerated Stalls
 - 16.20.6 Evil Zooms

17 Multi-Engine Flying

- 17.1 Engine Out Scenarios
 - 17.1.1 Takeoff
 - 17.1.2 Climb
 - 17.1.3 Slip String
 - 17.1.4 Coordination
 - 17.1.5 Perception and Initial Response
 - 17.1.6 Yaw Control at Reduced Speeds
 - 17.1.7 Minimum Control Speed --- Definitions
 - 17.1.8 Effect of Altitude, Weight, etc.
 - 17.1.9 Effect of Center of Mass
 - 17.1.10 Effect of Drag (e.g. Landing Gear)
 - 17.1.11 Roll Control
 - 17.1.12 Critical Engine
- 17.2 Engine Out Procedures
 - 17.2.1 Basic Takeoff Considerations
 - 17.2.2 Balanced Field Length; Takeoff Decision Speed
 - 17.2.3 Procedure: Low Altitude
 - 17.2.4 Procedure: Higher Altitude
 - 17.2.5 Airspeed Management
 - 17.2.6 Engine-Out Go-Arounds
 - 17.2.7 Low-Speed Engine-Out Demonstrations

18 Stalls and Spins

- 18.1 Stalls: Causes and Effects
- 18.2 Stalling Part vs. All of the Wing
- 18.3 Boundary Layers
 - 18.3.1 Separated versus Attached Flow
 - 18.3.2 Laminar versus Turbulent Flow
 - 18.3.3 Boundary Layer Control
 - 18.3.4 Summary
- 18.4 Coanda Effect, etc.
 - 18.4.1 Tissue-paper Demonstration
 - 18.4.2 Blowing the Boundary Layer
 - 18.4.3 Teaspoon Demonstration
 - 18.4.4 Fallacious Model of Lift Production
 - 18.4.5 Summary
- 18.5 Spin Entry
- 18.6 Types of Spin
 - 18.6.1 Spin Modes
 - 18.6.2 Samaras, Flat Spins, and Centrifugal Force
 - 18.6.3 NASA Spin Studies
 - 18.6.4 Effects of Changes in Orientation of Spin
- 18.7 Recovering from a Spin
- 18.8 Don't Mess With Spins

19 The Laws of Motion

- 19.1 Straight-Line Motion
- 19.2 Sitting in a Rotating Frame
- 19.3 Moving in a Rotating Frame
- 19.4 Centrifuges with and without Gravity
 - 19.4.1 The Centrifugal Field is Real
 - 19.4.2 Centrifuge
 - 19.4.3 Centrifuge and Gravity
- 19.5 Centrifugal Effects in a Turning Airplane
- 19.6 Angles and Axes
 - 19.6.1 Directions of Rotation: Yaw, Pitch, and Roll
 - 19.6.2 Attitude: Heading, Pitch, Bank
 - 19.6.3 Angle Terminology
 - 19.6.4 Yaw Does Not Commute with Pitch
 - 19.6.5 Yaw Does Not Commute with Bank
- 19.7 Torque and Moment
- 19.8 Angular Momentum

- 19.9 Gyroscopes
 - 19.9.1 Precession
 - 19.9.2 Precession: Which Way and How Much
 - 19.9.3 Inertial Platform
- 19.10 Gyroscopic Instruments
 - 19.10.1 Heading Indicator
 - 19.10.2 Artificial Horizon
 - 19.10.3 Rate-of-Turn Gyro

- 20 The Atmosphere
 - 20.1 Circulation Around Fronts and Low Pressure Centers
 - 20.1.1 Flow Around a Low
 - 20.1.2 Fronts and Troughs
 - 20.2 Pressure and Winds Aloft
 - 20.2.1 Thermal Gradient Wind
 - 20.2.2 Altimetry
 - 20.2.3 High Altimeter due to Low Temperature
 - 20.3 Prevailing Winds and Seasonal Winds
 - 20.3.1 Primary Circulation Patterns
 - 20.3.2 Continental / Oceanic Patterns
 - 20.4 Summary

- 21 Pilot-In-Command Decisionmaking
 - 21.1 Decisionmaking In General
 - 21.2 Decisionmaking during Lessons
 - 21.2.1 Please Act As PIC During Lessons
 - 21.2.2 How's Your Workload?
 - 21.2.3 I've Got It
 - 21.2.4 Hood Work
 - 21.3 Layers of Safety
 - 21.4 Example: Obstacle Clearance
 - 21.5 Flow Pattern
 - 21.6 Checklists
 - 21.7 Personal Minimums
 - 21.8 Skepticism, and Crisp Execution of Plan B
 - 21.9 Leadership and Resource Management
 - 21.10 Learn from the Experience of Others
 - 21.11 Try to Outdo Yourself

Introduction

This is a book about how to fly airplanes. As the subtitle suggests, the main topics are

- **Perceptions:** how to use your eyes, your ears, your fingertips, and the seat of your pants — to gather the information you need.
- **Procedures:** how to use your hands and feet — to make the airplane do what you want.
- **Principles:** how to organize your thinking — to make your flying easier and safer.

Several of the ideas in this book will seem new to most pilots. The ideas are actually quite old and straightforward, but they have been not been covered by traditional pilot training. Like so many basic truths, they will seem obvious in retrospect.

For example, consider the question: ``How does the altitude respond if you pull back on the yoke?" The key idea is there are two responses: pulling back causes a *short-term* response and a *long-term* response. It is easy (and exceedingly useful) to recognize the difference between the two.

Similarly, there is an important distinction between flight at cruising speed and flight at approach speed: procedures which are appropriate in one regime are inconvenient — or downright lethal — in the other regime. This book will tell you how to do things right at high speeds, low speeds, and everywhere in between.

As a third example, consider the ``pitch trim" wheel. What does it really do? Some pilots use it (as the name might suggest) to trim for a definite pitch attitude (which is a really bad idea). Other pilots use it to trim for a definite rate of climb (which is perhaps an even worse idea). Good pilots trim for a definite airspeed, or, better yet, a definite angle of attack.

The best pilots all seem to know these things implicitly. The purpose of this book is to make these things explicit — to give them names and to draw pictures of them.

Some people may still be wondering: is it really necessary to learn new procedures, perceptions, and principles? After all, there are 700,000 pilots out there, most of whom seem to get by OK. The answer is simple: 2000 of those pilots had accidents last year. Many of those accidents would not have occurred if people had been taught the ideas put forward in this book.

* Readership, Topics, and Objectives

This book is intended to appeal to pilots and everyone else who is interested in how airplanes behave. The idea is to concentrate on ideas that are useful in the cockpit, and to explain them as clearly as possible.

In addition to describing how the airplane behaves, this book describes in some detail *why* the airplane behaves that way. This may not be strictly necessary, but it is often very helpful, because: (1) Knowing why gives you more confidence that you are doing the right thing. (2) Knowing why helps you know what to expect in unusual situations. (3) Explanations that make sense are easier to remember than explanations that don't make sense. Human beings hate being told to do something without any explanation. If they are not told the true explanations, they will make up their own ``explanations". All too often these made-up ``explanations" cover only the everyday situations; they go haywire when applied to unusual situations, let alone emergencies.

Here are just a few of the topics to be covered:

- What happens if you push or pull on the yoke a little?
- What happens if you open or close the throttle a little?
- What does the trim wheel really control, and why?
- What is the best way to escape from a spiral dive?
- What happens if you go outside the weight & balance envelope?

- What do the airflow and pressure patterns look like near a wing?
- Why is a skid more dangerous than a proper slip?

This is not meant to be an aerodynamics book. If you want to *build* airplanes, go read an aerodynamics book. If you want to *fly* airplanes, read this book.

Actually, there are two kinds of aerodynamics books on the market:

1. ``Aerodynamics for engineers" — The good news is that these books are typically quite detailed and reliable. The bad news is that even the simplest ideas are expressed in mathematical terms; you will need years and years of study in order to understand what is being said. The other bad news is that even if you can follow the math, it won't do you any good during flight. I don't do calculations in the cockpit, and you shouldn't either.
2. ``Aerodynamics for pilots" — Many of these books are bad news all around. They don't really tell you how to build an airplane, and they don't really tell you how to fly it, either. They might tell you that angle of attack is important, but they don't tell you how to perceive angle of attack during flight, or how to control it. What's worse, many of the ideas in these books are just plain wrong.

For example, nearly all of the ``aerodynamics for pilots" books say a wing produces lift because it is curved on top and flat on the bottom. Alas, this isn't correct; it isn't even a useful approximation. We all know that airplanes can fly just fine upside down, which indicates that the difference in shape between top and bottom can't be all that crucial. Besides, some aircraft use symmetric airfoils (where the top is a mirror image of the bottom) and they work just fine.

Again, the purpose of this book is to concentrate on ideas that are useful in the cockpit. It explains things at a nontechnical level that should be accessible to almost everybody. Most people (including me) find the *picture* of an airflow pattern a lot easier to grasp than the equation that describes the airflow.

* How to Use this Book

I hope you will find these topics interesting... but this book is not *just* for entertainment: I find that the information presented here helps people fly the airplane better.

As illustrated in [figure 0.1](#), the main objective is to know what procedures to follow, and what behavior to expect from the airplane. This requires two things: background information and practice. Practice is the more important of the two. Background information (knowing *why* the airplane ought to behave as it does) tells you *what* to practice, and helps you understand and remember what you've practiced.



[Figure 0.1](#): Main Objective and its Supports

This book can tell you a few things about good procedures, and can buttress those procedures with whys and hows — but you will have to practice on your own.

There are many very fine pilots who do not understand how an airplane works, just as there are many very fine pianists who do not understand how a piano works. Proficient performance is what matters, in both cases. Good procedures and good habits are crucial.

During flight, you will probably not have time to do any deep reasoning about aerodynamics, especially in critical situations where doing the right thing matters most. As the fighter pilots say, in a tight situation your brains turn to water and run out your ears.¹

Therefore, read the book — enjoy the book — but if you want to get maximum value from it go fly with an instructor and practice the procedures until they become second nature.

8212; Instructions — Terms of Sale

The purpose of this book is to express some of my ideas and opinions. The suitability of this book for any other purpose is expressly disclaimed. This book comes with no warranty whatsoever.

It is foreseen that you may wish to take action based on some of these ideas and opinions. Such action is entirely at your own risk. You should be aware that aviation involves risks, some of which are irreducible, and some of which can be greatly reduced by careful piloting.

Some care has been taken with this book, in the hopes that it will dispel more errors and misconceptions than it creates. However, nothing in this world is perfect, and you are warned that this book is neither 100% complete nor 100% error-free.

Before taking any potentially hazardous action, obtain and understand all available information on the subject. Do not use this book as a substitute for skilled professional flight instruction.

In no case will the author or publisher be liable for any direct, indirect, secondary, or consequential damages.
I

1

Or something like that.

1 Energy Awareness and Energy Management

Note: You can buy a used airplane for about the same price as a new sports car.

Riddle: What's the main difference between the sports car and the airplane?

Answer: If you accelerate the sports car to about 75 miles per hour and pull back on the steering wheel, nothing very interesting happens.

When piloting an airplane, two of your most fundamental duties are (1) controlling the airplane's speed and (2) controlling its altitude.

Performing these duties would be easy if the airplane were equipped with ideal controls, so that you could (1) move a lever that would immediately change the airspeed by a few knots, with no change in altitude, or (2) move another lever that would immediately change the altitude by a few dozen feet, with no change in airspeed.

Alas, it is physically impossible to build an airplane with such ideal controls. One purpose of this chapter is to explain how *real* controls affect the airspeed and altitude of a *real* airplane.

For example, consider the seemingly simple maneuver of changing speed while maintaining a constant altitude. We will see that this requires a complex sequence of adjustments of several controls. There are two ways to deal with this maneuver. One way would be to discover (by trial and error) the required sequence of adjustments, and perform that sequence by rote forever after. A far easier and better way is to understand the fundamental relationships, so that the proper sequence seems logical and obvious.

Understanding how the airplane *really* responds to the controls makes your flying not only easier, but safer as well.

Generally, a pilot who tries to control airspeed and altitude separately winds up controlling one or the other rather poorly. Usually it is the airspeed that suffers. All too often, the airspeed gets too low, whereupon the wing stalls and the pilot rather abruptly loses control. This is the classic stall/spin accident. It is usually fatal. The good news is that if you understand what the controls *really* do, you can stay out of this sort of trouble.

The key to understanding the relationship between airspeed and altitude — and several other things — is the concept of *energy*.

Energy is not a new or complicated concept. Most pilots understand that being ``high and fast" is very, very different from being ``low and slow"; the concept of energy just makes this notion a little more precise and gives it an official name.

Good pilots think about energy all the time. The more critical the situation, the more carefully they evaluate the energy before reaching for the controls.

Once you grasp the basic concept of energy, you will be able to apply it in many ways, to many different situations. This is a big improvement over trying to figure out all possible situations one by one. Energy gives you the ``big picture".


1.1 Total Energy Cannot Change

As illustrated in [figure 1.1](#), there are four types of energy that are crucially important for airplanes, namely:

- potential energy, which is proportional to the airplane's altitude;
- kinetic energy, which is proportional to the square of the airspeed;
- the chemical energy in the fuel; and finally
- the energy left behind in the air as the plane passes through, stirring the air and leaving it slightly warmer.

There are of course other types of energy, but the four forms listed above are the ones pilots use all the time,

so let's concentrate on them for now.¹



[Figure 1.1](#): Total Energy Cannot Change

Energy has the remarkable property that it cannot be created or destroyed — it can only be converted from one form to another. This rule (which physicists call the law of conservation of energy) is not one of Newton's laws; it was not even known in Newton's day.

Consider the analogy with water: it can be converted to ice and back again, yet the amount of water doesn't change in the process. Even if the water leaks away and we lose track of where it is, the number of water molecules in the world hasn't changed.

The same notion applies to energy, as illustrated in [figure 1.1](#). Fuel energy can be converted to altitude; altitude can be exchanged for airspeed; altitude can be cashed in to pay for drag; et cetera. The amount of energy in the world doesn't change.


Some of these energy conversions are irreversible. Fuel burn, for example, is a one-way street; we cannot (alas) operate the engine backwards and replenish the fuel supply. Similarly, when energy is dissipated by drag, that energy can never be recaptured in a useful form.

The airspeed and altitude together are called the *mechanical energy*. Engine power increases the mechanical energy, while dissipation decreases the mechanical energy.



Altitude is being cashed in to pay for drag. The airspeed is not changing, and no energy is being taken from the fuel tank.

[Figure 1.2](#): Energy Conversion -- Glide




Fuel is being consumed to pay for drag and purchase altitude.

[Figure 1.3](#): Energy Conversion -- Climb



Fuel is being consumed to pay for drag. Altitude and airspeed are not changing much.

[Figure 1.4](#): Energy Conversion – Cruise



convert very much fuel.

[Figure 1.5](#): Energy Conversion – Zoom



Conversely, if you push forward on the yoke, the airplane will speed up and descend. Once again, if you do it quickly enough, drag and engine power will not affect the energy budget very much.

[Figure 1.6](#): Energy Conversion – Pushover



During the early part of the takeoff roll, drag is negligible. There is no change in altitude, so virtually all engine power goes toward building up airspeed.

[Figure 1.7](#): Energy Conversion – Initial Roll



An important conversion is the flare maneuver, which occurs at the end of every flight. It is possible to maintain altitude without using the engine, by gradually cashing in airspeed to pay for drag.


[Figure 1.8](#): Energy Conversion – Flare

1.2 Energy Conversion

[Figure 1.2](#) through [figure 1.8](#) show several examples of how one form of energy can be converted to another. We now investigate energy-conversion processes in a little more detail.

1.2.1 Converting Speed to Altitude and Back

An airplane (like any other object) has *potential energy* proportional to its altitude. Every increment of altitude represents an increment of energy. Similarly, any moving object has *kinetic energy* proportional to the square of its speed. We can easily convert back and forth between these two forms of energy. A roller-coaster is a well-known² example of this, as illustrated in [figure 1.9](#).



[Figure 1.9](#): The Law of the Roller Coaster

At the left of the figure, we have a roller-coaster at a low altitude, moving quickly. In the middle of the figure,

the roller-coaster has a higher altitude, but much less speed. At the right of the figure, the roller-coaster has returned to the lower altitude and regained its speed.

Since the roller-coaster carries no fuel and has very little friction, potential energy (altitude) and kinetic energy (speed) are the only forms of energy we need to take into account.

Here is the law of the roller-coaster:

Conversion factor = 9 feet per knot, per hundred knots

This law applies to airplanes, roller-coasters, or anything else that converts potential energy to or from kinetic energy. The altitude gain is proportional to (a) the amount of airspeed loss times (b) the average airspeed during³ the maneuver. Let's apply this to a couple of examples: if you are cruising straight and level at 201 knots, and you pull back on the yoke, when you reach 200 knots you will have zoomed up 18 feet. If you started at 101 knots and pulled back to 100 knots (once again a loss of one knot) you would only gain 9 feet.

This rule applies to aircraft, roller coasters, and indeed all objects in the world — in situations where friction can be neglected. The conversion factor, 9 feet per knot per hundred knots, is just the reciprocal of the acceleration of gravity⁴ expressed in aviation units.

The two forms of energy — altitude and airspeed squared — are deeply related, even though they are measured in different units. We need a conversion factor (9 feet per knot per hundred knots) so we can convert from one set of units to the other.

1.2.2 Energy Per Unit Mass

Since we are about to start comparing these mechanical forms of energy with other forms, we must start paying attention to an additional detail: an object's potential energy depends not only on its altitude but also on its mass. A 300-ton Boeing at any given altitude has 300 times more potential energy than a 1-ton Piper at the same altitude.

Similarly, an object's kinetic energy is also proportional to its mass. A 300-ton object at any given airspeed has 300 times more energy than a 1-ton object at the same airspeed.

Since the mass of an airplane does not usually change much during the course of a maneuver, we can often simplify the discussion by ignoring the distinction between "energy per unit mass" and genuine "energy". In cases where the distinction matters, I will remind you of it.

1.2.3 Converting Fuel to Altitude

Having understood the conversion between altitude and speed, let's bring fuel into the picture. Each pound of fuel contains a certain amount of chemical energy. The engine allows us to convert this chemical energy to mechanical energy. Assuming typical engine efficiency, the fuel-to-altitude conversion factor is:

Typical conversion factor = 6300 foot-tons per gallon

That is, climbing 6300 feet takes 1 gallon more fuel than level cruising for the same amount of time, in a typical one-ton airplane. A heavier plane would require proportionately more fuel for the same climb.




Figure 1.10: Converting Fuel to Altitude

To understand where this number comes from, and what it means, consider the experiment shown in [figure 1.10](#). First we fly straight and level for ten minutes, maintaining 90 knots; we observe the fuel flow gauge is reading 5 gallons per hour. Then we open the throttle and climb for the same amount of time at the same airspeed; we observe a vertical speed of 630 feet per minute and a fuel flow of 11 gallons per hour.

The experiment tells us that in the example airplane, climbing at 90 knots consumes 6 gallons per hour more fuel than level flight at the same speed. During 10 minutes (one sixth of an hour) the climb will eat up one extra gallon. The same 10 minutes at 630 fpm will gain us 6300 feet of altitude. The example plane weighs exactly one ton, so we get the conversion factor claimed above: 6300 foot-tons per gallon.

The exact value of the conversion factor will vary a little bit from airplane to airplane, depending on the efficiency of the engine, etc., but 6300 foot-tons per gallon is a good approximation in most cases.

To determine the fuel-to-altitude conversion factor for your airplane, you can (1) divide 6300 by the weight (in tons) of your airplane; (2) perform the experiment described above; or (3) work it out using the cruise-performance and climb-performance data in your airplane's Pilot's Operating Handbook (POH).

Here are the results for several airplanes from various manufacturers, using POH numbers:

Airplane	foot-tons per gallon
Two-place, carbureted, fixed gear, fixed prop	6172
Four-place, carbureted, fixed gear, fixed prop	6362
Four-place, fuel injected, retractable, constant speed prop	6410
Six-place twin, fuel injected, retractable, constant speed prop	6384

If the airplane were 100% efficient at converting fuel to altitude, the conversion factor would be higher — but it is hard to build a really efficient engine with a reasonable size, weight, and cost.

1.2.4 Power versus Energy

Since fuel corresponds to altitude, fuel flow rate must correspond to rate of climb. Airline crews use this fact routinely: to make the transition from level flight to a 500 fpm descent at constant airspeed, they just retard the throttles until they see a certain reduction on the fuel flow gauges.

This notion of "energy per unit time" is officially called *power*. You don't want to confuse power with energy, any more than you would want to confuse a vertical speed indicator with an altimeter; the former indicates altitude *per unit time*, while the latter indicates altitude.

The airplane has instruments that measure most — but not all — of the relevant forms of energy and power. The energy gauges include the altimeter, airspeed indicator, and fuel gauges. These tell you how much potential energy, kinetic energy, and chemical energy there is on board.

The most common power gauges include vertical speed indicators and fuel flow gauges; these tell you at a glance how much power is flowing in and out of the potential and chemical reservoirs. Sometimes other power gauges are installed; gliders often have a "total energy variometer", which measures the rate of change in mechanical energy (potential plus kinetic) by measuring a combination of altitude change and airspeed change. Such a device is more useful than an ordinary vertical speed indicator for detecting updrafts, for the following reason: Inadvertently pulling back on the yoke will cause a positive indication on the vertical speed indicator (by the law of the roller-coaster) which might be confused with a real updraft; pulling on the yoke will cause *no* indication on the TE variometer.

Since the glider has no engine power to worry about, the TE variometer gives a reasonably complete picture of how much power is flowing in or out of the aircraft (updraft = power in; dissipation = power out). In an airplane with an engine and without a TE variometer, it is somewhat trickier to visualize what is going on.

[Figure 1.11](#) summarizes this section by showing the various forms of energy and power, and some of the relationships between them. The ones that are circled are the ones for which aircraft gauges commonly exist.




Figure 1.11: Forms of Energy and Power

A reminder for the purists: a given quantity of gasoline contains a certain amount of chemical energy, period. In contrast, a given amount of altitude represents a certain amount of energy *per unit mass* of airplane. Therefore it is a slight oversimplification to suggest (as in [figure 1.11](#)) that the fuel gauge and the altimeter measure exactly the same thing, but there is no harm in it if the mass of the airplane isn't changing. Similar remarks apply to the airspeed indicator.

1.2.5 Drag and the Power Curve — Introduction

The time has come to bring drag into the picture.

The power dissipation due to drag is equal to the drag force times the airspeed.⁵ Power is energy per unit time, which should not be confused with energy itself.

The distinction between energy and power is emphasized in the following analogy:

Altitude (energy) is like money in the bank. You pay the cost of climbing to altitude only once. If desired, you can cash in the altitude energy to do useful things.

Drag (power) is like rent; you have to pay a certain amount of energy per unit time for the privilege of flying the airplane through the air. That energy can never be recovered.

The amount of drag — the amount of rent you have to pay — depends on your airspeed⁶ in a complicated way. The relationship is shown in [figure 1.12](#), and is called the *power curve*.




Figure 1.12: Power Curve (Engine Idle)

(You may be more familiar with this curve in an upside-down version called the "power required" curve. The orientation given here is preferable, for the following reason: Airplanes don't have "power required" meters but do have vertical speed indicators. Therefore this orientation is more meaningful in the cockpit. Also note that drag contributes a negative amount to our power budget, in contrast to the engine which contributes a positive amount.)

In the figure, airspeed is labelled in Knots of Indicated Air Speed (KIAS). A knot is a nautical mile per hour, as discussed in [section 14.2.2](#). The meaning of indicated (versus true) airspeed is discussed in [section 2.12](#).

This figure applies to straight-ahead gliding flight. The engine is producing zero power; for any particular airspeed, the airplane will descend at the rate specified by the power curve. Altitude — i.e. gravitational

potential energy — is being cashed in to pay for the frictional losses.

The traditional units for the vertical axis in this figure would be horsepower, but I have used feet per second instead. This is intended to clarify the equivalence of all four forms of energy by measuring them in a common set of units. We have seen how to think of airspeed in terms of altitude (9 feet per knot per hundred knots) and also how to think of fuel in terms of altitude (6300 foot-tons per gallon), so it is only logical that power should be measured as vertical speed; that is, altitude-change per unit time.

The terminology and basic applications of the power curve are presented in the next couple of paragraphs; some more advanced applications will be presented in [section 7.5](#).




Figure 1.13: Power Curve — Three Regimes

As shown in [figure 1.13](#), the power curve is divided into three regimes. The right-hand part of the curve (from moderate airspeeds on up) is called the *front side of the power curve*. Normal cruising flight is conducted in this range of airspeeds.

In this regime, the faster you go, the more power is consumed by friction. This is completely unsurprising — everybody knows that moving an object through the air quickly takes more force than doing it slowly. You can see in [figure 1.13](#) that if you glide at a very high airspeed, you will have a large rate of descent.

What is less obvious to non-pilots is that at low airspeeds there is another regime with very high drag. This is called the *mushing regime*, and is labelled in the figure. The logic here is that it is more efficient to visit a lot of air and yank it down gently than to visit a small amount of air and yank it down violently. In this regime the airplane must fly at a high angle of attack in order to support its weight. This creates strong wingtip vortices that in turn produce huge amounts of induced drag, as discussed in [section 3.12.3](#). Therefore if you are in the mushing regime, flying more slowly causes more descent rate, as can be seen in [figure 1.13](#). This is quite unlike cars — a car moving slowly incurs very little frictional loss. Of course, cars don't need to support their weight by pulling down on the air.

The dividing line between the mushing regime and the front side of the power curve is the highest point on the power curve. At this point, the airplane can fly with the minimal amount of dissipation; this is the "low-rent district". The airspeed where this occurs is called the best-rate-of-climb airspeed and denoted V_Y .

Finally, we consider the extreme lower-left part of the power curve. This is called the *stalled regime*, as indicated in [figure 1.13](#).⁸ Flight in this regime is very, very peculiar.

The mushing regime and the stalled regime are collectively referred to as the *back side of the power curve*.

Life would be simpler if manufacturers would explicitly show the power curve somewhere in the POH, but they don't. You have to figure it out for yourself. Fortunately, the general shape of the power curve is more-or-less⁹ the same for all airplanes, so the concepts discussed here are very widely applicable.

1.2.6 Rates of Energy Conversion

An airplane can very rapidly and efficiently convert airspeed to altitude, and vice versa. Because of this, these two forms of energy are often considered together, and are collectively referred to as the mechanical energy.

In contrast, it is difficult to convert fuel to mechanical energy quickly, and it is difficult to dissipate large amounts of mechanical energy via drag quickly (especially while maintaining a safe airspeed).

A rapid conversion of airspeed to altitude is called a *zoom* — a fairly common maneuver.¹⁰ You should always be careful when performing a zoom, because if the airspeed gets too low there could suddenly be very unpleasant consequences.

The airplane's ability to convert airspeed to altitude and back again is the key to many aerobatic maneuvers. There is no way you could perform a loop using engine power alone; you have to zoom. Bob Hoover's airshow routine typically closes with a spectacular energy management demonstration. After shutting down the engine, he performs a series of complex aerobatic maneuvers, including an eight-point roll and a hammerhead.¹¹ He then returns for landing and coasts to the reviewing stand, all without restarting the engine. It is quite a fascinating lesson in pilot technique.

1.3 Effect of Controls on Energy

The previous section introduced the main forms of energy that affect flight. The next step is to discuss how the pilot can control the energy in various ways. This section doesn't introduce very many additional concepts; it mainly just combines and applies the concepts introduced previously.

We continue to use the analogy between energy and money. Therefore, deciding how much power should flow from one reservoir to another is called the power budget.

1.3.1 Power Budget — Using the Engine

Figure 1.14 shows how engine power affects the power budget.¹² The bottom curve applies when the engine is operating at 1700 RPM, the middle curve applies at 2000 RPM, and the top curve applies at 2300 RPM.




Figure 1.14: Power Curve (Various Engine Power Settings)

Point A indicates a 500 fpm descent at 80 knots. Point B indicates level flight at the same airspeed, and point C indicates a 500 fpm climb still at the same airspeed. The rule is simple: if the engine produces more power, the airplane will descend at a lesser rate or even ascend.

Point D corresponds to level flight at 110 knots. The power setting is the same as at point C — but the energy that was being used to purchase altitude (point C) is now being used to pay for the higher drag at the new airspeed (point D). If engine power exactly equals frictional losses, the airplane will stay level — fuel energy is being used to pay for the friction.

The numbers in this example are consistent with a rule of thumb that applies to a wide range of light aircraft: starting from level flight, to set up a 500 fpm descent,

- Reduce power by 300 RPM (for a fixed-pitch prop), or
- Reduce power by 3" of manifold pressure (for a constant-speed prop).

This rule works surprisingly well over quite a range of different makes and models. Make a point of learning whatever version of this rule applies to your airplane. It is a big improvement over blindly guessing at throttle settings.

1.3.2 The Effects of the Throttle

I make sure all my students really understand the effects of a power change. In the first or second lesson, we get the airplane trimmed for straight and level flight (using a moderate power setting). We then push the throttle a little more open. The student may be expecting that the airplane will respond by accelerating, just like a car. But airplanes are not the same as cars! In most airplanes (including all the common trainers) the airplane will actually slow down slightly.¹³ This experiment — observing how power changes affect the trim speed of the airplane — is one of the first things I do not just for students but also for myself when I am learning to fly a new make & model of airplane. (It is also important to learn how flap extension affects the trim speed, and how the flaps and power interact.)

The throttle¹⁴ controls power. What could be simpler? The throttle controls power. (Remember, power is energy per unit time.)

There are three things this power could be used for:

1. Power is needed to overcome drag. Flight at speeds above or below V_Y requires more power than flight at V_Y .
2. Climbing requires more power than level flight, other things being equal.
3. Acceleration requires more power than unaccelerated flight along the same path.

It may seem obvious that engine thrust ``should" cause the airplane to accelerate, but usually that's not what happens. Although the airplane is being pulled forward, the trim mechanism notices what is going on and immediately converts the new energy to altitude. Therefore the throttle can be reliably used to control up/down motion. As discussed in [chapter 6](#), this is the normal, natural aerodynamic behavior.

Of course, if you defeat the trim mechanism, all bets are off. For instance:

- During the takeoff roll, the airplane is not free to move in the vertical dimension, so the trim has no effect. Therefore (in this special situation) energy coming from the engine is converted to speed, not altitude.
- Similarly, suppose your autopilot is manipulating the yoke so that the airplane maintains level flight. This means the natural aerodynamics of the trim mechanism is irrelevant. When you open the throttle (in this special situation) the added energy will be converted to airspeed, not altitude. Note that the autopilot has to move the yoke to make this happen — so we can reasonably say that the airspeed change is ``caused" by the yoke movement more directly than it is ``caused" by the added power.

I reiterate that in flight, if you (and the autopilot) leave the yoke and trim alone, opening the throttle just makes the airplane climb. If you want to accelerate or decelerate without an altitude excursion, you will need to adjust the throttle *and* the yoke, as discussed in [section 7.2](#).

A car, of course, will speed up when you open the throttle. But this has got nothing to do with the behavior of an airplane in flight.

An airplane is not the same as a car. Cars don't have trim. Cars aren't free to move in the third dimension.

Now that we understand the effects of opening the throttle, the effects of closing the throttle should be no surprise. The airplane will maintain its trim speed (or possibly speed up very slightly) and descend. This is easy to understand in terms of energy; compare points B and A in [figure 1.14](#). If engine power is reduced, the only way to pay the rent is to cash in altitude energy at a steady rate.

[1.3.3](#) The Effects of the Yoke

Now let's do a slightly different experiment: pulling on the yoke. As before, start with the airplane nicely trimmed in straight and level flight. Then pull the yoke back a little ways and hold it there. What happens next? Several things will happen, on various time-scales:

- The pitch attitude will change. This is important, but will not be discussed here.[15](#)
- You will slow down. Think of this as the primary effect.[16](#) The new, lower airspeed will persist throughout the short term and the long term.
- There will be a short term effect and a long term effect on altitude. That is:
 - Because of the decrease in airspeed, you will zoom upwards. This is a short-term, one-time increase in altitude, according to the law of the roller-coaster. You are trading in kinetic energy, exchanging it for potential energy.
 - At the new airspeed, you will be operating at a new point on the power curve.
 - If this is a more-efficient operating point, you will get a long-term climb.
 - If this is a less-efficient operating point, you will get a long-term descent.

Let's clarify the long-term behavior by considering two versions of this experiment. In the first version, as illustrated in [figure 1.15](#), the airplane is initially on the front side of the power curve — cruising at 105 knots, which is definitely on the front side of the power curve. Pull back on the yoke a little, and hold it.




Figure 1.15: Pulling on the Yoke — From Cruise

What happens to the airspeed and altitude?¹⁷ The first thing that happens is that the airplane slows down from 105 knots to 100 knots. You should think of this as the primary effect of moving the yoke. This is a short-term *and* long-term effect.

As a consequence of this speed change, the airplane will zoom up about 45 feet, according to the law of the roller coaster: 9 feet per knot, per hundred knots. This is a short-term, one-time effect.

Also, we can see from the power curve that 100 knots is a more-efficient airspeed. Less power will be consumed by drag, so the airplane will ascend. (Remember we've kept the engine power unchanged.) The airplane will continue to climb at a steady rate for a long time.

The short-term altitude change is governed by the law of the roller-coaster, while the long-term altitude change is governed by the power curve.

So far this all seems pretty normal — but the second version of the experiment is much more interesting, as shown in [figure 1.16](#). Let's reconfigure the airplane for flight on the back side of the power curve — say 58 knots. Trim the plane for straight-and-level flight, then pull back on the yoke a little and hold it there.




Figure 1.16: Pulling on the Yoke — From Slow Flight

The first part of the story is the same: you will slow down. Let's say the new speed is 53 knots. As always, you should think of this as the primary effect: if you pull back on the yoke you will slow down.

The second part of the story is also the same: there will be a one-time increase in altitude. This time it will be about 25 feet. The zoom is less than in the previous case, because the initial airspeed was less.

The final part of the story contains the surprise: because the new airspeed represents a higher-drag (less-efficient) point on the power curve, the airplane will enter a steady descent. At the new airspeed, it will descend and descend and descend.

As always, the short-term altitude change is governed by the law of the roller-coaster, while the long-term altitude change is governed by the power curve.

. This scenario (a short-term ascent followed by a long-term descent) is called a *zoom*.¹⁸ It is the bane of student pilots when they start learning to perform landings. Starting from a low airspeed a few feet above the runway, they pull back on the yoke. The airplane obediently zooms upward, then (alas) descends at a tremendous rate and makes an airplane-shaped hole in the runway.

Students who have not been taught the distinction between the short-term and long-term effects have a hard time figuring out this situation.

Note: this treacherous behavior (short term ascent followed by long-term descent) does not imply that the airplane is stalled or about to stall. As mentioned in [section 1.2.5](#), the mushing regime is not the same as the stalled regime. In the mushing regime, induced drag is the culprit; stalling is a completely different issue, which is discussed in [chapter 5](#).

Sometimes the mushing regime is called the ``*regime of reversed control*'', but this is not a very good term. The following table summarizes the actual effects of pulling back on the yoke:

	Front-side effect	Mushing effect	Reversal?
Airspeed	decrease	decrease	no
Short-term altitude	increase	increase	no ¹⁹
Long-term altitude	increase	decrease	yes

By two votes out of three, we conclude that the term ``*regime of reversed control*'' is not a good description of the mushing regime.²⁰

1.3.4 Sizes of Energy Reservoirs

The following observation may help put into perspective the sizes of the various energy reservoirs. First, consider normal cruising flight: the energy in the fuel tank is enough to ``pay the rent'' (overcome drag) for several hours. Second, consider a power-off glide: starting from a reasonable cruising altitude, altitude energy can be cashed in to pay the rent for several minutes. Finally, consider the flare maneuver: it is possible to arrest a power-off descent and maintain level flight by cashing in airspeed for a few seconds. To summarize:

You can pay for drag by cashing in fuel...for a few hours.
You can pay for drag by cashing in altitude...for a few minutes
You can pay for drag by cashing in airspeed...for a few seconds.

So, we see that the available energy reservoirs have very different sizes.

This difference in sizes has many consequences, but the most important one is this: you cannot make large altitude corrections (only small ones) by borrowing from the airspeed reservoir.

That is, suppose you are a few feet below your desired altitude. The quickest way to get back up is to pull back on the yoke. You thereby cash in some airspeed energy to buy altitude, according to the law of the roller-coaster. On the other hand, if you try to go up some more by pulling back some more, you will very soon run out of airspeed.

The bottom line is: you should feel guilty about borrowing energy from the airspeed reservoir. There just isn't very much energy there to begin with, and letting the airspeed get too low can have serious consequences.

The pros and cons of controlling altitude by borrowing airspeed are discussed in more detail in [chapter 7](#).

1.4 Energy Management Strategy

The next step is to combine what we know about energy and develop general rules for energy management. Let's consider the four situations depicted in [figure 1.17](#).




Figure 1.17: Energy Management — Four Situations

In the figure, as we go from left to right the kinetic energy increases; similarly as we go from bottom to top the potential energy of the situation increases.

Let's start by considering the situation in the upper-left corner: the altitude is a bit high and the airspeed is a bit low. If we're lucky, the total energy might be about right. Therefore, the obvious thing to do is to push on the yoke. That will get rid of some altitude by converting it to airspeed, which is basically what we want.

In the lower-right corner we have the complementary scenario: the altitude is a bit low and the airspeed is a bit high. Once again, if we are lucky the total energy might be about right. Therefore, the obvious thing to do is to pull on the yoke (in moderation). That will convert some of the excess airspeed into altitude, which is basically what we want.

The situation in the upper-right corner is more challenging: both the airspeed and the altitude are too high. Unlike in the previous two scenarios, we clearly have an energy problem: the total energy is too high. There is nothing you can do with the yoke that will make the altitude better without making the airspeed worse, [21](#) and vice versa, so we have to find something else to do. The first step is to retard the throttle, the sooner the better; every bit of power that the engine produces only adds to the energy problem. The other way to get rid of energy is to increase drag. This can be done by extending the landing gear, extending the flaps, slipping, et cetera. Over time, the increased drag will take energy out of the system, which is what you want. If drag is not taking energy out of the system fast enough, you may have to perform a 360 degree turn or something in order to buy some more time.

Finally, let's consider the lower-left corner of [figure 1.17](#). In this case, both the airspeed and the altitude are too low. This is the proverbial coffin corner. You have an energy problem, and having too little energy is even worse than having too much energy. You should open the throttle immediately; this will (over time) convert some fuel energy into new airspeed and/or altitude.

If no power is available, do not try to ``stretch the glide''. There is nothing you can do with the yoke that will add new energy to the system; all you can do is minimize the loss by maintaining the canonical best-glide airspeed. Since you are too slow, push on the yoke to re-establish that airspeed. Since you are too low, choose a closer place to land.

This discussion of how the yoke and the throttle are used together for energy management is continued in [chapter 7](#).

1.5 Summary: Energy Management

Question: What makes the airplane gain altitude? Answer: four things:

- Updraft.
- Zoom.
- More engine power.

- Less drag.

The most common way of reducing drag is by selecting an airspeed closer to V_Y . (Of course it also pays to get rid of any extraneous drag, perhaps by retracting the flaps, retracting the landing gear, and/or reducing the amount of slip.)

Suppose you are on final approach for landing. You notice that you are below the glideslope. What should you do? Add power?? Pull back on the yoke?? — *This is asking the wrong question.* The glideslope indication alone doesn't give you enough information to decide what to do.

You need to perceive the airspeed as well as the height. Think about your energy: potential energy plus kinetic energy. Being low and slow is very different from being low and fast.

Instructors: on final, ask your students ``Are we high or low, fast or slow?" Make sure they evaluate the energy situation continually and correctly.

Altitude and airspeed tell you your total mechanical energy. In the short run there is nothing that will change the total mechanical energy; all you can do is use the yoke to trade energy back and forth between altitude and airspeed. The conversion factor is nine feet per knot, per hundred knots.

In the long run, the throttle (engine power) and the power curve (drag power) control the rate at which energy is entering and leaving the ``airspeed plus altitude" system. To establish a long-term climb, add power and/or trim for a speed closer to V_Y . To overcome drag (in unaccelerated level flight) requires power. To climb (while maintaining constant airspeed) requires added power. To accelerate (while maintaining constant vertical speed) requires added power.

The amount of energy in the airspeed reservoir is very small compared to the energy in the altitude reservoir, which is in turn very small compared to the energy in the fuel reservoir.

If you value your life, look at the airspeed indicator before pulling on the yoke. Looking at just one indicator (altitude *or* airspeed) for making a decision about just one control (yoke *or* throttle) is poor pilot technique and could well lead to a stall/spin accident. You must look at both indicators, size up the energy situation, and then decide what to do with both controls.

1

For instance, solar energy can produce updrafts and windshears. Sometimes the airplane's ability to extract energy from these is important, as discussed in [section 7.5.6](#) and [section 16.14.2](#).

2

Langewiesche ([reference 1](#)) devotes an entire chapter to ``The Law of the Roller Coaster".

3

To be exact: Take the initial airspeed and final airspeed and average them.

4

... that is, $g = 9.807$ meters per second per second; $1/g = 8.8537$ feet per knot, per hundred knots.

5

The relationship between force and power is discussed in more detail in [section 4.4](#).

6

As we shall see, it would be more precise to say that the drag depends on angle of attack — but airspeed is often a convenient stand-in for angle of attack, as discussed in [section 2.12](#).

7

a more precise definition of V_Y will be given in [section 7.5](#).

8

[Section 5.3](#) gives a precise definition of stall, and [section 5.3.2](#) explains why the power curve hooks back to the right in the stalled regime.

9

[Section 7.6](#) explains the slight variations from plane to plane, and how to sketch the power curve for your particular airplane.

10

The reverse conversion, altitude to airspeed, is equally common but does not have a correspondingly colorful name.

11

A hammerhead involves flying vertically upward until the airspeed is practically zero, turning 180 degrees around the now-horizontal yaw axis, and then retracing your steps vertically downward.

[12](#)

This is slightly idealized. See [section 7.5](#) for more details.

[13](#)

The rare exceptions are discussed in [section 6.1.4](#).

[14](#)

... in conjunction with the RPM control if you have a propeller governor.

[15](#)

Having a particular pitch attitude is rarely an end in itself. Instead, you should use it as a good means of controlling other things, such as angle of attack; see [section 2.6](#) and [section 2.10](#). Also note that abrupt movement of the yoke will provoke phugoid oscillations, as discussed in [section 6.1.12](#).

[16](#)

The aerodynamics of how the yoke and trim govern airspeed is discussed in [chapter 6](#).

[17](#)

Again, note that discussion of pitch changes is being postponed until [section 2.6](#).

[18](#)

Some older books call it ``ballooning".

[19](#)

... unless you pull back very, very slowly, in which case the short-term ascent might be masked by the long-term descent.

[20](#)

Similarly, in the mushing regime, other controls (such as the ailerons) become less effective, but they do not reverse.

[21](#)

... in the short run, at least — but see [section 7.7.1](#).

2 Angle of Attack Awareness and Angle of Attack Management

- If you want to go up, pull back on the yoke.
- If you want to go down, pull back a little more.
- If you want to go down real fast and spin around and around and around, just keep pulling back.

— Aviation proverb.

2.1 The Importance of Angle of Attack

Angle of attack is a very important and useful concept. Most of the airplane's critical performance numbers are more closely related to angle of attack than to anything else. Let's explore what this means.

You've probably heard that it is good to fly the airplane "by the numbers". The question is, *what* numbers?

Suppose we wish to achieve the best rate of climb:

- A) You could try to control the airplane by reference to the "rate of climb" number shown on the vertical speed indicator. This is not recommended!
- B) It would be better to maintain V_Y , the nominal best-rate-of-climb speed, as shown on the airspeed indicator, and accept whatever rate of climb results. This is almost exactly the right idea.
- C) It would be even better to realize that the best rate of climb is achieved at a particular angle of attack. In particular, if the airplane is lightly loaded compared to what was anticipated in the handbook, the best rate of climb will be achieved at a lower speed than is reflected in the handbook's V_Y value.

This is not an isolated example. Many of the airplane's critical performance numbers are really angle of attack numbers:

- p The stall occurs at a particular angle of attack.
- p The smallest power-off descent rate occurs at a particular angle of attack.
- p The best power-off glide ratio occurs at a particular angle of attack.
- p The recommended "approach speed" is really an angle of attack recommendation.
- p The best rate of climb occurs at a particular angle of attack.
- p The best angle of climb occurs at a particular angle of attack.[1](#), [2](#)

Here is a summary of the main ideas that will be explained in this chapter:

- The airplane is trimmed for a definite angle of attack. The "pitch" trim wheel is really an angle of attack selector.
- Push/pull motion of the yoke can be viewed as an extension of the trim wheel — just another way of controlling angle of attack. It is very difficult to stall the airplane unless you pull back on the yoke and/or apply lots of nose-up trim. This idea could save your neck.
- Outside visual references also provide information about angle of attack, if you know what to look for.
- The airspeed indicator provides quantitative information about angle of attack, when the airspeed is not too low. Correction factors must be applied to correct for nonstandard weight and/or G -loads.
- Configuration and power changes have minor effects on the trimmed angle of attack.

2.2 Definition of Angle of Attack

I will now explain what angle of attack is, why it is important, and how it is related to things a pilot can actually observe and control.

The basic idea is simple: the angle of attack is the angle at which the air hits the wing. The Wright brothers had only one flight instrument on their first airplane — an angle of attack instrument. It was all they needed.

Their angle of attack indicator consisted of a stick attached to the wing, with a piece of yarn dangling from the front end, as indicated in [figure 2.1](#). The yarn aligns itself with the relative wind.³ The stick serves as a reference line, and also serves to locate the yarn in a region of air that has not been too badly disturbed by the wing.



[Figure 2.1](#): Simple Angle of Attack Instrument

The angle between the stick and the yarn indicates angle of attack.

The exact alignment of the indicator stick relative to the airplane is not critical. The most elegant scheme is to orient the stick in the *zero-lift direction* so that zero angle of attack corresponds to zero coefficient of lift. That choice will be used throughout this book; see [section 2.14](#) for a discussion of other possibilities.

Most aircraft do not have any instruments that give you a direct indication of angle of attack. Surprisingly, many airliners and other aircraft that *do* have fancy angle-of-attack sensors don't make the information available to the flight crew — only to the autopilot. The bottom line is that most pilots have to use a few tricks in order to perceive angle of attack. We now discuss how this is done.

It turns out to be easier to maintain *some* constant angle of attack than to know precisely what angle of attack you've got. The strategy is summarized in the following outline.

- 1
 - There are several ways to maintain a constant angle of attack.
 - 1.1
 - The airplane is trimmed for a definite angle of attack (see [section 2.3](#)).
 - 1.2
 - You can perceive the angle of attack and regulate it by hand. To perceive the angle of attack, you need to compare the pitch attitude to the relative wind.
 - 1.2(a)
 - There are at least four ways to perceive the pitch attitude (see [section 2.5](#)).
 - 1.2(b)
 - There are a couple of ways to estimate the direction of the relative wind (see [section 2.11](#)).
- 2
 - You can use the airspeed and other considerations to decide if you are maintaining the *right* angle of attack (see [section 2.12](#)).

Now let's investigate each of the items in this outline.

2.3 Trim for Angle of Attack!

The simplest and best way to get the airplane to fly at a constant angle of attack is to *leave it alone*! An airplane, by its very structure, is trimmed for a definite angle of attack. The reason for this is discussed in [chapter 6](#). Even a dime-store balsa-wood glider wants to fly at a definite angle of attack.

This concept is so important that it is the focal point of the first lesson I give student pilots, who sometimes arrive with the misconception that pilots must use great skill and continual intervention to keep the airplane under control. I trim the airplane for straight and level flight and then take my hands off the controls, demonstrating that the airplane will fly just fine for quite a while with no intervention at all. I emphasize a professional pilot does not grab the controls firmly and move them quickly; a real pro grabs them lightly and moves them smoothly.

The second lesson is this: I trim the airplane for a speed near V_Y , straight and level. I then roll the trim wheel back a little, which results in a decrease in the trim speed. It does not result in a steady climb. I explain that

the trim wheel controls angle of attack, and that airspeed is related to angle of attack. Trim for angle of attack!

To make changes in the angle of attack, you should adjust the pitch attitude using pressure the yoke, then trim to remove the pressure, as discussed in [section 2.6](#).

Configuration changes can affect the airplane's preferred angle of attack. In a Cessna 152, 172, or 182, if you extend the flaps *while the engine is at a high power setting* or if you increase the power *while the flaps are extended* it will cause a nasty decrease in the trim speed. This is highly undesirable and dangerous behavior. This means that when you perform a go-around, the airplane tends to pitch up drastically and lose airspeed; to maintain control you need to push on the yoke while you retract the flaps and retrim. This pitch-up behavior is particularly treacherous because it is not familiar. The trim speed changes very little if you extend the flaps at low power settings, and/or change the power with the flaps retracted, so if you haven't recently performed many go-arounds or similar maneuvers you might be in for a nasty surprise.

For a typical Cherokee, extending two notches of flaps lowers the trim speed ten or fifteen knots. This is discussed further in [section 5.5](#) and [section 12.10](#). Increasing or decreasing engine power affects the trimmed angle of attack only slightly. As discussed in [section 1.3.2](#), if you just reduce power the airplane should just descend. It should not slow down appreciably; in fact it will probably speed up a little.

An advanced lesson serves to demonstrate that constant angle of attack is not quite the same as constant airspeed. When the airplane is subjected to a high *G*-loading, as in a steep turn, the trim mechanism causes it to speed up, so that it can support the increased load at the same angle of attack. This is important, since (as discussed in [section 6.2](#)) it helps explain graveyard spirals, and why it is a bit tricky to recover from them safely.

Conclusions:

Do not trim for pitch attitude.
Do not trim for rate of climb.
Trim for airspeed, approximately.
Trim for angle of attack!

[2.4](#) Three Contributions to Angle of Attack

As mentioned earlier, it is difficult to directly perceive angle of attack. Fortunately, there are three other quantities that can be perceived, and together they determine the angle of attack. They are:

- Pitch attitude, which is defined⁴ to be the angle that the fuselage makes relative to the horizontal.
- Angle of climb, which is just the angle between the flight path and the horizontal.
- Angle of incidence, which is the angle at which the wings are attached to the fuselage.

These quantities are related to the angle of attack by a very simple formula:

$$\text{Pitch Attitude} + \text{Incidence} = \text{Angle of Climb} + \text{Angle of Attack}$$

This relationship is illustrated in [figure 2.2](#). Perhaps the simplest case is straight and level flight at cruise airspeed. In this case, the pitch attitude is zero, the angle of climb is zero, and the angle of attack is equal to the angle of incidence. Some more examples, with specific numbers for a typical airplane, are included in [table 2.1](#).




Figure 2.2: Pitch + Incidence = Climb + Attack

Extending the flaps has the effect of increasing the incidence⁵ by several degrees. You need to be always

aware of what flap setting you are using, and to recognize the distinction between "pitch attitude" and "pitch attitude plus incidence". For any given flap setting, you can take the incidence to be constant, whereupon angle of attack depends only on pitch attitude and direction of flight.

The table mentions V_X and V_Y , which denote the airspeeds for best angle of climb and best rate of climb, respectively, as discussed in [section 7.5](#). The relationship of airspeed to angle of attack will be discussed in [section 2.12](#).

	Airspeed (KCAS)	Pitch Attitude	Incidence	Angle of Climb	Angle of Attack
stall	59	14.0	4.5	0	18.5
level at V_X	64	8.5	4.5	0	13.0
level at V_Y	76	4.0	4.5	0	8.5
climbing at V_Y	76	7.0	4.5	3	8.5
cruise	115	0.0	4.5	0	4.5


[Table 2.1](#): Angles in various situations

2.5 Perceiving Pitch Angle

In straight and level flight you can control angle of attack by controlling pitch attitude. You won't be able to pick a particular angle of attack such as 6.37 degrees, but whatever angle of attack you've got can be maintained.⁶

There are at least four ways of perceiving pitch attitude. Perhaps the best way is to use a mark on the windshield, as shown in [figure 2.3](#). The line of sight from your eye through the mark makes a good pointer. (Try not to move your head up and down too much!) If you can't find a scratch or bug-corpse in exactly the right place, you can *make* a mark, or a pair of marks, as discussed in [section 11.5.2](#). It is even simpler to rest your hand atop the instrument panel, holding the tip of your finger in the right place, as shown in [figure 11.2](#).

Suppose you identify (or make) the mark when the airplane is flying at the angle of attack that corresponds to V_Y . Then if you re-trim for a higher angle of attack⁷ the sight line through that mark will point two or three degrees above the horizon. Similarly, if you re-trim for high-speed cruise, the sight mark will appear three or four degrees below the horizon.




[Figure 2.3](#): Perceiving Pitch Using The Forward Horizon

The second way of perceiving pitch attitude also involves looking out the front, but uses a sight line through a point on the cowling. This is also indicated in [figure 2.3](#). Be sure you chose a point on the cowling directly ahead of your dominant⁸ eye; if your seat is way over on one side of the airplane and you choose a sight mark on the middle of the cowling, your sight line will be angled sideways, which will mess up your pitch attitude perception as soon as you try to bank the airplane. A Cessna 152 or 172 has a bolt on the cowling, directly ahead of the pilot, that makes a good sight mark.

A sight mark on the cowling has the advantage that it is farther away from your eye, so it is easier to keep both it and the horizon in focus at the same time. The disadvantage is that the sight line constructed this way sometimes points quite a ways below the horizon. This means the angle you are trying to perceive — the angle between this reference line and the relative wind — is larger. It is always harder to perceive a small change in something large than a small change in something that was small to begin with.

Using the cowling has one big advantage over using marks on the windshield: the cowling is a permanent part of the airplane and is in the same place on all airplanes of that make and model.

The third way to perceive pitch attitude is to observe the angle between the wing and the lateral horizon, as shown in [figure 2.4](#). On a high-wing airplane, the bottom surface of the wing makes a good reference. In particular, on a Cessna 152 / 172 / 182, the bottom surface has a rather large flat section, which makes an ideal reference — and this reference is very nearly aligned with the horizon at cruise angle of attack (in level flight).



[Figure 2.4](#): Perceiving Pitch Using The Lateral Horizon

On a low-wing airplane, you typically have to use a little more imagination to use the wing as a reference pointer — but it is definitely possible and definitely worth the effort. Sometimes it helps to envision the chord line with your mind's eye. If you control the angle between the chord line and the lateral horizon, you are controlling pitch attitude.

The idea that you can control pitch attitude while looking out the side window is very important. Aerobatics pilots often attach crosshair-like pointers to their wings, just so they can be sure to have an easy-to-use pitch attitude reference when they're looking out the side. Conversely, it is common to find students who (although they can fly OK while looking out the front) lose control of pitch as soon as they try to look out the side; this makes it tough to check landmarks or scan for traffic.

The fourth way of perceiving pitch attitude is to use the attitude indicator instrument — the artificial horizon. This has the drawback that it is much too close to your eye; you can't look at the attitude indicator and look for traffic at the same time. You should use outside pitch references whenever possible.

[2.6](#) Making Changes in Angle of Attack

The push/pull motion of the yoke and the trim wheel are part of the same system, jointly controlling the angle of attack. They also jointly control airspeed, as discussed in [section 2.12](#).

If you want to make a temporary increase in angle of attack, just raise the nose by applying a little back pressure on the yoke. When you reach the new pitch attitude, you can release most of the pressure, and for the first few moments the airplane will maintain the new pitch attitude. Then, as it decelerates, you will need to maintain progressively more back pressure in order to maintain the new pitch attitude (and new angle of attack). After a few seconds things will stabilize at a new pitch attitude, a new angle of attack, and a new airspeed. At this point, if you release the back pressure, the airplane will want to drop its nose so it can return to its trimmed angle of attack.

If you push or pull the airplane off its trim speed and then suddenly let go of the yoke, the airplane will not return smoothly and immediately to its trim speed; there will be some phugoid oscillation (as discussed in [section 6.1.12](#)). To undo a temporary change in angle of attack, the proper technique requires observing and controlling the pitch attitude. Let the nose drop to the correct pitch attitude, then apply enough back pressure to keep it from dropping farther. Then, as the airplane gradually returns to its trim speed, you will need progressively less pressure.

Similar logic applies to making long-term changes in angle of attack. Use the yoke to change the pitch attitude. At first very little pressure will be required to maintain the new pitch attitude. Then, as the airspeed changes, use pressure on the yoke to keep the attitude where you want it. Make the change permanent by using the trim wheel to trim off the applied pressure.

Let's see how these ideas apply to a typical maneuver: levelling off from a climb. Initially let's suppose you start out nicely trimmed, climbing at 475 feet per minute at 90 knots true airspeed.⁹ As discussed in [section 2.11](#), that means your direction of flight is 3 degrees above the horizon. As shown in [figure 2.5](#), the first step in the level-off is to change your direction of flight so it becomes horizontal. During the brief time that the direction of flight is changing, the aircraft will be out of equilibrium; lift will be less than weight. The load on the aircraft and its occupants will be slightly less than one *G*.




Figure 2.5: Level-Off Maneuver

As the direction of flight changes, you will need to lower the nose the same amount (three degrees). At this point, since the direction of flight and the pitch attitude have changed together, the angle of attack is (for the moment) the same as it was during the climb. This can be seen by comparing the top two parts of [figure 2.6](#). The airspeed is still 90 knots, which is the trim speed, so no yoke force will be needed to maintain the new attitude (for the moment). So far so good.




Figure 2.6: Angle of Attack during Level-Off

Since the airplane is no longer climbing, the engine power that had previously been devoted to increasing the altitude is now being devoted to increasing the airspeed. (See [section 1.3.1](#).)

As the airplane gradually accelerates from climb speed to cruise speed, the direction of flight remains horizontal, so the pitch attitude gradually decreases as the angle of attack decreases. This is shown in the bottom part of [figure 2.6](#). You need to apply progressively more forward pressure. In a trainer you might trim off this pressure all at once, but in a high-powered airplane you need to re-trim repeatedly, in stages, as the airplane keeps accelerating and accelerating.

Eventually, the airplane will reach cruise speed. At this point, the airplane has all the altitude (potential energy) and airspeed (kinetic energy) that it needs, so you should throttle back to cruise power. Now [10](#) you make your final trim adjustment and the level-off maneuver is complete.

Here is a useful trick: make a note of how much trim change is required in your favorite airplane to make the transition from climb to cruise. It is some definite amount, and remembering this amount obviates a lot of guessing and fiddling. I remember the amounts in terms of ``sectors'' and ``bumps''. That is, on most airplanes only a certain sector of the trim wheel is exposed, and this defines how much trim change can be achieved with a single hand motion; I call this one sector. Similarly, the trim wheel typically features a series of bumps, to make it easier to grasp. Each bump represents 1/4th or 1/5th of a sector.

Suppose after cruising in level flight for a while, you decide to climb to a higher altitude. If you roll in three sectors of nose-up trim as you start the climb, you can bet that you will need roll those three sectors back out to return to cruise airspeed afterwards. [11](#)

Similarly, suppose you are cruising along and encounter an updraft. If you roll in half a bump of nose-down trim to help you maintain altitude, you can bet that you will need to roll that half-bump back out when you exit the updraft and return to normal airspeed. Keep track of the amount! Say to yourself, ``I'm carrying a half-bump of nose-down trim which I'll have to get rid of sooner or later''.

2.7 Fly with a Light Touch

Here's some really important advice: Fly with a light touch. You need to feel if you are pulling the airplane off its trim speed.

Some airplanes have such heavy control forces that it's hard to imagine anyone accidentally pulling the plane off its trim speed. You need to trim it properly lest you wear yourself out trying to hold the yoke. For some maneuvers (such as the landing flare) you need to apply a significant amount of force — but please be aware of what you are doing.

Some other planes have such light control forces that if you have the proverbial death grip on the yoke, you could pull the airplane ten knots away from its trim speed without feeling it.

I once flew with a pilot who held the yoke so tightly that his knuckles turned white, literally. Every time he looked to the right, the airplane pitched down 10 or 15 degrees. Every time he looked to the left, the airplane pitched up 10 or 15 degrees. It's a good thing he didn't look to the left very long; otherwise we might have stalled.

For any plane, C-152 or Airbus, if you trim it properly you will be able to fly most maneuvers using just your thumb and one or two fingertips.

The yoke is not just a control carrying commands from you to the airplane — it is also a valuable sensor carrying information from the airplane to you. This is discussed in more detail in [section 12.12](#).

You should make sure the airplane is at all times trimmed for the right airspeed (or, rather, angle of attack). You should be aware of (and wary of) any force you apply to the yoke, forcing the airplane off its trim speed.

Fly with a light touch!

[2.8](#) Trim Won't Solve All The World's Problems

Although the airplane's tendency to return to its trimmed angle of attack is very powerful, very important, and usually very helpful, there is more to the story.

If the airplane is disturbed from its trimmed angle of attack, it will not just return; it will overshoot. It will oscillate a few times before settling down. These phugoid oscillations are slow enough that you can easily extinguish them by timely pressure on the yoke, as discussed in [section 6.1.12](#).

In smooth air, you can trim the airplane and let it fly itself. However, turbulent air will frequently provoke new phugoid oscillations so you will frequently need to apply small nudges to the yoke.

For similar reasons, it is not normal procedure to use the trim wheel to *initiate* a change in pitch attitude, airspeed, or angle of attack. That would just provoke an oscillation. Initiate the change with the yoke as described above. Put the pitch attitude where it belongs, keep it there with the yoke, and then trim off the pressure.

Finally, in some airplanes the trim speed is perturbed when you add power, when you extend flaps, and especially when you have power and flaps at the same time. See [section 5.5](#) and [section 12.10](#).

[2.9](#) Pitch Attitude versus Angle of Attack

The previous sections pointed out that while pitch attitude and angle of attack are related, they are not quite the same. Pitch attitude is measured relative to the horizon, but angle of attack involves the direction of the relative wind. In any situation where the relative wind is not horizontal, we have to be careful.

I forgot the distinction once; let me tell you the story. One summer I spent several weeks at the Aspen Center for Physics. This was my first opportunity to do any mountain flying, so I arranged for a lesson from the flight school at Aspen. The lesson included flying over the continental divide and landing at Leadville. Leadville is famous for being the highest airport in the United States — 9900 feet above sea level. On the day in question, it was about 90°F in the shade, so the density altitude at Leadville was around 13,000 feet, and I knew takeoff performance would be critical.

I used my best short-field procedure, even though the runway was 5000 feet long. I accelerated on the runway to the proper climb-out speed (75 knots indicated, 90 knots true) and then rotated to what I assumed was the correct climb-out attitude. Based on my experience at lowland airports, I knew that 11 degrees of nose-up attitude was usually just right for climb out. Following my usual habit, I scanned the airspeed indicator after we had climbed a few feet. To my horror, I observed that the airspeed was decreasing rapidly. I immediately lowered the nose, and flew the airplane in ground effect while it regained speed. (What had been intended as a short-field procedure ended with a peculiar imitation of soft-field procedure.) I used up

almost the entire runway getting back to 75 knots. At 75 knots I rotated again, choosing a much lower pitch attitude this time. We climbed out at 75 K_IAS and the rest of the lesson was relatively uneventful.

This situation is depicted in the following figures and table. [Figure 2.7](#) shows the normal takeoff procedure at a low-altitude airport. [Figure 2.8](#) shows that using the normal pitch attitude does not produce the normal angle of attack at a high-altitude airport, because the angle of climb is an indispensable part of the equation. [Figure 2.9](#) shows how to do it right. [Table 2.2](#) summarizes the arithmetic.



[Figure 2.7](#): Climbing Out From Lowland International



[Figure 2.8](#): Climbing Out From Leadville (Wrong)



[Figure 2.9](#): Climbing Out From Leadville (Right)

	Calib. Airspeed	Pitch Attitude	Incidence	Climb Rate @ True Airspeed	Angle of Climb	Angle of Attack
sea level	76 KCAS	11.0	4.5	900 fpm @ 76 KTAS	7	8.5
Leadville (wrong)	dropping rapidly	11.0	4.5	200 fpm @ 90 KTAS	1	14.5
Leadville (right)	76 KCAS	5.0	4.5	200 fpm @ 90 KTAS	1	8.5

[Table 2.2](#): Right versus wrong climb attitude

Understanding what went wrong in this scenario is very instructive. The main difference between a sea-level takeoff and a mountain takeoff is that the airplane does not climb nearly so steeply. The direction of flight is much more nearly horizontal. As can be seen by comparing [figure 2.8](#) with [figure 2.9](#), this means a much lower pitch attitude is needed to achieve the same angle of attack.

The really embarrassing part of my story is that I had actually calculated the climb gradient as part of my preflight preparation, to make sure I could clear obstructions. I just didn't make the connection between the climb gradient (which I calculated) the best-climb angle of attack (which I knew) and the pitch attitude (which I used for controlling the airplane). Fortunately I did know the connection between airspeed and angle of attack, and I scanned the airspeed indicator before the situation got too far out of hand.

[2.10](#) Power plus Attitude does not equal Performance

You may have heard the assertion that "Power plus Attitude equals Performance". Well, that assertion is not quite right, and has caused all sorts of unnecessary confusion.

Consider the following scenario: You are cruising along in a typical 180 horsepower, one-ton aircraft. You have constant power and constant attitude, so you expect constant performance. You do indeed get constant performance, and everything seems just fine.

Now, just raise the nose to a 15 degree nose-up attitude, and hold that attitude as accurately as you can. You will once again have constant power and constant attitude, so you might expect constant performance — but that is definitely not what you will get. Instead, you will get decreasing airspeed and increasing angle of attack. The initial climb that looked so promising will peter out and you will wind up on the edge of a stall.

If you think about this situation in terms of *energy* and *angle of attack*, the airplane's behavior is completely predictable.

First of all, we need to remember that not all climbs are steady climbs. As portrayed in [figure 2.10](#), it is possible for a roller-coaster with no engine at all to zoom up a little ways by cashing in its initial kinetic energy. Just because it starts out on a certain climb trajectory doesn't mean it can continue.



[Figure 2.10](#): Climb Powered by Speed

Airplanes, too, can be placed on climb trajectories that cannot be sustained by the available engine power. The initial climb succeeds only because airspeed is being cashed in to purchase altitude.

Unlike a roller-coaster, the airplane will not stay on its initial trajectory until it runs out of speed altogether. As the airspeed decays, the airplane will have to fly at a higher angle of attack in order to support its weight. Since, as discussed above, the angle of attack depends on the angle between the pitch attitude and the direction of flight, a constant attitude implies a non-constant direction of flight, as indicated in [figure 2.11](#).



[Figure 2.11](#): Constant Power & Attitude but Changing Performance

If you are lucky, the changing flight path will result in a trajectory where the rate of climb and the drag budget can be sustained by engine power, with no further decrease of airspeed; otherwise the maneuver will end in a stall.

One of the maneuvers you have to perform in order to get a commercial pilot certificate is called a *chandelle*. As discussed in [section 16.11](#), it involves turning as well as climbing, but if you disregard the turning part, the maneuver is exactly what is portrayed in [figure 2.11](#). This maneuver is an important part of the syllabus because it forces people to learn that constant power and constant attitude do not imply constant performance.

As discussed in [section 2.6](#), a pitch attitude excursion is not necessarily the same as an angle of attack excursion. Suppose due to turbulence or whatever, the pitch attitude and direction of flight both increase by 15 degrees. If you correct the situation promptly, the airspeed and altitude will not have time to change much. If on the other hand you allow the pitch excursion to persist, the airplane will begin to follow the chandelle trajectory shown in [figure 2.11](#). The altitude will increase (at least at first), the airspeed will decrease, and the angle of attack will increase. It is good pilot technique to correct pitch attitude excursions before they turn into altitude / airspeed / angle of attack excursions.

To summarize: the Leadville scenario and the chandelle scenario prove that angle of attack is far more important than pitch attitude in determining performance. But this does not mean you disregard pitch attitude — far from it. I recommend that you use pitch attitude as a *means* of controlling angle of attack — just don't use pitch attitude as a *substitute* for controlling angle of attack.

[2.11](#) Estimating the Relative Wind

As discussed above, to control the angle of attack you need to know both the pitch attitude and the direction of flight. [12](#) I have given several methods for estimating the pitch attitude. Now it is time to explain how to estimate the direction of the relative wind. This is almost the same thing as estimating direction of flight.

In level flight, the task is easy: The relative wind is coming at you horizontally. (Again, I am assuming there are no major updrafts or downdrafts.)

If the airplane is climbing or descending, the origin of the relative wind will be above or below the horizon, respectively. The amount above or below depends on the ratio of your vertical speed to your airspeed. I have

committed some of the numbers to memory; for instance, I know that flying a standard 3 degree glideslope at 90 knots involves a 480 fpm descent. Using the same little fact in reverse tells me that if I am climbing out at 90 knots and the vertical speed indicator (VSI) is reporting 480 fpm, I must be flying toward a point 3 degrees above the horizon; to say it the other way, the relative wind must be coming toward me from that point, three degrees above the horizon. That means that I can relabel the VSI as a ``direction of flight" indicator, as shown in [figure 2.12](#). Any particular relabeling is only valid for one airspeed.¹³



[Figure 2.12](#): Vertical, Horizontal Speed Gauges Determine Angle

If you maintain 90 knots and transition from level flight to a 480 fpm climb, you will have to raise the pitch attitude 3 degrees in order to maintain the same angle of attack.¹⁴

If you want to know the vertical speed that corresponds to some other angle and/or some other horizontal speed, you can refer to [table 2.3](#); a similar table appears in every ``instrument approach procedures" booklet published by the US government. The inverse table (finding the angle, given horizontal and vertical speeds) is shown in [table 2.4](#).

Horizontal speed / knots	
60	
75	
90	
105	
120	
3°	
320	
400	
480	
555	
635	
4°	
425	
530	
635	
745	
850	
5°	
530	
665	
795	
930	
1065	
6°	
640	
800	
960	
1120	
1275	
7°	
Angle	

745
935
1120
1305
1490

8°
855
1065
1280
1495
1710

[Table 2.3](#): Vertical Speed vs. Angle and Horizontal Speed

Horizontal speed / knots	
Vertical Speed	60
	75
	90
	105
	120
	250
	2.4
	1.9
	1.6
	1.3
	1.2
	500
	4.7
	3.8
	3.1
	2.7
	2.4
	750
	7.0
	5.6
	4.7
	4.0
	3.5
	1000
	9.3
	7.5
	6.3
	5.4
	4.7

[Table 2.4](#): Angle vs. Vertical Speed and Horizontal Speed

The VSI is not the only way of determining the direction of flight. If you are established on an ILS approach, as long as the glideslope needle stays centered you are descending at a known angle (usually three degrees). Similarly, there might be a VASI or other approach slope indicator that you could follow. As always, it is better to use outside references instead of instruments.

Perhaps the best way to judge the angle of descent is to use the "rule of thumb" as discussed in [section 12.3](#). That frees you from relying on any fancy equipment.

If you control the direction of flight using any of these techniques, and control the pitch attitude using the techniques discussed elsewhere in this chapter, then you are also controlling the angle of attack.

Actually, there is one more ingredient in this recipe: the wind. Three of the methods just mentioned (VASI, electronic glideslope, and rule of thumb) give you information about your direction of flight relative to the ground, but the angle of attack depends on your direction of flight *through the air*. In the presence of wind, the two are not quite the same. This is discussed in [section 12.4.3](#).

The scheme of estimating the direction of flight using the VSI gives the correct answer even when nature's wind is blowing (provided, again, there are no major updrafts or downdrafts).

Outside references should be your primary means of controlling angle of attack. Every so often you should look at the airspeed indicator to make sure you have got the *right* angle of attack (as discussed in [section 2.12](#)), but you should maintain that angle of attack by outside references.

Suggestion: one look out of ten, look at the instruments; nine looks out of ten, look at the outside references.

[2.12](#) Airspeed Is Related to Angle of Attack

[2.12.1](#) Airspeed versus Coefficient of Lift

So far in this chapter I have mentioned that the critical performance numbers usually specified by airspeeds such as V_Y are really angle of attack recommendations.

- I have mentioned that the trim wheel really controls angle of attack but to a good approximation controls airspeed.
- I have mentioned that the airspeed indicator saved my bacon when I had an angle of attack problem at Leadville.

Therefore you are probably beginning to suspect that there might be a relationship between angle of attack and airspeed. That's right! The purpose of this section is to tell you why you can use the airspeed indicator to control angle of attack, when you have to compensate for its imperfections, and when you can't trust it at all.

The basic line of reasoning is this: the amount of lift produced by the wing depends on angle of attack and calibrated airspeed. We can turn this around to get a simple relationship between airspeed and angle of attack (assuming lift is known, as it usually is). The key formula is

$$\text{lift} = \frac{1}{2} \rho v^2 \times \text{coefficient of lift} \times \text{area} \quad (2.1)$$

The coefficient of lift will be discussed below, and (in more detail) in [section 4.4](#). The quantity $\frac{1}{2} \rho v^2$ is called the *dynamic pressure*, also called Q for short, but more often than not people just call it one-half rho vee squared.

You don't need to calculate $\frac{1}{2} \rho v^2$ because your airspeed indicator does it for you. You may have thought that an airspeed indicator would ideally measure the *true airspeed* (which is simply the genuine speed of the air relative to the aircraft, denoted V in all the formulas). However, the airspeed indicator doesn't even try to measure V (i.e. the square root of v^2); instead it tries to measure something called *calibrated airspeed* (CAS), which is proportional to the square root of $\frac{1}{2} \rho v^2$. Note the factor of ρ in the CAS formula.¹⁵ While we're on the subject, *indicated airspeed* (IAS) refers to whatever is indicated on your airspeed indicator. It is the same as calibrated airspeed, plus whatever errors there are in the mechanism. This discussion assumes that your instrument is not too wildly inaccurate, so that formulas that apply to CAS exactly also apply to IAS accurately enough for present purposes.¹⁶

In flight, the lift is nearly always equal to the weight times load factor.¹⁷ The weight is presumably not changing much from moment to moment. This leads us to rearrange the lift equation as follows:

$$\text{coefficient of lift} = (\text{weight} \times \text{load factor}) / (\frac{1}{2} \rho v^2 \times \text{area}) \quad (2.2)$$

If the airspeed goes down, the coefficient of lift must go up. This relationship is illustrated in [figure 2.13](#).




Figure 2.13: Airspeed versus Coefficient of Lift

Three of the critical V-numbers are marked in [figure 2.13](#); each corresponds to a particular coefficient of lift.

2.12.2 Coefficient of Lift versus Angle of Attack

Now we bring in a new fact: The coefficient of lift is a simple function of the angle of attack. This dependence is shown in [figure 2.14](#). Note that for small angles of attack, the coefficient of lift is essentially proportional to the angle of attack. The angle of attack that gives the maximum coefficient of lift is called the "critical angle of attack" and is marked in the figure.




Figure 2.14: Coefficient of Lift versus Angle of Attack




Figure 2.15: Airspeed Is Related to Angle of Attack

By combining this fact with what we already know, we can establish the relationship between angle of attack and indicated airspeed. We combine [figure 2.13](#) with [figure 2.14](#), as is done in [figure 2.15](#). We see that a particular V-number, such as V_{NE} , corresponds to a particular coefficient of lift, which in turn corresponds to a certain angle of attack. The same goes for most of the other V-numbers, such as V_Y . The argument works in reverse, too: any particular angle of attack corresponds to a particular airspeed (assuming we know how much lift is being produced).

We conclude that the airspeed indicator is really a pretty good angle of attack indicator — with one major exception.

Here's the exception: There is a whole range of angles of attack near the critical angle of attack that all produce about the same coefficient of lift (because the coefficient of lift versus angle of attack curve is quite flat on top). All the coefficient-of-lift values in this small range correspond to nearly the same airspeed — namely V_S , the stalling airspeed.

The stall is a very critical flight regime. This is a regime where you would very much like to have an accurate instrument to indicate angle of attack, and alas it is the one regime where the airspeed indicator doesn't tell you anything you need to know.

You want to land the airplane at a very high angle of attack. You will have to perceive the angle of attack

using outside visual cues, as discussed in the previous sections. During the flare, the airspeed indicator doesn't tell you anything you need to know. I once asked an airline captain to tell me at what airspeed his airliner touched down. He said ``I don't know; I never looked. I've always had more important things to look at''. That was a good pilot's honest answer.

2.12.3 Correcting for Reduced Density

In all non-stalling regimes of flight, including (especially) final approach, the airspeed indicator provides your most quantitative information about angle of attack. We now discuss some corrections that may be needed.

The airspeed indicator is basically a pressure gauge; the pressure that moves the airspeed needle is the same dynamic pressure that holds up the wings in accordance with the lift formula ([equation 2.1](#)). Knowing the pressure that holds up the wing is more important than knowing your true airspeed.

The airspeed indicator is doing you a favor by not measuring speed *per se*. For instance, on final approach you should fly the proper indicated airspeed. At high density altitudes this will be a higher-than-normal true airspeed. Remember that calibrated airspeed is what holds up your wings.

In other words: do not correct V_X , V_Y , V_S , or approach speed ($1.3 V_{S0}$) for density altitude. Trust the calibrated airspeed.

The true airspeed that corresponds to this calibrated airspeed will be higher, by about 2% per thousand feet of density altitude. Your groundspeed will also be greater.

When landing at a high-altitude airport, the greater groundspeed means you will need more runway length, by about 4% per thousand feet of density altitude. Check the charts in your POH.

A high-altitude takeoff is even worse than the landing, because the engine (unless turbocharged) will be producing less power. Check the charts in your POH. Apply a generous safety factor, since many of the handbooks are disgracefully overoptimistic. Do the takeoff planning (not just the landing planning) before you land, lest you go into an airport you can't get out of.

2.12.4 Correcting for Reduced Lift Requirements

So far we have been assuming the weight was equal to some standard value. Let's relax that assumption and see what happens.

As mentioned in [section 7.5.7](#), it is easy to imagine flying a Cherokee Six at half of its maximum legal weight.

The problem is that the Pilot's Operating Handbook for the airplane specifies all the critical angle of attack information in terms of speeds — speeds that only apply at max weight. We know that the airplane stalls at a definite angle of attack, not at a definite airspeed or anything else.

In general, if you keep the angle of attack constant and lower the weight of the airplane by 10%, the airspeed needed to support that weight goes down by 5%. This is because the lift depends on the *square* of the airspeed in [equation 2.1](#); the square root of 0.90 is 0.95 and the square root of 1.10 is 1.05. For really large changes in weight, the speed change is even somewhat greater; the square root of 0.50 is not 0.75 but rather 0.707.

At reduced weights, your approach speed and all the other critical performance speeds must be reduced below their standard-weight handbook values.

The percentage change in speed
is half of the percentage change in weight.

Since the cruise speed depends on power and hardly depends on angle of attack, it does *not* decrease as the weight is decreased; the situation is depicted in [figure 7.13](#) in [section 7.5.7](#). The maneuvering speed *does* decrease as weight is decreased, as discussed in [section 2.13.2](#).

2.12.5 Correcting for Increased Lift Requirements

There is one fairly common situation where maintaining a given angle of attack requires flying at airspeeds

above the V-numbers given in the Pilot's Operating Handbook.

In a steep turn, the wings are required to produce enough lift not only to support the airplane's weight, but also to shove it around the corner. In a 60 degree bank, the lift requirement is doubled. We say there is a load factor of 2.0. The airspeed necessary to produce this lift at a given angle of attack is increased by a factor of $\sqrt{2}$, which is 1.41.

If you are going to use the airspeed indicator as a source of angle of attack information, you have to take this into account. If you fly at a speed near the bottom of the green arc in a steep turn, the airplane will stall. For example, if the airplane stalls at 60 knots in unaccelerated flight, it will stall at 85 knots in a 60 degree banked turn (since $60 \times 1.41 = 85$).

Also remember that the airplane is trimmed for a definite angle of attack, and it really wants to maintain that angle of attack. If you are cruising along, trimmed for 120 knots in straight and level flight, and the airplane gets into a 60 degree bank, it will accelerate to 169 knots (120 times the square root of 2) in order to meet the increased lift requirement at the same angle of attack. This situation is described in more detail in [section 6.2](#).

2.12.6 Compute with Calibrated not Indicated Airspeed

In a wide range of airplanes, it turns out that a good airspeed for final approach is 1.3 times the stalling speed.¹⁸

When applying this rule, a little sophistication is necessary, or you might get into trouble. In particular, you must not just look at the indicated stalling speed on the airspeed indicator, multiply by 1.3, and then try to use the result as your indicated airspeed on final.

The only safe way to calculate the approach speed is to multiply the *calibrated* stalling speed by 1.3, and then convert the result an indicated airspeed. That is, if you know the indicated stalling speed, the correct procedure is:

- a) convert the indicated speed to a calibrated speed, using the conversion information in the Pilot's Operating Handbook;
- b) multiply the calibrated speed by 1.3; and
- c) convert this calibrated approach speed back to an indicated airspeed you can use in the cockpit.

[Table 2.5](#) shows an example which contrasts the right and wrong calculations.

	CAS	IAS	safe approach speed?
stall	50	43	
$1.3 \times \text{indicated stall}$	58	56	no
$1.3 \times \text{calibrated stall}$	65	65	yes

[Table 2.5](#): Calibrated versus Indicated Approach Speed

The origin of the problem is this: It is possible to position the Pitot tube and static port so that the IAS is a few knots higher than the CAS in cruise conditions, yet a few knots lower than the CAS near the stall. Manufacturers commonly do this, presumably in hopes of making pilots think the airplane performs better than it really does.

These errors would not be much of a problem if the the IAS were simply proportional to the CAS. The constant of proportionality would drop out of the calculation, and you could skip steps (a) and (c). Alas, in many airplanes the errors are highly nonlinear. The indicated airspeeds are much too low at the low end of the scale. If you multiply such a low number by 1.3, you get a number that is still much too low, but falls at a

place where the gauge is more accurate, so you wind up with a real airspeed that is dangerously low.

You may be wondering about other calculations, such as the corrections for nonstandard weight. Should calculations also be done using calibrated airspeed? The answer, alas, is not 100% obvious. It depends on whether you think the errors in the system depend on airspeed itself, or depend on angle of attack.

- In the rather unlikely case that the error is in the gauge itself, it would be better to convert to CAS, multiply, and convert back to IAS.
- More often, the airspeed instrument itself is a very accurate pressure gauge, but the Pitot and static ports are positioned so that they pick up bogus pressures at high angles of attack. In such a case you should calculate the weight-correction using indicated airspeed directly. That is, if you are at 64% of standard weight, your indicated airspeed should be 80% of the standard indicated airspeed. The point is that you want to fly the maneuver at the correct angle of attack. If the errors depend only on angle of attack, they drop out of this calculation.

You might want to measure your airplane, as follows: Fly it at its maximum weight, at a safe altitude, and observe the indicated airspeed at which the stall warning horn comes on. Do this in the clean configuration¹⁹ and in the landing configuration. Then repeat the measurements at the lowest convenient weight. Then you will know for sure how the indicated airspeed varies with weight, at particular angles of attack.

2.12.7 Correcting for Slip

It is easy to get into situations where the indicated airspeed is wildly inaccurate. In some airplanes the opening that is supposed to measure the static pressure is located on the side of the fuselage, and during a slip that point is subject to some dynamic pressure in addition to the static pressure.²⁰

In such a case you must remember that it is angle of attack that really matters. You can use the airspeed indicator if you wish *before* the slip to help figure out what angle of attack you want, but *during* the slip you must maintain that angle of attack by looking at the angles themselves (pitch angle and direction of flight). See [section 11.2](#) for more on this.

2.12.8 Drag and Lift-to-Drag Ratio

Let's return to the scenario of the airplane flying at half of its standard weight, and ask (a) what is the best glide speed, and (b) how well will it glide at that speed.

To answer these questions we need to think about drag as well as lift. ([Section 2.12.4](#) concentrated on topics like V_S and V_A which depend on total lift, not lift-to-drag ratio.) Fortunately, the answer comes out the same. This is because the formula for drag,

$$\text{drag} = \frac{1}{2} \rho V^2 \times \text{coefficient of drag} \times \text{area} \quad (2.3)$$

has the same form as the famous formula for lift:

$$\text{lift} = \frac{1}{2} \rho V^2 \times \text{coefficient of lift} \times \text{area} \quad (2.4)$$

The key idea is that the coefficient of drag depends on angle of attack; at any particular angle of attack the coefficient does not perceptibly depend on weight or airspeed. The same is true of the coefficient of lift and the lift-to-drag ratio.

If you want to glide from point A to point B in no-wind conditions,²¹ the main thing you care about is lift-to-drag ratio. For example, if your airplane is capable of a 10-to-1 lift-to-drag ratio, then you can glide to a point that is 1/10th of a radian (i.e. six degrees) below the horizon.

The optimal lift-to-drag ratio is achieved at a definite angle of attack. To support the weight of the airplane at that angle of attack, you will need to fly at a speed proportional to the square root of the weight, for the reasons given in [section 2.12.4](#).

The lightly-loaded gliding airplane will have the same angle of descent, the same direction of flight, and the same total gliding distance, as indicated in [figure 2.16](#). The only difference is that it will have a slower descent rate and a slower forward speed; this is indicated in the figure by stopwatches that show how long it takes the plane to reach a particular point.



Figure 2.16: Angle of Glide Independent of Weight

The moral of the story is if you are flying a lightly-loaded airplane, you should fly it "by the numbers", namely the angle of attack numbers. The critical airspeed numbers (climb speed, approach speed, stalling speed, etc.) are all reduced according to half the weight-change percentage. That is, if you are 10% light, reduce the handbook speeds by 5%.

There is one well-known exception to the rule of thumb that says important performance speeds decrease as the weight decreases. That is, the cruising speed actually increases at reduced weights. This is not an exception to the real rule that speeds should vary with weight *at a given angle of attack*, because cruising speed is not tied to a particular angle of attack. If the airplane is lightly loaded, you can cruise at a lower angle of attack and a higher airspeed, since the wings need to do less work to support the weight of airplane.

2.13 Not Everything Depends on Angle of Attack

Some of the airplane's critical performance numbers depend directly on angle of attack, while others don't. It's somewhat useful to know which are which, so you can know which ones change with the weight of the airplane and which ones don't.

2.13.1 Explicit Airspeed Limits

There is a *normal-operations* airspeed, V_{NO} . This is indicated by the top of the green arc on the airspeed indicator. You should not exceed this speed except in smooth air, and then only with caution. The idea here is that you don't want to break the wing. There is a maximum coefficient of lift, and the lift force depends on this coefficient times calibrated airspeed squared. By limiting the airspeed, you limit the maximum force that the wing can produce. This is typically what determines V_{NO} .

There is also a *never-exceed* airspeed, V_{NE} . This is indicated by the top of the yellow arc, and by a red radial line on the airspeed indicator. As the name suggests, you should never exceed this speed under any circumstances. This limit depends on many things, including drag force on the primary structure (wings, tail, landing gear etc.); drag force on secondary items (antennas, fairings, etc.); instability of the structure and control systems due to flutter; and other nasty complications.

2.13.2 Maneuvering Speed

If you are flying in moderate or severe turbulence, you should keep your airspeed below the *maneuvering airspeed*, V_A . By the same token, you should avoid large, sudden deflections of the controls unless your airspeed is below V_A . The idea behind V_A is that you want the wing to *stall* before anything *breaks*. You may think that a stall is bad, but remember that you can recover from a stall much more easily than you can recover from a broken airplane.

Maneuvering speed means the wing is supposed to stall before it produces enough Gs to break any part of the airplane.

We say it is *supposed* to stall, not guaranteed to stall, because the formal definition of V_A takes into account only certain types of rough control usage, and only certain types of turbulence (namely purely vertical updrafts and downdrafts). In real life, other possibilities must be considered. For instance, if you start out at V_A and encounter a large horizontal wind shear, arbitrarily large forces can be developed. For this and several other reasons, the exact value of V_A should not be taken too literally.

Still, the general idea of V_A makes sense: If you observe or anticipate a situation that imposes large G loads on the airplane, you should slow down and/or confine yourself to gentler maneuvers.

Unlike V_{NO} , the maneuvering speed varies in proportion to the square root of the mass of the airplane. The

reason for this is a bit tricky. The trick is that V_A is not a force limit but rather an *acceleration* limit. When the manufacturers determine a value for V_A , they are not worried about breaking the wing, but are worried about breaking *other* important parts of the airplane, such as the engine mounts. These items don't directly care how much force the wing is producing; they just care about the acceleration they are undergoing.

By increasing the mass of the airplane, you decrease the overall acceleration that results from any overall force. (Of course, if you increase the mass of cargo, it increases the stress on the cargo-compartment floor — but it decreases the stress on unrelated components such as engine mounts, because the acceleration is less.)

This means you should put V_A along with V_S and V_Y etc. on your list of critical airspeeds that vary in proportion to the square root of the mass of the airplane. However, V_A depends on real *mass* not on weight, so unlike the others it does *not* increase with load factor.

To illustrate this point, consider what happens when the airplane is in a steep turn. Compared to unaccelerated flight:

1. The stalling speed increases (because the stalling angle of attack stays the same), and
2. the airspeed for best rate of climb increases (because the optimum angle of attack stays the same), but
3. the maneuvering speed remains the same (since it doesn't directly depend on angle of attack).

Finally, we should note that there are two different concepts that, loosely speaking, are called maneuvering speeds.

- The *design* maneuvering speed, which we can denote $V_{A(D)}$, is primarily of interest to aircraft designers, not pilots. The designer must choose a value for $V_{A(D)}$ and then build an aircraft strong enough to withstand certain stressful maneuvers at that speed. Higher values of $V_{A(D)}$ promote safety, by forcing the design to be stronger.
- The maneuvering speed *limitation*, which we can denote $V_{A(L)}$, is of interest to pilots. It is an operating limitation. It appears on a placard in the cockpit. Lower values of $V_{A(L)}$ promote safety, by restricting certain operations to lower, less-stressful airspeeds.

This is a book for pilots, not designers, so when we use V_A it always means $V_{A(L)}$. But you should be careful when reading the FARs and other books, because they sometimes use the same symbol to mean two different things, which makes it very hard to think clearly.

2.13.3 Overview of Limits and Performance Numbers

We see that there are four main classes of numbers:

- The low-speed limits and most of the "optimum" numbers (including the stall, best rate of climb, best lift/drag ratio, best endurance, normal approach, short field approach, etc.) involve, to an excellent approximation, definite angles of attack. The corresponding airspeeds vary in proportion to the square root of weight. (To a second approximation each of these angles of attack will change slightly because of propwash over the wings, propeller efficiency, and other factors that depend on speed and power.)
- The high-speed limits (including never-exceed speed, the normal-operations limit, the landing-gear operation limit etc.) are, for all practical purposes, definite indicated airspeeds.
- The maneuvering limit is not exactly a definite speed (since the limiting speed varies in proportion to the square root of mass) but it is not exactly a definite angle of attack either (since V_A does not depend on load factor in a steep turn).
- Top speed and normal cruise speed depend on weight and engine power, as discussed in [section 7.6.5](#). They also depend on density altitude, as discussed in [section 7.5.4](#). Best angle of climb depends on weight, power, and wind, as discussed in [section 7.5.3](#).

2.14 Relative versus Absolute Angle of Attack

You can skip this section unless you are trying to compare this book with another book that uses a different

definition of ``the" angle of attack.

As mentioned in connection with [figure 2.1](#), we are free to choose how the angle-of-attack reference stick is aligned relative to the rest of the wing. Throughout this book, we choose to align it so that zero angle of attack corresponds to zero coefficient of lift. According to the standard terminology, the angle measured relative to the zero-lift direction is called the *relative* angle of attack.

Some other books try to align the reference with the chord line²² of the wing. According to the standard terminology, the angle measured relative to chord line is called the *absolute* angle of attack.

If you try to compare books, there is potential for confusion, because this book uses ``angle of attack" as shorthand for relative angle of attack, while some other books use the same term as shorthand for absolute angle of attack. To make sense when comparing books, you must do away with all the shorthand and use the fully explicit terms. In particular:

$$\begin{aligned} \text{relative angle of attack} &= \\ \text{absolute angle of attack} + k & \\ \text{absolute angle of attack} &= \text{(2.5)} \\ \text{relative angle of attack} - k & \end{aligned}$$

where $-k$ is the X -intercept of the graph of coefficient of lift versus absolute angle of attack, when graphed using whatever reference the other book is using to define absolute angle of attack. The X -intercept is always zero in this book.

Note that the names ``absolute" versus ``relative" don't mean anything. Neither quantity is more absolute or more relative than the other.

Also note that are not just two possibilities; the choice of reference is really quite arbitrary. It is perfectly valid to measure angles relative to the bottom surface of the wing, or the long axis of the fuselage, or any other reference you choose, provided you are consistent about it.

Using the chord as a reference works OK if you are only talking about one *section* of a *plain* wing. On the other hand:

- On typical airplanes, the chord of the wing tip points in a different direction from the chord of the wing root. Which one should be considered ``the" reference?
- When you extend the flaps, the chord line obviously changes. (See [section 5.4.3](#) for more on this.) Oddly, most books continue to measure angles relative to where the chord of the unflapped wing would have been. This makes the experiments easier. It is mostly harmless. It certainly highlights the arbitrariness of the choice of reference.

The simple rule ``pitch plus incidence equals angle of attack plus angle of incidence" ([figure 2.2](#)) is always technically valid (because the arbitrariness in the angle of incidence cancels the arbitrariness in the angle of attack). But if you want the rule to be *useful* in situations where flap settings are changing, you need to choose *relative* angle of attack or something very similar; otherwise you don't know what angle of attack produces a given amount of lift.

[2.15](#) Summary

- Trim for angle of attack! *Trim for angle of attack!*
- Push/pull motion of the yoke is an extension of the trim wheel, convenient for making changes in angle of attack.
- Fly with a light touch, so you can feel if you are pulling the aircraft off its trim speed.
- Pitch attitude is not the same as angle of attack. Angle of attack is what really matters.
- You can observe pitch attitude and direction of flight as a *means* for controlling angle of attack.
- The airspeed indicator gives you quantitative information about angle of attack (except near the stall).

- If the aircraft is producing a non-standard amount of lift, many of the critical V-numbers must be corrected. The percentage change in speed is half the percentage change in weight.

1

Wind may affect what angle of attack gives the best angle of climb or best angle of glide; see [section 7.5.6](#).

2

Changes in engine performance may slightly affect what angle of attack gives the best climb rate or best climb angle.

3

The *relative wind* is defined to be the speed and direction that the air is moving relative to the airplane. (It is very, very different from the velocity of the wind relative to the ground.) Unless otherwise specified, the relative wind is measured at a place where the airmass has not been greatly disturbed by the passage of the airplane. See [section 2.11](#) for more details.

4

See [section 19.6.2](#) for a more formal definition.

5

The effect of flaps is discussed in more detail in [section 5.5](#). See also the table in [section 12.11](#).

6

I am imagining a day without appreciable updrafts and downdrafts, so that the relative wind is horizontal; otherwise the story gets a little more complicated.

7

adjusting power if necessary to maintain level flight

8

Most people are right-eye dominant. If neither eye is strongly dominant, you can choose one arbitrarily and close the other when checking the sight line.

9

The meaning of true (versus calibrated) airspeed will be discussed in [section 2.12](#).

10

Not sooner, since on most airplanes, power changes slightly affect trim speed.

11

Maybe not exactly, because your indicated airspeed at the new cruise altitude may be slightly different. Apply the expected amount of trim, see how it works, and then trim off any remaining yoke force.

12

Incidence is important too, but you rarely need to worry about it. It only changes when you are changing the flap setting.

13

I am glossing over the distinction between horizontal speed and total speed. The tables tabulate the tangent and arctangent, based on true horizontal and vertical motion. In a really steep dive, the airspeed indicator would indicate the total motion, which is the resultant of the horizontal and vertical components. If you really wanted to calibrate the angles against the VSI and the (total) airspeed, you should use the sine and arcsine. However, at any halfway reasonable angle, the difference is less than a percent or two. Don't worry about it.

14

See [section 12.4.3](#), including [figure 12.13](#), for an analogous discussion of angles relative to the ground.

15

The constant of proportionality is arranged so that at sea level, in standard conditions, the calibrated airspeed is identical with the true airspeed. Since you are almost always flying at altitudes above sea level, your true airspeed will almost always be larger than your calibrated airspeed.

16

... but beware: as discussed in [section 2.12.6](#), there are cases where it is quite important to keep track of the distinction between IAS and CAS.

17

Load factor is defined in [section 6.2.3](#). You don't need to worry about it except during aggressive maneuvers.

18

See [section 12.11.3](#) for more on this and related topics.

19

In general, *clean* refers to something with small parasite drag. The *clean configuration* refers to flaps

retracted, landing gear retracted (if possible), et cetera.

20

In other airplanes the static pressure is measured on a mast, far from the side of the fuselage, so this problem does not occur.

21

If you want to maximize gliding time instead of distance, or if you want to account for tailwinds and/or updrafts, see [section 7.5](#).

22


The chord line is the straight line drawn from the leading edge to the trailing edge, as defined in [figure 3.12](#).

3 Airfoils and Airflow

- Have you heard how to make a small fortune in the aviation business?
- Start with a large one.

3.1 Flow Patterns Near a Wing


In this chapter I will explain a few things about how air behaves as it flows past a wing. There will be lots of illustrations, such as [figure 3.1](#), produced by a wind-tunnel simulation¹ program that I wrote for my computer. The wing is stationary in the middle of the wind tunnel; air flows past it from left to right. A little ways upstream of the wing (near the left edge of the figure) I have arranged a number of smoke injectors. Seven of them are on all the time, injecting thin streams of purple smoke. The smoke is carried past the wing by the airflow, making visible *stream lines*.



[Figure 3.1](#): Flow Past a Wing

In addition, on a five-times closer vertical spacing, I inject *pulsed streamers*. The smoke is turned on for 10 milliseconds out of every 20. In the figure, the blue smoke was injected starting 70 milliseconds ago, the green smoke was injected starting 50 milliseconds ago, the orange smoke was injected starting 30 milliseconds ago, and the red smoke was injected starting 10 milliseconds ago. The injection of the red smoke was ending just as the snapshot was taken.

The set of all points that passed the injector array at a given time defines a *timeline*. The right-hand edge of the orange smoke is the "30 millisecond" timeline.



[Figure 3.2](#): Upwash and Downwash


[Figure 3.2](#) points out some important properties of the airflow pattern. For one thing, we notice that the air just ahead of the wing is moving not just left to right but also upward; this is called *upwash*. Similarly, the air just aft of the wing is moving not just left to right but also downward; this is called *downwash*. Downwash behind the wing is relatively easy to understand; the whole purpose of the wing is to impart some downward motion to the air.

The upwash in front of the wing is a bit more interesting. As discussed in [section 3.6](#), air is a fluid, which means it can exert pressure on itself as well as other things. The air pressure strongly affects the air, even the air well in front of the wing.

Along the leading edge of the wing there is something called a *stagnation line*, which is the dividing line between air that flows over the top of the wing and air that flows under the bottom of the wing. On an airplane, the stagnation line runs the length of the wingspan, but since [figure 3.2](#) shows only a cross section of the wing, all we see of the stagnation line is a single point.

Another stagnation line runs spanwise along the trailing edge. It marks the place where air that passed above the wing rejoins air that passed below the wing.

We see that at moderate or high angles of attack, the forward stagnation line is found well *below* and *aft* of the leading edge of the wing. The air that meets the wing just above the stagnation line will backtrack toward the nose of the airplane, flow up over the leading edge, and then flow aft along the top of the wing.



[Figure 3.3](#): Velocity Field of a Wing

[Figure 3.3](#) introduces some additional useful concepts. Since the air near the wing is flowing at all sorts of different speeds and directions, the question arises of what is the “true” airspeed in the wind tunnel. The logical thing to do is to measure the velocity of the *free stream*; that is, at a point well upstream, before it has been disturbed by the wing.

The pulsed streamers give us a lot of information. Regions where the pulsed streamers have been stretched out are high velocity regions. This is pretty easy to see; each pulsed streamer lasts exactly 10 milliseconds, so if it covers a long distance in that time it must be moving quickly. The maximum velocity produced by this wing at this angle of attack is about *twice* the free-stream velocity. Airfoils can be very effective at speeding up the air.

Conversely, regions where the pulsed streamers cover a small distance in those 10 milliseconds must be low-velocity regions. The minimum velocity is zero. That occurs near the front and rear stagnation lines.

The relative wind vanishes on the stagnation lines. A small bug walking on the wing of an airplane in flight could walk along the stagnation line without feeling any wind.^{[2](#)}

Stream lines have a remarkable property: the air can never cross a stream line. That is because of the way the stream lines were defined: by the smoke. If any air tried to flow past a point where the smoke was, it would carry the smoke with it. Therefore a particular parcel of air bounded by a pair of stream lines (above and below) and a pair of timelines (front and rear) never loses its identity. It can change shape, but it cannot mix with another such parcel.^{[3](#)}

Another thing we should notice is that in low velocity regions, the stream lines are *farther apart* from each other. This is no accident. At reasonable airspeeds, the wing doesn’t push or pull on the air hard enough to change its density significantly (see [section 3.4.3](#) for more on this). Therefore the air parcels mentioned in the previous paragraph do not change in area when they change their shape. In one region, we have a long, skinny parcel of air flowing past a particular point at a high velocity. (If the same amount of fluid flows through a smaller region, it must be flowing faster.) In another region, we have a short fat parcel flowing by at a low velocity.


The most remarkable thing about this figure is that the blue smoke that passed slightly above the wing got to the trailing edge 10 or 15 milliseconds *earlier* than the corresponding smoke that passed slightly below the wing.

This is not a mistake. Indeed, we shall see in [section 3.10.3](#) that if this were not true, it would be impossible for the wing to produce lift.

This may come as a shock to many readers, because all sorts of standard references claim that the air is somehow required to pass above and below the wing in the same amount of time. I have seen this erroneous statement in elementary-school textbooks, advanced physics textbooks, encyclopedias, and well-regarded pilot training handbooks. Bear with me for a moment, and I’ll convince you that [figure 3.3](#) tells the true story.

First, I must convince you that there is no law of physics that prevents one bit of fluid from being delayed

relative to another.



[Figure 3.4](#): Delay is Not Forbidden

Consider the scenario depicted in [figure 3.4](#). A river of water is flowing left to right. Using a piece of garden hose, I siphon some water out of the river, let it waste some time going through several feet of coiled-up hose, and then return it to the river. The water that went through the hose will be delayed. The delayed parcel of water will never catch up with its former neighbors; it will not even try to catch up.

Note that delaying the water did not require compressing the water, nor did it require friction.


Let's now discuss the behavior of air near a wing. We will see that there are two parts to the story: The *obstacle* effect, and the *circulation* effect.

The first part of the story is that the wing is an obstacle to the air. Air that passes near such an obstacle will be delayed. In fact, air that comes arbitrarily close to a stagnation line will be delayed an arbitrarily long time. The air molecules just hang around in the vicinity of the stagnation line, like the proverbial donkey midway between two bales of hay, unable to decide which alternative to choose.

Air near the wing is delayed relative to an undisturbed parcel of air. The obstacle effect is about the same for a parcel passing above the wing as it is for the parcel passing a corresponding distance below the wing. This effect falls off very quickly as a function of distance from the wing. You can see that the air that hits the stagnation line dead-on (the middle blue streamer) never makes it to the trailing edge, as you can see in all three panels of [figure 3.5](#).

When the wing is producing zero lift, the obstacle effect is pretty much the whole story, as shown in the top panel of [figure 3.5](#).

Now we turn to the second part of the story, the circulation effect. In [figure 3.5](#) the panels are labelled as to angle of attack. Lift is proportional to angle of attack whenever the angle is not too large. In particular, the zero-lift case is what we are calling zero angle of attack, even for cambered wings, as discussed in [section 2.2](#).



[Figure 3.5](#): Airflow at Various Angles of Attack

For the rest of this section, we assume the wing is producing a positive amount of lift. This makes the airflow patterns much more interesting, as you can see from the second and third panels of [figure 3.5](#). An air parcel

that passes above the wing arrives at the trailing edge early. It arrives early compared to the parcel a corresponding distance below the wing, with no exceptions. This is because of something called circulation, as will be discussed in [section 3.10](#).

We can also see that *most* of the air passing above the wing arrives early in absolute terms, early compared to an undisturbed parcel of air. The exception occurs very close to the wing, where the obstacle effect (as previously discussed) overwhelms the circulation effect.


Unlike the obstacle effect, the circulation effect drops off quite slowly. It extends for quite a distance above and below the wing -- a distance comparable to the wingspan.

A wing is amazingly effective at producing circulation, which speeds up the air above it. Even though the air that passes above the wing has a longer path, it gets to the back *earlier* than the corresponding air that passes below the wing. Note the contrast:

- The change in speed is temporary. As the air reaches the trailing edge and thereafter, it quickly returns to its original, free-stream velocity (plus a slight downward component). This can be seen in the figures, such as [figure 3.3](#) --- the spacing between successive smoke pulses returns to its original value.
- The change in relative position is permanent. If we follow the air far downstream of the wing, we find that the air that passed below the wing will never catch up with the corresponding air that passed above the wing. It will not even try to catch up.

3.2 Pressure Patterns Near a Wing

[Figure 3.6](#) is a contour plot that shows what the pressure is doing in the vicinity of the wing. All pressures will be measured relative to the ambient atmospheric pressure in the free stream. The blue-shaded regions indicate suction, i.e. negative pressure relative to ambient, while the red-shaded regions indicate positive pressure relative to ambient. The dividing line between pressure and suction is also indicated in the figure.



[Figure 3.6](#): Pressure Near a Wing

The pressure and suction created by the wing are conveniently measured in multiples of the dynamic pressure.⁴ It is usually represented by the symbol Q . For a typical general-aviation flight situation, Q is about half a pound per square inch. The maximum positive pressure on the airfoil is exactly equal to Q ; this occurs right at the stagnation lines.⁵ The maximum suction depends on the angle of attack, and on the detailed shape of the airfoil; for the situation in [figure 3.6](#) the max suction is just over $0.8 Q$. Each contour in the figure represents exactly $0.2 Q$ (roughly 0.1 psi).

There is a lot we can learn from studying this figure. For one thing, we see that the front quarter or so of the wing does half of the lifting. Another thing to notice is that suction acting on the top of the wing is vastly more important than pressure acting on the bottom of the wing. In [figure 3.6](#), the wing is flying at an angle of attack of 3 degrees, a reasonable "cruise" value.


At this angle of attack, there is almost no high pressure on the bottom of the wing; indeed there is mostly *suction* there.⁶ The only reason the wing can support the weight of the airplane is that there is *more* suction on the top of the wing. (There is a tiny amount of positive pressure on the rear portion of the bottom surface, but the fact remains that suction above the wing does more than 100% of the job of lifting the airplane.)⁷

This pressure pattern would be really hard to explain in terms of bullets bouncing off the wing. Remember, the air is a fluid, as discussed in [section 3.6](#). It has a well-defined pressure everywhere in space. When this pressure field meets the wing, it exerts a force: pressure times area equals force.

At higher angles of attack, above-atmospheric pressure does develop below the wing, but it is always less pronounced than the below-atmospheric pressure above the wing.

3.3 Stream Line Curvature

[Figure 3.7](#) shows what happens near the wing when we change the angle of attack. You can see that as the velocity changes, the pressure changes also.



[Figure 3.7](#): Airflow and Pressure Near Wings

It turns out that given the velocity field, it is rather straightforward to calculate the pressure field. Indeed there are two ways to do this; we discuss one of them here, and the other in [section 3.4](#).

We know that air has mass. Moving air has momentum. If the air parcel follows a curved path, there must be a net force on it, as required by Newton's laws.⁸

Pressure alone does not make a net force; you need a pressure *difference* so that one side of the air parcel is being pressed harder than the other. Therefore the rule is this: If at any place the stream lines are curved, the pressure at nearby places is different.

You can see in the figures that tightly-curved streamlines correspond to big pressure gradients and vice versa.

If you want to know the pressure everywhere, you can start somewhere and just add up all the changes as you move from place to place to place. This is mathematically tedious, but it works. It works even in situations where Bernoulli's principle isn't immediately applicable.

[3.4 Bernoulli's Principle](#)

We now discuss a second way in which pressure is related to velocity, namely Bernoulli's principle. In situations where this principle can be applied (which includes most situations), this is by far the slickest way to do it.

Bernoulli's principle is derived from the law of conservation of energy. It involves the kinetic energy of moving air and the potential energy stored in the "springiness" of the air. Just as energy can be stored in a wound-up spring, energy is stored in pressurized air.

Pressure, denoted P , is (by definition) a force per unit area, which is the same thing as an energy per unit volume:⁹

$$P = \text{Potential Energy per volume} \quad (3.1)$$

Meanwhile, moving air contains kinetic energy just like any other moving object:

$$\frac{1}{2}\rho v^2 = \text{Kinetic Energy per volume} \quad (3.2)$$

where v is the local velocity, and ρ (the Greek letter "rho") is the density, i.e. the mass per unit volume. Combining these, we conclude:

$$P + \frac{1}{2}\rho v^2 = \text{Mechanical Energy per volume} \quad (3.3)$$

Next, we make the approximation that we can ignore non-mechanical forms of energy (such as chemical reactions or heat produced by friction), and that we are not adding energy to the air using pumps, pistons, or whatever. Then, using the law that total energy cannot change (see [chapter 1](#)), we conclude that a given air parcel's mechanical energy remains constant as it flows past the wing.

Now, if the right-hand side of [equation 3.3](#) is a constant, it tells us that whenever a given parcel of air increases its velocity, it must decrease its pressure, and vice versa. This relationship is called Bernoulli's principle.

Higher velocity means lower pressure, and vice versa
(assuming constant mechanical energy).

Oftentimes¹⁰ it turns out that all the air parcels start out with the same mechanical energy. In such a case we can even make a Bernoulli-like statement comparing *different* parcels of air: Any fast-moving air must have lower pressure than any slow-moving air with the same mechanical energy.

Bernoulli's principle cannot be trusted if processes other than kinetic energy and pressure energy are important. In particular, in the "boundary layer" very near the surface of a wing, energy is constantly being dissipated (converted to heat) by friction. Fortunately, the boundary layer is usually very thin (except near the stall), and if we ignore it entirely Bernoulli's principle gives essentially the right answer.


3.4.1 Magnitude

It makes sense to measure the local velocity (lower-case v) at each point as a multiple of the free-stream velocity (capital V) since they vary in proportion to each other. Similarly it makes sense to measure relative pressures in terms of the dynamic pressure:

$$Q = \frac{1}{2}\rho V^2 \quad (3.4)$$

which is always small compared to atmospheric pressure (assuming V is small compared to the speed of sound). The pressure versus velocity relationship is shown graphically in [figure 3.8](#). The highest possible pressure (corresponding to completely stopped air) is one Q above atmospheric, while fast-moving air can have pressure several Q below atmospheric.

It doesn't matter whether we measure P as an absolute pressure or as a relative pressure (relative to atmospheric). If you change from absolute to relative pressure it just shifts both sides of Bernoulli's equation by a constant, and the new value (just as before) remains constant as the air parcel flows past the wing. Similarly, if we use relative pressure in [figure 3.8](#), we can drop the word "Atm" from the pressure axis and just speak of "positive one Q " and "negative two Q " --- keeping in mind that all the pressures are only slightly above or below one atmosphere.



[Figure 3.8](#): Pressure versus Velocity

Bernoulli's principle allows us to understand why there is a positive pressure bubble right at the trailing edge of the wing (which is the last place you would expect if you thought of the air as a bunch of bullets). The air at the stagnation line is the slowest-moving air in the whole system; it is not moving at all. It has the highest possible pressure, namely $\text{Atm} + Q$.

As we saw in the bottom panel of [figure 3.7](#), at high angles of attack a wing is extremely effective at speeding up the air above the wing and retarding the air below the wing. The maximum local velocity above the wing can be more than *twice* the free-stream velocity. This creates a negative pressure (suction) of more than $3Q$.

3.4.2 Altimeters; Static versus Stagnation Pressure

Consider the following line of reasoning:

1. The airplane's altimeter operates by measuring the pressure at the *static port*. See [section 20.2.2](#) for more on this.
2. The static port is oriented sideways to the airflow, at a point chosen so that the air flows past with a local velocity just equal to the free-stream velocity.

3. In accordance with Bernoulli's principle, this velocity must be associated with a "lower" pressure there.
4. You might think this lower pressure would cause huge errors in the altimeter, depending on airspeed. In fact, though, there are no such errors. The question is, why not?

The answer has to do with the notion of "lower" pressure. You have to ask, lower than what? Indeed the pressure there is $1 Q$ lower than the mechanical energy (per unit volume) of the air. However, in your reference frame, the mechanical energy of the air is $1 \text{ Atm} + 1 Q$. When we subtract $1 Q$ from that, we see that the pressure in the static port is just equal to atmospheric. Therefore the altimeter gets the right answer, independent of airspeed.

Another way of saying it is that the air near the static port has 1 Atm of potential energy (pressure) and $1 Q$ of kinetic energy.

In contrast, the air in the Pitot tube has the same mechanical energy, $1 \text{ Atm} + 1 Q$, but it is all in the form of potential energy since (in your reference frame) it has no kinetic energy.

The mechanical energy per unit volume is officially called the *stagnation pressure*, since it is the pressure that you observe in the Pitot tube or any other place where the air is stagnant, i.e. where the local velocity v is zero (relative to the airplane).

In ordinary language "static" and "stagnant" mean almost the same thing, but in aerodynamics they designate two very different concepts. The *static pressure* is the pressure you would measure in the reference frame of the air, for instance if you were in a balloon comoving with the free stream. As you increase your airspeed, the stagnation pressure goes up, but the static pressure does not.

Also: we can contrast this with what happens in a carburetor. There is no change of reference frames, so the mechanical energy (per unit volume) remains 1 Atm . The high-speed air in the throat of the Venturi has a pressure below the ambient atmospheric pressure.

3.4.3 Compressibility

First, a bit of terminology:

- *Pressure* denotes a force per unit area.
- *Compressibility* denotes a change in density in response to pressure.

Non-experts may not make much distinction between a "pressurized" fluid and a "compressed" fluid, but in the engineering literature there is a world of difference between the two concepts.

Every substance on earth is compressible --- be it air, water, cast iron, or anything else. It *must* increase its density when you apply pressure; otherwise there would be no way to balance the energy equations.

However, changes in density are not very important to understanding how wings work, as long as the airspeed is not near or above the speed of sound. Typical general aviation airspeeds correspond to Mach 0.2 or 0.3 or thereabouts (even when we account for the fact that the wing speeds up the air locally), and at those speeds the density never changes more than a few percent.

For an ideal gas such as air, density is proportional to pressure, so you may be wondering why pressure-changes are important but density-changes are not. Here's why:

- Lift depends on a pressure *difference* between the top and bottom of the wing. Similarly pressure drag depends on pressure *differences*. Therefore the relevant differential pressures are *zero* plus important terms proportional to $\frac{1}{2}\rho v^2$.
- Meanwhile, the density is some *big number* plus or minus unimportant terms proportional to $\frac{1}{2}\rho v^2$.

To say it again: Flight depends directly on *total* density but not directly on *total* atmospheric pressure, just differences in pressure.

Many books say the air is "incompressible" in the subsonic regime. That's bizarrely misleading. In fact, when those books use the words "incompressible flow" it generally means that the density undergoes only small-percentage changes. This has got nothing to do with whether the fluid has a high or low compressibility. The real explanation is that the density-changes are small because the pressure-changes are small compared

to the total atmospheric pressure.

Similarly, many books say that [equation 3.3](#) only applies to an “incompressible” fluid. Again, that’s bizarrely misleading. Here’s the real story:

1. Compressibility specifies to first order how density depends on pressure. [Equation 3.1](#) specifies to first order how the energy depends on pressure. It already accounts for the effects of compressibility and all other first-order quantities. Therefore [equation 3.3](#) is valid whenever the pressure-changes are a small percentage of the total pressure, regardless of compressibility.
2. At high airspeeds, the pressure changes are bigger, and you need a more sophisticated form of Bernoulli’s equation. As shown below, it is straightforward to include second-order terms --- which, by the way, don’t depend on compressibility, either. Indeed you can use the full equation of state, to derive Bernoulli’s equation in a form that is valid even for large-percentage changes in pressure. See [reference 2](#), page 29, equation 11.

Here is Bernoulli’s equation including the second-order term. I have rewritten it in terms of energy per mass (rather than energy per volume), to make it clear that compression doesn’t matter, since a parcel’s mass doesn’t change even if its volume and energy are changing:

$$\frac{P}{\rho_0} + \frac{1}{2} v^2 = \text{constant}$$

where ρ_0 is the density of air at atmospheric pressure, and where γ (gamma) is a constant that appears in the equation of state for the fluid. Its value ranges from 1.666 for helium, to 1.4 for air, to 1.0 for cool liquid water. It’s ironic that the correction is actually smaller for air (which has a high compressibility) than it is for water (which has a much lower compressibility).

In [equation 3.5](#), when the pressure P is near atmospheric, the term in square brackets approaches unity, and the expression becomes equivalent to the elementary version, [equation 3.3](#), as it should. This in turn tells us that the constant on the RHS of [equation 3.5](#) is equal to the mechanical energy that the parcel has whenever it is in free-stream conditions.

Don’t let anybody tell you that Bernoulli’s principle can’t cope with compressibility. Even the elementary version ([equation 3.3](#)) accounts for compressibility to first order.

3.5 Stall Warning Devices

We are now in a position to understand how stall warning devices work. There are two types of stall-warning devices commonly used on light aircraft. The first type (used on most Pipers, Mooneys, and Beechcraft) uses a small vane mounted slightly below and aft of the leading edge of the wing as shown in the left panel of [figure 3.9](#). The warning is actuated when the vane is blown up and forward. At low angles of attack (e.g. cruise) the stagnation line is forward of the vane, so the vane gets blown backward and everybody is happy. As the angle of attack increases, the stagnation line moves farther and farther aft underneath the wing. When it has moved farther aft than the vane, the air will blow the vane forward and upward and the stall warning will be activated.

The second type of stall-warning device (used on the Cessna 152, 172, and some others, not including the 182) operates on a different principle. It is sensitive to suction at the surface rather than flow along the surface. It is positioned just below the leading edge of the wing, as indicated in the right panel of [figure 3.9](#). At low angles of attack, the leading edge is a low-velocity, high-pressure region; at high angles of attack it becomes a high-velocity, low-pressure region. When the low-pressure region extends far enough down around the leading edge, it will suck air out of the opening. The air flows through a harmonica reed, producing an audible warning.




Figure 3.9: Stall Warning Devices

Note that neither device actually detects the stall. Each one really just measures angle of attack. It is designed to give you a warning a few degrees *before* the wing reaches the angle of attack where the stall is expected. Of course if there is something wrong, such as frost on the wings (see [section 3.13](#)), the stall will occur at a lower-than-expected angle of attack, and you will get no warning from the so-called stall warning device.

3.6 Air Is A Fluid, Not A Bunch of Bullets

We all know that at the submicroscopic level, air consists of particles, namely molecules of nitrogen, oxygen, water, and various other substances. Starting from the properties of these molecules and their interactions, it is possible to calculate macroscopic properties such as pressure, velocity, viscosity, speed of sound, et cetera.

However, for ordinary purposes such as understanding how wings work, you can pretty much forget about the individual particles, since the relevant information is well summarized by the macroscopic properties of the fluid. This is called the *hydrodynamic approximation*.

In fact, when people try to think about the individual particles, it is a common mistake to overestimate the size of the particles and to underestimate the importance of the interactions between particles.




Figure 3.10: The Bullet Fallacy

If you erroneously imagine that air particles are large and non-interacting, perhaps like the bullets shown in [figure 3.10](#), you will never understand how wings work. Consider the following comparisons. There is only one important thing bullets and air molecules have in common:

Bullets hit the bottom of the wing, transferring upward momentum to it.

Similarly, air molecules hit the bottom of the wing, transferring upward momentum to it.

Otherwise, all the important parts of the story are different:

No bullets hit the top of the wing.

Air pressure on top of the wing is only a few percent lower than the pressure on the bottom.

The shape of the top of the wing doesn't matter to the bullets.

The shape of the top of the wing is crucial. A spoiler at location "X" in [figure 3.10](#) could easily double the drag of the entire airplane.

The bullets don't hit each other, and even if they did, it wouldn't affect lift production.

Each air molecule collides with one or another of its neighbors 10,000,000,000 times per second. This is crucial.

Each bullet weighs a few grams.

Each nitrogen molecule weighs 0.000000000000000000000005 grams.

Bullets that pass above or below the wing are undeflected.

The wing creates a pressure field that strongly deflects even far-away bits of fluid, out to a distance of a wingspan or so in every direction.

Bullets could not possibly knock a stall-warning vane forward.

Fluid flow nicely explains how such a vane gets blown forward and upward. See [section 3.5](#).

The list goes on and on, but you get the idea. Interactions between air molecules are a big part of the story. It is a much better approximation to think of the air as a *continuous fluid* than as a bunch of bullet-like particles.

3.7 Other Fallacies

You may have heard stories that try to use the *Coanda effect* or the *teaspoon effect* to explain how wings produce lift. These stories are completely fallacious, as discussed in [section 18.4.4](#) and [section 18.4.3](#).

There are dozens of other fallacies besides. It is beyond the scope of this book to discuss them, or even to catalog them all.

3.8 Inverted Flight, Cambered vs. Symmetric Airfoils

Almost everybody has been told that an airfoil produces lift because it is curved on top and flat on the bottom. But aren't you also aware that airshow pilots routinely fly for extended periods of time upside down? Doesn't that make you suspicious that there might be something wrong with the story about curved on top and flat on the bottom?

Here is a list of things you need in an airplane intended for upside-down flight:


- You need super-duper seatbelts to keep the pilot from flopping around.
- You need to make sure the airframe is strong enough to withstand extra stress, including stress in new directions.
- You need to make sure that the fuel, engine oil, and battery acid stay where they are supposed to be.

You will notice that changing the cross-sectional shape of the wing is not on this list. Any ordinary wing flies just fine inverted. Even a wing that is flat on one side and curved on the other flies just fine inverted, as shown in [figure 3.11](#). It may look a bit peculiar, but it works.



[Figure 3.11](#): Inverted Flight

The misconception that wings must be curved on top and flat on the bottom is commonly associated with the previously-discussed misconception that the air is required to pass above and below the wing in equal amounts of time. In fact, an upside-down wing produces lift by exactly the same principle as a rightside-up wing.




[Figure 3.12](#): Airfoil Terminology

To help us discuss airfoil shapes, [figure 3.12](#) illustrates some useful terminology.

1. The *chord line* is the straight line drawn from the leading edge to the trailing edge.
2. The term *camber* in general means "bend". If you want to quantify the amount of camber, draw a curved line from the leading edge to the trailing edge, staying always halfway between the upper surface and the lower surface; this is called the *mean camber line*. The maximum difference between this and the chord line is the amount of camber. It can be expressed as a distance or (more

commonly) as a percentage of the chord length.


A *symmetric airfoil*, where the top surface is a mirror image of the bottom surface, has zero camber. The airflow and pressure patterns for such an airfoil are shown in [figure 3.13](#).



[Figure 3.13](#): Symmetric Airfoil

This figure could be considered the side view of a symmetric wing, or the top view of a rudder. Rudders are airfoils, too, and work by the same principles.

At small angles of attack, a symmetric airfoil works better than a highly cambered airfoil. Conversely, at high angles of attack, a cambered airfoil works better than the corresponding symmetric airfoil. An example of this is shown in [figure 3.14](#). The airfoil designated "631-012" is symmetric, while the airfoil designated "631-412" airfoil is cambered; otherwise the two are pretty much the same.¹¹ At any normal angle of attack (up to about 12 degrees), the two airfoils produce virtually identical amounts of lift. Beyond that point the cambered airfoil has a big advantage because it does not stall until a much higher relative angle of attack. As a consequence, its maximum coefficient of lift is much greater.



[Figure 3.14](#): Camber Fends Off The Stall

At high angles of attack, the leading edge of a cambered wing will slice into the wind at less of an angle compared to the corresponding symmetric wing. This doesn't prove anything, but it provides an intuitive feeling for why the cambered wing has more resistance to stalling.

On some airplanes, the airfoils have no camber at all, and on most of the rest the camber is barely perceptible (maybe 1 or 2 percent). One reason wings are not more cambered is that any increase would require the bottom surface to be concave --- which would be a pain to manufacture. A more profound reason is that large camber is only really beneficial near the stall, and it suffices to create lots of camber by extending the flaps when needed, i.e. for takeoff and landing.

Reverse camber is clearly a bad idea (since it causes earlier stall) so aircraft that are expected to perform well upside down (e.g. Pitts or Decathlon) have symmetric (zero-camber) airfoils.

We have seen that under ordinary conditions, the amount of lift produced by a wing depends on the angle of attack, but hardly depends at all on the amount of camber. This makes sense. In fact, the airplane would be unflyable if the coefficient of lift were determined solely by the shape of the wing. Since the amount of camber doesn't often change in flight, there would be no way to change the coefficient of lift. The airplane could only support its weight at one special airspeed, and would be unstable and uncontrollable. In reality, the pilot (and the trim system) continually regulate the amount of lift by regulating the all-important angle of attack; see [chapter 2](#) and [chapter 6](#).

[3.9 Thin Wings](#)

The wing used on the Wright brothers' first airplane is thin, highly cambered, and quite concave on the bottom. This is shown in [figure 3.15](#). There is no significant difference between the top surface and the bottom surface --- same length, same curvature. Still, the wing produces lift, using the same lift-producing principle as any other airfoil. This should further dispel the notion that wings produce lift because of a difference in length between the upper and lower surfaces.

Similar remarks apply to the sail of a sailboat. It is a very thin wing, oriented more-or-less vertically, producing sideways lift.




Figure 3.15: The Wrights' 1903 Airfoil

Even a thin flat object such as a barn door will produce lift, if the wind strikes it at an appropriate angle of attack. The airflow pattern (somewhat idealized) for a barn door (or the wing on a dime-store balsa glider) is shown in [figure 3.16](#). Once again, the lift-producing mechanism is the same.




Figure 3.16: Barn Door --- Natural Airflow

3.10 Circulation

3.10.1 Visualizing the circulation

You may be wondering whether the flow patterns shown in [figure 3.16](#) or the earlier figures are the only ones allowed by the laws of hydrodynamics. The answer is: almost, but not quite. [Figure 3.17](#) shows the barn door operating with the same angle of attack (and the same airspeed) as in [figure 3.16](#), but the airflow pattern is different.




Figure 3.17: Barn Door --- Unnatural Stream Lines




Figure 3.18: Barn Door --- Pure Circulation




Figure 3.19: Barn Door --- Natural Stream Lines


The new airflow pattern ([figure 3.17](#)) is highly symmetric. I have deleted the timing information, to make it clear that the stream lines are unchanged if you flip the figure right/left and top/bottom. The front stagnation line is a certain distance behind the leading edge; the rear stagnation line is the same distance ahead of the trailing edge. This airflow pattern produces no lift. (There will be a lot of torque --- the so-called Rayleigh torque --- but no lift.)

The difference between these figures is *circulation* --- [figure 3.16](#) has circulation while [figure 3.17](#) does not. ([Figure 3.19](#) is the same as [figure 3.16](#) without the timing information.)


To understand circulation and its effects, first imagine an airplane with barn-door wings, parked on the ramp on a day with no wind. Then imagine stirring the air with a paddle, setting up a circulatory flow pattern, flowing nose-to-tail over the top of the wing and tail-to-nose under the bottom (clockwise in this figure). This is the flow pattern for pure circulation, as shown in [figure 3.18](#). Then imagine that a headwind springs up (left to right in the figure). At each point in space, the velocity fields will add. The circulatory flow and the wind will *add* above the wing, producing high velocity and low pressure there. The circulatory flow will partially *cancel* the wind below the wing, producing low velocity and high pressure there.

If we take the noncirculatory left-to-right flow in [figure 3.17](#) and add various amounts of circulation, we can generate *all* the flow patterns consistent with the laws of hydrodynamics --- including the actual natural airflow shown in [figure 3.16](#) and [figure 3.19.12](#)


There is nothing special about barn doors; real airfoils have analogous airflow patterns, as shown in [figure 3.20](#), [figure 3.21](#), and [figure 3.22](#).



[Figure 3.20](#): Unnatural Airflow --- Angle of Attack but No Circulation



[Figure 3.21](#): Pure Circulation



[Figure 3.22](#): Normal, Natural Airflow

If you suddenly accelerate a wing from a standing start, the initial airflow pattern will be noncirculatory, as shown in [figure 3.20](#). Fortunately for us, the air absolutely hates this airflow pattern, and by the time the wing has traveled a short distance (a couple of chord-lengths or so) it develops enough circulation to produce the normal airflow pattern shown in [figure 3.22](#).

[3.10.2](#) How Much Circulation? The Kutta Condition

In real flight situations, precisely enough circulation will be established so that the rear stagnation line is right at the trailing edge, so no air needs to turn the corner there. The counterclockwise flow at the trailing edge in [figure 3.17](#) is cancelled by the clockwise flow in [figure 3.18](#). Meanwhile, at the leading edge, both [figure 3.17](#) and [figure 3.18](#) contribute clockwise flow, so the real flow pattern ([figure 3.19](#)) has lots and lots of flow around the leading edge.

The general rule --- called the *Kutta condition* --- is that the air hates to turn the corner at a sharp trailing edge. To a first approximation, the air hates to turn the corner at *any* sharp edge, because the high velocity there creates a lot of friction. For ordinary wings, that's all we need to know, because the trailing edge is the only sharp edge.

The funny thing is that if the trailing edge is sharp, an airfoil will work even if the leading edge is sharp, too. This explains why dime-store balsa-wood gliders work, even with sharp leading edges.

It is a bit of a mystery why the air hates turning a corner at the trailing edge, and doesn't mind so much turning a sharp corner at the leading edge --- but that's the way it is.¹³ This is related to the well-known fact that blowing is different from sucking. (Even though you can blow out a candle from more than a foot away, you cannot suck out a candle from more than an inch or two away.) In any case, the rule is:
The air wants to flow cleanly off the trailing edge.

As the angle of attack increases, the amount of circulation needed to meet the Kutta condition increases.

Here is a nice, direct way of demonstrating the Kutta condition: At a safe altitude, start with the airplane¹⁴ in the clean configuration in level flight, a couple of knots above the speed where the stall warning horn comes on. Maintaining constant pitch attitude and maintaining level flight, extend the flaps. The stall warning horn will come on. There is no need to stall the airplane; the warning horn itself makes the point.

This demonstration makes it clear that the flap (which is at the back of the wing) is having a big effect on the airflow around the entire wing, including the stall-warning detector (which is near the front).

Extending the flaps (while maintaining constant pitch and constant direction of flight) increases the angle of attack. This increases the circulation, which trips the stall-warning detector as described in [section 3.5](#).


3.10.3 How Much Lift? The Kutta-Zhukovsky Theorem

Here is a beautifully simple and powerful result: The lift is equal to the airspeed, times the circulation, times the density of the air, times the span of the wing. This is called the *Kutta-Zhukovsky theorem*.¹⁵

$$\text{Lift} = \text{airspeed} \times \text{circulation} \times \text{density} \times \text{span}$$

Since circulation is proportional to the coefficient of lift and to the airspeed, this new notion is consistent with our previous knowledge that the lift should be proportional to the coefficient of lift times airspeed squared.

You can look at a velocity field and visualize the circulation. In [figure 3.23](#), the right-hand edge of the blue streamers shows where the air is 70 milliseconds after passing the reference point. For comparison, the vertical black line shows where the 70 millisecond timeline would have been if the wing had been completely absent. However, this comparison is not important; you should be comparing each air parcel above the wing with the corresponding parcel below the wing.



[Figure 3.23](#): Circulation Advances Upper & Retards Lower Streamers

Because of the circulatory contribution to the velocity, the streamers above the wing are at a relatively advanced position, while the streamers below the wing are at a relatively retarded position.

If you refer back to [figure 3.7](#), you can see that circulation is proportional to angle of attack. In particular, note that when the airfoil is not producing lift there is no circulation --- the upper streamers are not advanced relative to the lower streamers.

The same thing can be seen by comparing [figure 3.20](#) to [figure 3.22](#) --- when there is no circulation the upper streamers are not advanced relative to the lower streamers.

3.10.4 Quantifying the Circulation

Circulation can be measured, according to the following procedure. Set up an imaginary loop around the wing. Go around the loop clockwise, dividing it into a large number of small segments. For each segment, multiply

the length of that segment times the speed of the air *along the direction of the loop* at that point. (If the airflow direction is opposite to the direction of the loop, the product will be negative.) Add up all the products. The total velocity-times-length will be the circulation. This is the official definition.


Interestingly, the answer is essentially independent of the size and shape of the loop.¹⁶ For instance, if you go farther away, the velocity will be lower but the loop will be longer, so the velocity-times-length will be unchanged.

3.11 Mechanically-Induced Circulation

There is a widely-held misconception that it is the velocity *relative to the skin of the wing* that produces lift. This causes no end of confusion.

Remember that the air has a well defined velocity and pressure everywhere, not just at the surface of the wing. You can go anywhere near or far from the wing with a windmill and a pressure gauge and measure the velocity and pressure. The circulatory flow set up by the wing creates low pressure in a huge region extending far above the wing. The velocity at each point determines the pressure at that point.

The circulation near a wing is normally set up by the interaction of the wind with the shape of the wing. But there are other ways of setting up circulatory flow. In [figure 3.24](#), the wings are not airfoil-shaped but paddle-shaped. By rotating the paddle-wings we can set up a circulatory airflow pattern by brute force.




[Figure 3.24](#): Paddle-Wing Airplane

Bernoulli's principle would apply point by point in the air near the wing, creating low pressure that would pull up on the wings, even though the air near the wing would have no velocity relative to the wing since it would be "stuck" between the vanes of the paddle. The Kutta-Zhukovsky theorem would apply, as stated above: lift is equal to the airspeed, times the circulation, times the density of the air, times the span of the wing.

This phenomenon --- creating the circulation needed for lift by mechanically stirring the air --- is called the *Magnus effect*.

The airplane in [figure 3.24](#) would have definite controllability problems, since the notion of angle of attack would not exist (see [chapter 2](#) and [chapter 6](#)). The concept, though, is not as ridiculous as might seem. The famous aerodynamicist Flettner once built a ship that "sailed" all the way across the Atlantic using huge rotating cylinders rather than sails to catch the wind.

Also, it is easier than you might think to demonstrate this important concept. You don't need four vanes on the rotating paddle; a single flat surface will do. A business card works fairly well. Drop the card from shoulder height, with its long axis horizontal. As you release it, give it a little bit of backspin around the long axis. It will fly surprisingly well; the lift-to-drag ratio is not enormous, but it is not zero either. The motion is depicted in [figure 3.25](#).



[Figure 3.25](#): Fluttering Card --- Lift Created by Circulation

You can improve the performance by giving the wing a finer aspect ratio (more span and/or less chord). I once took a manila folder and cut out several pieces an inch wide and 11 inches long; they work great.

As an experiment, try giving the wing the wrong direction of circulation (i.e. topspin) as you release it. What do you think will happen?

I strongly urge you to try this demonstration yourself. It will improve your intuition about the relationship of circulation and lift.




Figure 3.26: Curve Balls

We can use these ideas to understand some (but not all) of the aerodynamics of tennis balls and similar objects. As portrayed in [figure 3.26](#), if a ball is hit with a lot of backspin, the surface of the spinning ball will create the circulatory flow pattern necessary to produce lift, and it will be a “floater”. Conversely, the classic “smash” involves topspin, which produces negative lift, causing the ball to “fly” into the ground faster than it would under the influence of gravity alone. Similar words apply to leftward and rightward curve balls.

To get even close to the right answer, we must ask where the relative wind is fast or slow, relative to the center of the ball --- not relative to the rotating surface of the ball. Remember that the fluid has a velocity and a pressure everywhere, not just at the surface of the ball. Air moving past a surface is what creates *drag*, not lift. Bernoulli says that when an air parcel changes speed, it exchanges its kinetic energy (airspeed) for potential energy (pressure). For the floater, the circulatory flow created by the backspin combines with the free-stream flow created by the ball’s forward motion to create high-velocity, low-pressure air above the ball --- that is, lift.

This simple picture of mechanically-induced circulation applies best to balls that have evenly-distributed roughness. Cricket balls are in a different category, since they have a prominent equatorial seam. If you spin-stabilize the orientation of the seam, and fly the seam at an “angle of attack”, airflow over the seam causes extra turbulence which promotes attached flow on one side of the ball. See [section 18.3](#) for some discussion of attached versus separated flow. Such effects can overwhelm the mechanically-induced circulation.

To really understand flying balls or cylinders, you would need to account for the direct effect of spin on circulation, the effect of spin on separation, the effect of seams on separation, et cetera. That would go beyond the scope of this book. A wing is actually easier to understand.

3.12 Lift Requires Circulation & Vortices

3.12.1 Vortices

A *vortex* is a bunch of air circulating around itself. The axis around which the air is rotating is called a vortex line. It is mathematically impossible for a vortex line to have loose ends. A smoke ring is an example of a vortex. It closes on itself so it has no loose ends.

The circulation necessary to produce lift can be attributed to a *bound vortex* line. It binds to the wing and travels with the airplane. The question arises, what happens to this vortex line at the wingtips?

The answer is that the vortex spills off each wingtip. Each wing forms a *trailing vortex* (also called *wake vortex*) that extends for miles behind the airplane. These trailing vortices constitute the continuation of the bound vortex. See [figure 3.27](#). Far behind the airplane, possibly all the way back at the place where the plane left ground effect, the two trailing vortices join up to form an unbroken ¹⁷vortex line.




Figure 3.27: Bound Vortex, Trailing Vortices

The air rotates around the vortex line in the direction indicated in the figure. We know that the airplane, in order to support its weight, has to yank down on the air. The air that has been visited by the airplane will have a descending motion relative to the rest of the air. The trailing vortices mark the boundary of this region of descending air.

It doesn't matter whether you consider the vorticity to be the cause or the effect of the descending air --- you can't have one without the other.

Lift must equal weight times load factor, and we can't easily change the weight, or the air density, or the wingspan. Therefore, when the airplane flies at a low airspeed, it must generate lots of circulation.

* Winglets, etc.

It is a common misconception that the wingtip vortices are somehow associated with unnecessary spanwise flow, and that they can be eliminated using fences, winglets, et cetera. The reality is that the vortices are completely necessary; you cannot produce lift without producing vortices. By fiddling with the shape of the wing the designer can control *where* along the span the vortices are shed, but there is no way to get rid of the vorticity without getting rid of the lift.

Winglets encourage the vortices to be shed at the wingtips, not somewhere else along the span. This produces more lift, since it is only the amount of span *that carries circulation* that produces lift according to the Kutta-Zhukovsky theorem. Still, as a general rule, adding a pair of six-foot-tall winglets has no aerodynamic advantage compared to adding six feet of regular, horizontal wing on each side.¹⁸

The bound vortex that produces the circulation that supports the weight of the airplane should not be confused with the little vortices produced by vortex generators (to re-energize the boundary layer) as discussed in [section 18.3](#).

3.12.2 Wake Turbulence

When air traffic control (ATC) tells you "caution --- wake turbulence" they are really telling you that some previous airplane has left a wake vortex in your path. The wake vortex from a large, heavy aircraft can easily flip a small aircraft upside down.

A heavy airplane like a C5-A flying slowly is the biggest threat, because it needs lots of circulation to support all that weight at a low airspeed. You would think that a C5-A with flaps extended would be the absolute worst, but that is not quite true. The flaps do increase the circulation-producing capability of the wing, but they do not extend over the full span. Therefore a *part* of the circulation is shed where the flaps end, and another part is shed at the wingtips. If you fly into the wake of another plane, two medium-strength vortices will cause you less grief than a single full-strength vortex. Therefore, you should expect that the threat from wake vortices is greatest behind an airplane that is *heavy, slow, and unflapped*.

Like a common smoke ring, the wake vortex does not just sit there, it moves. In this case it moves downward. A common rule of thumb says they normally descend at about 500 feet per minute, but the actual rate will depend on the wingspan and coefficient of lift of the airplane that produced the vortex.

Vortices are part of the air. A vortex in a moving airmass will be carried along with the air. In fact, the reason wake vortices descend is that the right vortex is carried downward by the flow field associated with the left vortex, and the left vortex is carried downward by the flow field associated with the right vortex. Superimposed on this vertical motion, the ordinary wind blows the vortices downwind, usually more-or-less

horizontally.

When a vortex line gets close to the ground, it “sees its reflection”. That is, a vortex at height H moves as if it were being acted on by a mirror-image vortex a distance H below ground. This causes wake vortices to spread out --- the left vortex starts moving to the left, and the right vortex starts moving to the right.

* **Avoiding Wake Turbulence Problems**

If you are flying a light aircraft, avoid the airspace below and behind a large aircraft. Avoiding the area for a minute or two suffices, because a vortex that is older than that will have lost enough intensity that it is probably not a serious problem.

If you are landing on the same runway as a preceding large aircraft, you can avoid its wake vortices by flying a high, steep approach, and landing at a point well beyond the point where it landed. Remember, it doesn't produce vortices unless it is producing lift. Assuming you are landing into the wind, the wind can only help clear out the vortices for you.

If you are departing from the same runway as a preceding large aircraft, you can avoid its vortices --- in theory --- if you leave the runway at a point *well* before the point where it did, and if you make sure that your climb-out profile stays above and/or behind its. In practice, this might be hard to do, since the other aircraft might be able to climb more steeply than you can. Also, since you are presumably taking off into the wind, you need to worry that the wind might blow the other plane's vortices toward you.

A light crosswind might keep a vortex on the runway longer, by opposing its spreading motion. A less common problem is that a crosswind might blow vortices from a parallel runway onto your runway.

The technique that requires the least sophistication is to delay your takeoff a few minutes, so the vortices can spread out and be weakened by friction.

3.12.3 Induced Drag

Here are some more benefits of understanding circulation and vortices: it explains induced drag, and explains why gliders have long skinny wings. Induced drag is commonly said to be the “cost” of producing lift. But there is no law of physics that requires a definite cost. If you could take a very large amount of air and pull it downward very gently, you could support your weight at very little cost. The cost you absolutely must pay is the cost of making that trailing vortex. For every mile that the airplane flies, each wingtip makes another mile of vortex. The circulatory motion in that vortex involves nontrivial amounts of kinetic energy, and that's why you have induced drag. A long skinny wing will need less circulation than a short fat wing producing the same lift. Gliders (which need to fly slowly with minimum drag) therefore have very long skinny wings (limited only by strength; it's hard to build something long, skinny, and strong).

3.12.4 Soft-Field Takeoff

We can now understand why soft-field takeoff procedure works. When the aircraft is in ground effect, it “sees its reflection” in the ground. If you are flying 10 feet above the ground, the effect is the same as having a mirror-image aircraft flying 10 feet below the ground. Its wingtip vortices spin in the opposite direction and largely cancel your wingtip vortices --- greatly reducing induced drag.

As discussed in [section 13.4](#), in a soft-field takeoff, you leave the ground at a very low airspeed, and then fly in ground effect for a while. There will be no wheel friction (or damage) because the wheels are not touching the ground. There will be very little induced drag because of the ground effect, and there will be very little parasite drag because you are going slowly. The airplane will accelerate like crazy. When you reach normal flying speed, you raise the nose and fly away.

3.13 Frost on the Wings

The Federal Aviation Regulations prohibit takeoff when there is frost adhering to the wings or control surfaces, unless it is polished smooth.


It is interesting that they do not require it to be entirely removed, just polished smooth.

There are very good aerodynamic reasons for this rule:

- The most obvious effect of roughness on the wings is to create a lot more drag, as seen in the right panel in [figure 3.28](#), which shows wind-tunnel data for a real airfoil (the NACA 631-412 airfoil; see

[reference 5](#)). At cruise angle of attack, the drag is approximately doubled; at higher angles of attack (corresponding to lower airspeeds) it is even worse.

- The less obvious (yet more critical) problem is that roughness causes the wing to stall at a considerably lower angle of attack, lower coefficient of lift, and higher airspeed. This can be seen in the left panel of [figure 3.28](#). The pilot of the frosty airplane could get a very nasty surprise.



[Figure 3.28](#): Roughness Degrades Wing Performance

As mentioned in [section 3.4](#), Bernoulli's principle cannot be trusted when energy is being removed from the system by friction. Frost, by sticking up into the breeze, is very effective in removing energy from the system. This tends to de-energize the boundary layer, leading to separation which produces the stall.[19](#)

It is interesting that at moderate and low angles of attack (cruise airspeed and above) the frost has hardly any effect on the coefficient of lift. This reinforces the point made in [section 3.11](#) that the velocity of the air *right at the surface*, relative to the surface, is not what produces the lift.

Sometimes the air temperature is just above freezing, but due to history or due to radiative cooling, the skin of the airplane is much colder and covered with frost. A jug of warm water works wonders.

[3.14](#) Consistent (Not Cumulative) Laws of Physics

We have seen that several physical principles are involved in producing lift. Each of the following statements is correct as far as it goes:

- The wing produces lift "because" it is flying at an angle of attack.
- The wing produces lift "because" of circulation.
- The wing produces lift "because" of Bernoulli's principle.
- The wing produces lift "because" of Newton's law of action and reaction.

We now examine the relationship between these physical principles. Do we get a little bit of lift because of Bernoulli, and a little bit more because of Newton? No, the laws of physics are not cumulative in this way.

There is only one lift-producing process. Each of the explanations itemized above concentrates on a different aspect of this one process. The wing produces circulation in proportion to its angle of attack (and its airspeed). This circulation means the air above the wing is moving faster. This in turn produces low pressure in accordance with Bernoulli's principle. The low pressure pulls up on the wing and pulls down on the air in accordance with all of Newton's laws.

[3.15](#) Momentum in the Air

For an airplane in steady flight, the forces must balance. We know from the Newton's third law[20](#) that for every force there must be an equal and opposite force *somewhere*, but the special idea here is that there must be an equal and opposite force *locally* to maintain equilibrium.

The earth pulls down on the airplane (by gravity). This force is balanced locally because the air pulls up on the airplane (by means of pressure near the wings). Of course the same pressure that pulls up on the airplane pulls down on the air. This air exerts a pressure on neighboring parcels of air, which act on other parcels, and so forth all the way to the earth's surface. At the earth's surface, pressure pushes up on the air and pushes down on the earth. The downward force on the earth is just enough to balance the fact that the airplane is pulling up on the earth (by gravity).

Since force is just momentum per unit time, the same process can be described by a big "closed circuit" of momentum flow. The earth transfers downward momentum to the airplane (by gravity). The airplane transfers downward momentum to the air (by pressure near the wings). The momentum is then transferred from air parcel to air parcel to air parcel. Finally the momentum is transferred back to the earth (by pressure

at the surface), completing the cycle. There is no net accumulation of momentum anywhere (in long-term steady flight).

You need to look at [figure 3.27](#) to get the whole story. If you look only at things like [figure 3.2](#), you will never understand how the momentum balance works, because that figure doesn't tell the whole story. You might be tempted to make the following erroneous argument:

- In [figure 3.2](#), there is some upward momentum ahead of the wing, and some downward momentum behind the wing.
- As the wing moves along, it carries the pattern of upwash and downwash along with it.
- Therefore the total amount of upward and downward momentum in the air is not changing as the wing moves along. No momentum is being transferred to the air. Therefore no lift is being produced. This is nonsense!

To solve this paradox, remember that [figure 3.2](#) shows only the flow associated with the bound vortex that runs along the wing, and does not show the flow associated with the trailing vortices. That is, it is only valid relatively close to the wing and relatively far from the wingtips.

Look at that figure and choose a point, say, half a chord ahead of the wing. You will see that the air has some upward momentum at that point. All points above and below that point *within the frame of the figure* also have upward momentum. But it turns out that if you go up or down from that point more than a wingspan or so, you will find that all the air has downward momentum. This downward flow is associated with the trailing vortices. Near the wing the bound vortex dominates, but if you go higher or lower the trailing vortices dominate.

In fact, if you add up all the momentum in an entire column of air, for any column ahead of the wing, you will find that the total vertical momentum is zero. The total momentum associated with the trailing vortices exactly cancels the total momentum associated with the bound vortex.

If you consider points directly ahead of the wing (not above or below), a slightly different sort of cancellation occurs. The flow associated with the trailing vortices is never enough to actually reverse the flow associated with the bound vortex; there is always *some* upwash directly ahead of the wing, no matter far ahead. But the contribution associated with the trailing vortices greatly reduces the magnitude, so the upwash pretty soon becomes negligible. This is why it is reasonable to speak of "undisturbed" air ahead of the airplane.

Behind the wing there is no cancellation of any kind; the downwash of the wing is only reinforced by the downward flow associated with the trailing vortices. There is plenty of downward momentum in any air column behind the wing.

This gives us a simple picture of the airplane's interaction with the air: There is downward momentum in any air column that passes through the vortex loop (which is shown in [figure 3.27](#)). There is no such momentum in any air column that is ahead of the wing, outboard of the trailing vortices, or aft of the starting vortex.

So now we can understand the momentum balance: (1) As the airplane flies along minute by minute, it imparts more and more downward momentum to the air, by enlarging the region of downward-moving air behind it. (2) The air imparts downward momentum to the earth. (3) The gravitational interaction between earth and airplane completes the circuit.

3.16 Summary: How a Wing Produces Lift

- The flow pattern created by a wing is the sum of the obstacle effect (which is significant only very near the wing, and is the same whether or not the wing is producing lift) plus the circulation effect (which extends for huge distances above and below the wing, and is proportional to the amount of lift, other things being equal).
- A wing is very effective at changing the speed of the air. The air above is speeded up relative to the corresponding air below. Each air parcel gets a temporary change in speed and a permanent offset in position.
- Bernoulli's principle asserts that a given parcel of air has high velocity when it has low pressure, and vice versa.

- Below-atmospheric pressure above the wing is much more pronounced than above-atmospheric pressure below the wing.
- There is significant upwash ahead of the wing and even more downwash behind the wing.
- The front stagnation line is well below and behind the leading edge.
- The rear stagnation line is at or very near the trailing edge. The Kutta condition says the air wants to flow cleanly off the sharp trailing edge. This determines the amount of circulation.
- An airfoil does *not* have to be curved on top and/or flat on the bottom in order to work. A rounded leading edge is a good idea, but even a barn door will fly.
- Air passing above and below the wing does *not* do so in equal time. When lift is being produced, every air parcel passing above the wing arrives substantially *early* (compared to corresponding parcel below the wing) even though it has a longer path.
- Most of the air above the wing arrives early in absolute terms (compared to undisturbed air), but this is not important, and the exceptions are doubly unimportant.
- Lift is equal to circulation, times airspeed, times density, times wingspan.
- Well below the stalling angle of attack, the coefficient of lift is proportional to the angle of attack; the circulation is proportional to the coefficient of lift times the airspeed.
- Air is a fluid, not a bunch of bullets. The fluid has pressure and velocity everywhere, not just where it meets the surface of the wing.
- There is downward momentum in any air column behind the wing. There is zero momentum in any air column ahead of the wing, outboard of the trailing vortices, or aft of the starting vortex.
- Vortex lines cannot have loose ends; therefore you cannot produce lift without producing wake vortices.
- Induced drag arises when you have low speed and/or short span, because you are visiting a small amount of air and yanking it down violently, producing strong wake vortices. In contrast there is very little induced drag when you have high speed and/or long span, because you are visiting a large amount of air, pulling it down gently, producing weak wake vortices.

1

These simulations are based on a number of assumptions, including that the viscosity is small (but not zero), the airspeed is small compared to the speed of sound, the airflow is not significantly turbulent, no fluid can flow through the surface of the wing, and the points of interest are close to the wing and not too close to either wingtip.

2

To be more precise: there is no wind in either of the two dimensions that show up in [figure 3.3](#). There might be some flow in the third dimension (i.e. spanwise along the stagnation line) but that isn't relevant to the present discussion.

3

... although for turbulent flow, the stream lines can get so tangled that they lose any useful meaning.

4

This was defined in [section 2.12](#); see also [section 3.4](#).

5

By Bernoulli's principle, the slowest air has the highest pressure. At the stagnation lines, the air is stopped --- which as slow as it can get! See [section 3.4](#), especially [figure 3.8](#).

6

This low pressure is associated with fast-moving air in this region. You may be wondering why some of this fast-moving air arrives at the trailing edge late. The answer is that it spent a lot of time

hanging around near the leading-edge stagnation line, moving much slower than the ambient air. Then as it passes the wing, it moves faster than ambient, but not faster enough to make up for the lost time.

Of course, if there were no atmospheric pressure below the wing, there would be no way to have reduced pressure above the wing. Fundamentally, atmospheric pressure below the wing is responsible for supporting the weight of the airplane. The point is that pressure changes above the wing are more pronounced than the pressure changes below the wing.

Newton's laws are discussed in [section 19.1](#).

This is a first-order equation, valid whenever the pressure changes are a small percentage of the total atmospheric pressure. See [section 3.4.3](#) for more on this.

... but not always. See [section 18.4](#) for a counterexample.

The airfoil designations aren't just serial numbers; the digits actually contain information about the shape of the airfoil. For details see [reference 5](#).

We are still assuming negligible viscosity, small percentage pressure changes, no turbulence in the fluid, no fluid flowing through the surface of the wing, and a few other reasonable assumptions.

Actually, you never get 100% of the circulation predicted by the Kutta condition, especially for crummy airfoils like barn doors. For nice airfoils with a rounded leading edge, you get something like 99% of the Kutta circulation.

Choose an airplane where the stall warning indicator is on the flapped section of the wing. This includes the Cessna C-152 and C-172, but not the C-182. It includes most Mooneys and the Grumman Tiger, but excludes Piper Cherokees and the Beech Bonanza.

The second author's name is properly spelled Жуковский. When Russian scientists write this name in English, they almost always spell it Zhukovsky ... which is the spelling used in this book. Not coincidentally, that conforms to standard transliteration rules and is a reasonable guide to the pronunciation. Other plausible transliterations include Zhukovskii, Zhukovskiy, and perhaps Zhukovskij. The French prefer Joukowski, which was also popular in the US at one time even though it wasn't very logical. In any case, Joukowski is the same person as Zhukovsky. He is sometimes called "The Father of Russian Aviation".

This assumes that the loop is big enough to include the places where circulation is being produced (i.e. the wing and the boundary layer).

There is a rule that says vortex lines can *never* have loose ends. They form closed loops, like magnetic field lines. This is not a mere law of physics; it is a mathematical identity.

However, the winglet solution may provide a practical advantage when taxiing and parking. This is why Boeing put winglets (instead of additional span) on the 747-400 --- they wanted to be able to park in a standard slot at the airport.

Boundary layers, separation, etc. are discussed in more detail in [section 18.3](#).

See [section 19.1](#) for a discussion of the laws of motion.

4 Lift, Thrust, Weight, and Drag

It is better to be on the ground wishing you were flying, rather than up in the air wishing you were on the ground.

— Aviation proverb.

4.1 Definitions

The main purpose of this chapter is to clarify the concepts of *lift*, *drag*, *thrust*, and *weight*. Pilot books call these the *four forces*.


It is not necessary for pilots to have a super-precise understanding of the four forces. The concept of *energy* (discussed in [chapter 1](#)) is considerably more important. In the cockpit (especially in critical situations like final approach) I think about the energy budget a lot, and think about forces hardly at all. Still, there are a few situations that can be usefully discussed in terms of forces, so we might as well learn the terminology.

The relative wind acting on the airplane produces a certain amount of force which is called (unsurprisingly) the *total aerodynamic force*. This force can be resolved into components, called lift and drag.


- *Lift* is the component of aerodynamic force perpendicular to the relative wind.
- *Drag* is the component of aerodynamic force parallel to the relative wind.
- *Weight* is the force directed downward from the center of mass of the airplane towards the center of the earth. It is proportional to the mass of the airplane times the strength of the gravitational field.
- *Thrust* is the force produced by the engine. It is directed forward along the axis of the engine (which is usually more or less parallel to the long axis of the airplane).

These are the official definitions.


[Figure 4.1](#) shows the orientation of the four forces when the airplane in "slow flight" — descending with a nose-high attitude, with the engine producing some power. Similarly, [figure 4.2](#) shows the four forces when airplane in a high-speed descent. The angle of attack is much lower, which is consistent with the higher airspeed. Finally, [figure 4.3](#) shows the four forces when the airplane is in a climb. The angle of attack, the lift, and the drag have the same magnitude as in [figure 4.2](#).



[Figure 4.1](#): The Four Forces — Low Speed Descent



[Figure 4.2](#): The Four Forces — High Speed Dive



[Figure 4.3](#): The Four Forces — Climb

Note that the four forces are defined with respect to three different coordinate systems: lift and drag are defined relative to the wind, gravity is defined relative to the earth, and thrust is defined relative to the orientation of the airplane. In level flight these coordinate systems usually don't differ too much, but in [figure 4.1](#) you can see the difference.

This situation seems a little complicated, and it is: for instance, thrust, lift and drag all have vertical components that combine to oppose the weight; similarly the thrust and lift both have forward horizontal components.


In ordinary cruising flight, the situation is simpler. When all three coordinate systems coincide, lift must balance weight and thrust must balance drag.

[4.2](#) Balance of Forces

[Figure 4.4](#) proves that it isn't safe to assume that lift always matches weight, or thrust exactly matches drag. A bomb falling straight down has no lift and no thrust; when it reaches terminal velocity its weight is supported purely by drag. Another interesting case is a moon lander ([figure 4.5](#)) hovering on its rocket plume — it has no lift and no drag; its weight is supported by its thrust.



[Figure 4.4](#): Bomb Has Weight = Drag



[Figure 4.5](#): Moon Lander Has Weight = Thrust

You may think lift, thrust weight, and drag are defined in a crazy way, but the definitions aren't going to change anytime soon. They have too much history behind them, and they actually have advantages when analyzing complex situations.

The good news is that these subtleties are really quite unimportant. First of all, the angles in [figure 4.1](#) are greatly exaggerated. In ordinary transportation (as opposed to aerobatics), even in climbs and descents, deck angles are generally quite small, so thrust is always *nearly* horizontal. Also, the relative wind differs from horizontal by only a few degrees, so drag is always *nearly* horizontal, and lift is *nearly* vertical except in turns.

The most important situation where you have to worry about the forces on the airplane is during a turn. In a steeply-banked turn, the lift vector is inclined quite a bit to the left or right of vertical. In order to support the

weight of the airplane *and* pull the airplane around the turn, the lift must be significantly greater than the weight. This leads us to the notion of *load factor*, which is discussed in [section 6.2.3](#).

The bottom line is that thrust is usually nearly equal (and opposite) to drag, and lift is usually nearly equal (and opposite) to weight times load factor.

If we don't like the technical definitions of lift, drag, thrust and weight, we are free to use other terms. In particular, we can make the following sweeping statement: in unaccelerated flight, the upward forces balance the downward forces, and the forward forces balance the rearward forces. This statement is true whether or not we calculate separately the contributions of lift, drag, thrust and weight.

Before going on, let me mention a couple of petty paradoxes. (1) In a low-speed, high-power climb, lift is less than weight — because thrust is supporting part of the weight. It sounds crazy to say that lift is less than weight during climb, but it is technically true. (2) In a low-power, high-speed descent, lift is once again less than weight — because drag is supporting part of the weight.


These paradoxes are pure technicalities, consequences of the peculiar definitions of the four forces. They have no impact on pilot technique.

There is some additional discussion of the balance of forces in [section 19.1](#).


4.3 Types of Drag

We have seen that the total force on the airplane can be divided into lift and drag. We now explore various ways of subdividing and classifying the drag.

When a force acts on a surface, it is often useful to distinguish processes that act parallel to the surface (friction along the surface) versus forces that act perpendicular to the surface (pressure against the surface).



[Figure 4.6](#): Pressure Drag




[Figure 4.7](#): Friction Drag

[Figure 4.6](#) illustrates the idea of *pressure drag*. If the tea table is moving from right to left, you can oppose its motion by putting your hand against the front vertical surface and pushing horizontally.

[Figure 4.7](#) illustrates the idea of *friction drag*. Another way to oppose the motion of the tea-table is to put your hand in the middle of the horizontal surface and use friction to create a force *along* the surface. This might not work too well if your hand is wet and slippery.

[Figure 4.8](#) shows a situation where air flowing along a surface will create lots of friction drag. There is a large area where fast-moving air is next to the non-moving surface. In contrast, there will be very little pressure drag because there is very little frontal area for anything to push against.




[Figure 4.8](#): Friction Drag Airflow Orientation

Friction drag is proportional to viscosity (roughly, the “stickiness” of the fluid). Fortunately, air has a rather low viscosity, so in most situations friction drag is small compared to pressure drag. In contrast, pressure drag does not depend very strongly on viscosity. Instead, it depends on the mass density of the air.


Friction drag and pressure drag both create a force in proportion to the area involved, and to the square of the airspeed. Part of the pressure drag that a wing produces depends on the amount of lift it is producing. This part of the drag is called *induced drag*. The rest of the drag — everything except induced drag — is called *parasite drag*.

The part of the parasite drag that is not due to friction is called *form drag*. That is because it is extremely sensitive to the detailed form and shape of airplane, as we now discuss.

A streamlined object such as the one shown in [figure 4.9](#) can have ten times less form drag than a non-streamlined object of comparable frontal area (e.g. the flat plate in [figure 4.10](#)). The peak pressure in front of the two shapes will be the same, but (1) the streamlined shape causes the air to accelerate, so the region of highest pressure is smaller, and more importantly, (2) the streamlined shape cultivates high pressure *behind* the object that pushes it forward, canceling most of the pressure drag, as shown in [figure 4.9](#). This is called *pressure recovery*.



[Figure 4.9](#): Pressure Recovery



[Figure 4.10](#): Form Drag

Any object moving through the air will have a high-pressure region in front, but a properly streamlined object will have a high-pressure region in back as well.

The flow pattern¹ near a non-streamlined object is not symmetric fore-and-aft because the stream lines *separate* from the object as they go around the sharp corners of the plate. Separation is discussed at more length in [chapter 18](#).

Streamlining is never perfect; there is always at least some net pressure drag. Induced drag also contributes to the pressure drag whenever lift is being produced (even for perfectly streamlined objects in the absence of

separation).

Except for very small objects and/or very low speeds, pressure drag is larger than friction drag (even for well-streamlined objects). The pressure drag of a non-streamlined object is much larger still. This is why on high-performance aircraft, people go to so much trouble to ensure that even the smallest things (e.g. fuel-cap handles) are perfectly aligned with the airflow.

An important exception involves the air that has to flow through the engine compartment to cool the engine. A lot of air has to flow through narrow channels. The resulting friction drag — called *cooling drag* — amounts to 30% of the total drag of some airplanes.

Unlike pressure drag, friction drag cannot possibly be canceled, even partially. Once energy is lost to friction, it is gone forever.

The various categories of drag are summarized in [figure 4.11](#). The way to reduce induced drag (while maintaining the same amount of lift) is to have a longer wingspan and/or to fly faster. The way to minimize friction drag is to minimize the total wetted area (i.e. the total area that has high-speed air flowing along it). The way to reduce form drag is to minimize separation, by making everything streamlined.



[Figure 4.11](#): Categories of Drag

4.4 Coefficients, Forces, and Power

The word "drag", by itself, usually refers to a force (the force of drag). Similarly, the word "lift", by itself, usually refers to a force. But there are other ways of looking at things.

* Coefficients

It is often convenient to write the drag force as a dimensionless number (the *coefficient of drag*) times a bunch of factors that characterize the situation:

$$\text{drag force} = \frac{1}{2} \rho V^2 \times \text{coefficient of drag} \times \text{area} \quad (4.1)$$

where ρ (the Greek letter "rho") is the density of the air, V is your true airspeed, and the relevant area is typically taken to be the wing area (excluding the surface area of the fuselage, et cetera).

Similarly, there is a *coefficient of lift*:

$$\text{lift force} = \frac{1}{2} \rho V^2 \times \text{coefficient of lift} \times \text{area} \quad (4.2)$$

We used these equations back in [section 2.12](#) to explain why the airspeed indicator is a good source of information about angle of attack.

One nice thing about these equations is that the coefficient of lift and the coefficient of drag depend on the angle of attack and not much else. If you could (by magic) hold the angle of attack constant, the coefficient of lift and the coefficient of drag would be remarkably independent of airspeed, density, temperature, or whatever.

The coefficient of lift is a ratio² that basically measures how effectively the wing turns the available dynamic pressure into useful average suction over the wing. A typical airfoil can achieve a coefficient of lift around 1.5 without flaps; even with flaps it is hard to achieve a coefficient of lift bigger than 2.5 or so. For data on real airfoils, see [figure 3.14](#) and/or [reference 5](#).


[Figure 4.12](#) shows how the various coefficients depend on angle of attack. The left side of the figure corresponds to the highest airspeeds (lowest angles of attack). Note that the coefficient-of-lift curve has been scaled down by a factor of ten to make it fit on the same graph as the other curves. Airplanes are really good at making lots of lift with little drag.

In the range corresponding to normal flight (say 10 degrees angle of attack or less) we can make the following approximations, which I will call the *basic lift/drag model*:

- the coefficient of lift is proportional to the angle of attack,

- the coefficient of induced drag is proportional to the square of the angle of attack, and
- the coefficient of parasite drag is essentially constant.

At higher angles of attack (approaching or exceeding the critical angle of attack) these approximations break down. The coefficient of parasite drag will rapidly become quite large, and the induced drag will probably be quite large also. There will be no simple proportionality relationships. The details aren't of much interest to most pilots, for the following reason: Typically you recover from a stall as soon as you notice it, so you don't spend much time in the stalled regime. If you do happen to be interested in stalled flight and spins, see [chapter 18](#).




[Figure 4.12](#): Coefficients versus Angle of Attack

In flight, we are not free to make any amount of lift we want. The lift is nearly always equal to the weight times the load factor. This leads us to rearrange the lift equation as follows:

$$\text{coefficient of lift} = (\text{weight} \times \text{load factor}) / (1/2 \rho v^2 \times \text{area}) \quad (4.3)$$

where the airspeed and load factor are more-or-less the only variables on the right-hand side of the equation. Because of the factor of airspeed squared, the airplane must fly at a very high coefficient of lift in order to support its weight at low airspeeds.


[Figure 4.13](#) plots the same four curves against airspeed. Now the left side of the plot corresponds to the lowest airspeeds (highest angles of attack).



[Figure 4.13](#): Coefficients versus Airspeed

* Forces

[Figure 4.14](#) shows the corresponding forces. We see that whereas the *coefficient* of parasite drag was more or less constant, the *force* of parasite drag increases with airspeed. If somebody says "the drag is a ... function of airspeed" you have to ask whether "drag" refers to the drag coefficient, the drag force, or (as discussed below) the drag power.




[Figure 4.14](#): Forces versus Airspeed

We can also see in the figure that the lift force curve is perfectly constant, which is reassuring, since the figure was constructed using the principle that the lift force must equal the weight of the airplane; this is how I converted angle of attack to airspeed.

The lowest point in the total drag force curve corresponds to $V_{L/D}$, and gives the best lift-to-drag ratio. Using the standard lift/drag model and a little calculus, it can be shown that this occurs right at the point where the induced drag force curve crosses the parasite drag force curve.

* Powers

[Figure 4.15](#) shows the amount of dissipation due to drag, for the various types of drag. Dissipation is a form of power, i.e. energy per unit time.



[Figure 4.15](#): Powers versus Airspeed

Dissipation is related to force by the simple rule:

$$\text{power} = \text{force} \cdot \text{velocity} \quad (4.4)$$

In this equation, we are multiplying two vectors using the *dot product* ($\cdot$),³ which means that only the velocity component in the direction of the force counts.

In the case of drag, we have specifically:

$$\text{dissipation} = \text{force of drag} \cdot \text{airspeed} \quad (4.5)$$

The lowest point in the curve for total drag power corresponds to V_Y , and gives the best rate of climb. Using the standard lift/drag model and a little calculus, it can be shown that at this speed, the occurs right at the point where the induced drag power is $3/4$ ths of the total, and the parasite drag power is $1/4$ th of the total. Actually, in the airplane represented in these figures, V_Y is so close to the stalling speed that the standard lift/drag model is starting to break down, and the 3:1 ratio is not exactly accurate.

In the case of lift, the lift force is (by its definition) perpendicular to the relative wind, so there is no such thing as dissipation due to lift. (Of course the physical process that produces lift also produces induced drag, but the part of the force properly called lift isn't the part that contributes to the power budget.)

4.5 Induced vs. Parasite Drag

There are several useful conclusions we can draw from these curves. For starters, we see that the curve of total power required to overcome dissipation has a familiar shape; it is just an upside-down version of the power curve that appears in [section 1.2.5](#) and elsewhere throughout this book.

We can also see why the distinction between induced drag and parasite drag is significant to pilots:

- In the mushing regime, most of the drag is induced drag. As you go slower and slower, induced drag increases dramatically and parasite drag becomes almost negligible.
- At high airspeeds, parasite drag is dominant and induced drag becomes almost negligible.

In the high-speed regime (which includes normal cruise), the power required increases rapidly with increasing airspeed. Eventually it grows almost like the *cube* of the airspeed. The reason is easy to see: parasite drag is the dominant contribution to the coefficient of drag in this regime, and is more-or-less independent of airspeed.⁴ We pick up two factors of V from [equation 4.2](#) and one from [equation 4.4](#). Knowing this cube law is useful for figuring out the shape of your airplane's power curve ([section 7.6.2](#)), and for figuring out how big an engine you need as a function of speed ([section 7.6.4](#)) and altitude ([section 7.6.5](#)).

¹

[Figure 4.10](#) is not as precise as the other airflow diagrams in the book. My flow software is not capable of properly modeling the wake of the flat plate, so I had to take some liberties.

²

It is a dimensionless number, not measured in pounds or seconds or anything, just a pure number.

³

Contrast this with the cross product used in [section 19.7](#).

⁴

Induced drag decreases as the airspeed increases, but this is a relatively minor contribution in this regime.

5 Vertical Damping, Roll Damping, and Stalls

The really good pilots use their superior judgment to keep them out of situations where they might be required to demonstrate their superior skill.

5.1 Introduction and Overview

The purpose of this chapter is to examine how the airplane responds to pure vertical motions and to pure rolling motions. We will see that (except near the stall) the airplane vigorously resists such motions.

For a non-aerodynamic object like a pompom, if you wave it through the air, it will resist the motion, due to ordinary air friction. An airplane has friction, too, but we will see that there is another process ("aerodynamic damping") that is enormously more powerful than friction.


This strong aerodynamic damping should not be taken for granted, since you can certainly get an airplane into situations where the damping goes to zero or becomes negative. This is why the airplane is hard to fly near the stall. We will discuss how to deal with and/or prevent such situations.

5.2 Vertical Damping

5.2.1 Origins of Vertical Damping

Normally, the airplane is in equilibrium — all forces are in balance.¹ Let's consider the vertical forces in particular, and see how the airplane *maintains* its equilibrium.

To see how the wing reacts initially² to eliminate any unbalanced vertical force, consider the scenario in [figure 5.1](#).



[Figure 5.1](#): Sudden Increase in Weight

Initially, the airplane is buzzing along in straight-and-level flight and is nicely trimmed. Vertical forces are in balance. Then we imagine there is sudden change in the weight of the airplane, relative to the lift. A sudden excess of lift over weight could happen in several ways, such as the departure of a skydiver. Conversely, a sudden excess of weight over lift could happen in at least three ways:

- The lift decreases if you lose airspeed because of a sudden windshear.
- The load on the airplane (the effective weight) increases in a steeply banked turn.
- The weight increases if an albatross flies in the window and sits on the seat beside you.

Since we are analyzing the *initial* reaction, so we will assume there has not yet been any change in the pitch attitude.

For a brief instant after the weight increase, there will be an unbalanced downward force. According to Newton's second law, this will result in a downward acceleration. This in turn means the airplane will begin to descend.

If the downward force remained unbalanced, the airplane would continue to accelerate downward. It would not just go down, it would go down faster and faster and faster. This is not what happens, for a very interesting reason. As soon as the wing picks up an appreciable downward velocity, its angle of attack will be different.




Figure 5.2: Vertical Damping

As we discussed in [section 2.2](#), angle of attack is just the angle at which the air hits the wing. In [figure 5.2](#) we see that the air hits the wing at a larger angle during the descent; the pitch attitude of the airplane has not changed, but the relative wind is coming from a new direction, ahead of and below the airplane. This increase in angle of attack normally results in an increase in coefficient of lift. The extra lift balances the new weight, and equilibrium is restored.

This phenomenon is called *vertical damping*. As we shall see, an airplane normally has very strong vertical damping, and this is crucial for normal flight.³




Figure 5.3: Vertical Damping Mechanism

The process is depicted in [figure 5.3](#). The steps involved are:

- Unbalanced downward force makes downward acceleration.
- Downward acceleration leads to downward velocity.
- Downward velocity causes increased angle of attack.
- Increased angle of attack causes increased upward force.
- This continues until upward force equals downward force. The final state is a steady descent, with no further acceleration because the forces are once again in balance.

If the extra weight were removed, the airplane would return to level flight at the original angle of attack.

This strong vertical damping is the reason why we almost always assume that lift equals weight.⁴ If the forces were out of balance, the airplane would accelerate upward or downward, and the angle of attack would change until balance was restored. In practice, balance is restored so quickly that weight is never significantly different from lift.

I reiterate that this chapter considers only the initial response of the wing alone; the longer-term response of the airplane as a whole (including the horizontal stabilizer) is discussed in [chapter 6](#).

5.2.2 Loss of Vertical Damping

Vertical damping may seem obvious — but it is not. You should not take vertical damping for granted, because it doesn't always exist. It goes away at the stall.




Figure 5.4: Loss of Vertical Damping

Let's repeat the previous experiment, but this time let's imagine that the airplane was flying at a rather low airspeed (higher angle of attack) when it picked up the added weight. This is analyzed in [figure 5.4](#); note the higher angle of attack when compared with [figure 5.3](#).

- As before, the added weight causes a downward acceleration.
- As before, this leads to downward velocity.
- As before, this causes increased angle of attack.
- Surprise! the increased angle of attack causes *no* increase in upward force, because the coefficient of lift does not increase forever as a function of angle of attack.
- Equilibrium is *not* restored. The airplane continues to accelerate, descending faster and faster....

Now we find ourselves in a really nasty situation. Even if the extra weight were removed, the airplane would continue to fly at the new angle of attack — the very high angle of attack depicted on the right-hand side of [figure 5.4](#). The airplane would continue to descend, and even accelerate downward.

5.3 The Stall

5.3.1 Definition of Stall

The situation just described is called a *stall*. A number of peculiar things happen at the stall, including a loss of vertical damping.

* Definitions:

- The *stall* occurs at the critical angle of attack.
- The *critical angle of attack* is the point where further increases in angle of attack do not result in a further increase in coefficient of lift.
- The *unstalled regime* refers to angles of attack below the critical angle of attack; the *stalled regime* refers to angles of attack beyond the critical angle of attack.

The stall occurs at a special point on the coefficient of lift curve. Not coincidentally, this corresponds to a special point on the power curve, as was indicated back in [figure 1.13](#). It is worth exploring this relationship. [Figure 5.5](#) shows two curves: the vertical speed versus airspeed, and the coefficient of lift versus angle of attack. As discussed in [section 2.12](#), there is a deep relationship between airspeed and coefficient of lift; if the coefficient of lift is small the airplane has to fly at a higher speed to support its weight.




Figure 5.5: Power Curve related to Coefficient of Lift Curve

The coefficient of lift has a maximum; therefore the airspeed has a minimum. This minimum airspeed is called the stalling speed, and is denoted V_S .

5.3.2 Flying Beyond the Stall?

We are now in a position to answer a question that used to cause a lot of confusion: can you fly ``beyond'' the stall? Some people say yes, some people say no. The answer depends on which curve you are talking about.

- Yes, it is definitely possible to fly at an angle of attack higher than the critical angle of attack. (It may require super-human skill to overcome the pathological handling characteristics in this regime, but it is possible in principle to maintain a stalled angle of attack indefinitely.)
- No, it is not possible to sustain flight at airspeeds below the stalling speed.

In [figure 5.5](#), the coefficient of lift curve does not end at the stall, it just turns downward. Similarly, the power curve does not end at the stall, it just hooks back under. The rightward hook in the latter is related to the downturn in the former.

To say it another way: in the stalled regime, the coefficient of lift decreases with increasing angle of attack, so the airspeed required to support the weight of the airplane must actually increase as the airplane becomes more and more deeply stalled.

In the stalled regime the aircraft has a high and increasing coefficient of drag.⁵ Therefore it takes a lot of power to maintain level flight in this regime. At constant power, the rate of climb decreases (or, more likely, becomes more negative) as the aircraft becomes more deeply stalled.

A typical point in the stalled regime is indicated by the black six-pointed star in [figure 5.5](#). Flight in this regime — flying beyond the stall — is very peculiar. If some disturbance gives the airplane a slight upward velocity, it will accelerate upward and become less and less stalled. Conversely, if some disturbance produces a slight downward velocity, the airplane will accelerate downward and become more and more stalled. If you lower the nose the airplane will ascend; if you raise the nose it will descend. This sort of flying is no fun at all.

However, even in the stalled regime, the wings are producing enough lift to support the weight of the airplane. Lift does not go to zero at the stall. Indeed, the coefficient of lift is maximized at the stall!

The stall is a problem not because of loss of lift, but because of loss of vertical damping. Vertical damping is very important.

In some aircraft, the stall occurs quite suddenly, because there is a rather sharp corner in the coefficient of lift curve, as shown in [figure 5.6](#). Just below the critical angle of attack, there is good vertical damping; just above the critical angle of attack there is strongly negative vertical damping.




Figure 5.6: Nasty Stalling Behavior

Most aircraft are not so nasty. For the coefficient of lift curve shown in [figure 5.5](#), the vertical damping goes away gradually as you approach the stall. The aircraft will handle about the same one degree below the critical angle of attack or one degree beyond the critical angle of attack.

We will defer until [chapter 18](#) a discussion of what causes the stall, i.e. what properties of the airflow cause the coefficient of lift curve to bend over.

5.4 Roll Damping

5.4.1 Origins of Roll Damping

In [section 5.2](#) we considered how the airplane would respond to an unbalanced force that was purely vertical. Now let's consider how it responds to an unbalanced force that causes a torque around the roll axis.⁶ For instance, imagine several large passengers suddenly got up and moved to the left side of the airplane. This scenario is depicted in [figure 5.7](#).




Figure 5.7: Roll Damping

To understand this situation, we use the same logic as in the previous section. Remember, the angle of attack is defined to be the angle between a reference pointer (welded onto the wing, shown in red in the figure) and the direction of flight through the air. Although the nose of the airplane is still moving straight ahead, the left wingtip is moving ahead *and down*, while the right wingtip is moving ahead *and up*. This means the left wingtip is operating at an increased angle of attack, while the right wingtip is operating at a reduced angle of attack.

In any normal (i.e. unstalled) situation, this difference in angle of attack results in the downgoing wingtip producing more lift than the upgoing wingtip. These forces oppose the rolling motion. We describe this situation by saying the rolling motion of the airplane is heavily damped

Consider the following contrast:

The front wheel of a bicycle is only lightly damped (assuming the bearings are in good shape and the wheel is not touching the ground). If you give the wheel a shove, it will keep going around and around for a long time.

If you keep shoving on the bicycle wheel, it will accelerate: rolling faster and faster. This is the rotational version of Newton's second law: angular acceleration is proportional to angular force (i.e. torque).

The rolling motion of an airplane is quite heavily damped (in the unstalled regime). If you give an airplane a shove (e.g. by briefly deflecting the ailerons) it will not keep rolling. The rate of roll goes away almost as soon as the roll-inducing forces go away.

If you deflect the ailerons, the airplane will accelerate around the roll axis, but only for a short time. Thereafter, if you maintain the same deflection, the wingtip-to-wingtip difference in angle of attack will generate

forces that prevent any further angular acceleration. The airplane will roll at a steady rate, proportional to the aileron deflection.

A light, single-engine airplane has so little inertia around the roll axis that you can hardly notice it. Therefore the roll rate is essentially always proportional to aileron deflection. The ailerons' job is to overcome the roll damping.

A twin-engine airplane has a lot more rotational inertia (because it has those heavy engines mounted way out on the wing, and maybe tip-tanks also). You may notice that it does not respond as quickly to aileron deflection. To initiate a roll, you have to overcome inertia; during this time the rotational acceleration is proportional to the aileron deflection. Eventually the roll rate builds up to the point where roll damping becomes effective — that is, the wingtip-to-wingtip difference in angle of attack prevents further acceleration and the final roll-rate is proportional to the aileron deflection.

5.4.2 Loss of Roll Damping

Roll damping is crucial to normal flight. It should not be taken for granted, since it goes away at or near the stall. (This is analogous to the loss of vertical damping discussed in [section 5.2.](#))




Figure 5.8: Loss of Roll Damping

As depicted in [figure 5.8](#), the damping goes away as we approach the stall:

The left side of the figure shows an airplane in a roll, at a normal airspeed. Because of the rolling motion, the left (downgoing) wingtip is flying at a higher angle of attack, which (in this regime) produces more lift, compared to the right wingtip. Large forces are generated opposing the rolling motion.

The right side shows the same thing, except that the angle of attack is much higher — at or just beyond the critical angle. Once again because of the rolling motion, the left (downgoing) wingtip is flying at a higher angle of attack, but alas this no longer means it is producing more lift. In fact, it could well be producing less lift than the right wingtip. The aerodynamic forces do not oppose the initial rolling motion, but could well amplify it.

The loss of roll damping that occurs at the stall is quite a departure from normal behavior. This is precisely how you enter a spin or a snap roll: you arrange that one wingtip is flying above the critical angle of attack while the other is flying below the critical angle of attack.

5.4.3 Schemes to Increase Roll Damping

Since an unintentional loss of roll damping (spin or snap roll) is even more obnoxious and dangerous than an unintentional loss of vertical damping (straight-ahead stall), aircraft designers go to some trouble to increase the roll damping. They make use of the following two facts:

- 1) All bits of wing contribute equally to the lift, and to the vertical damping.
- 2) Bits of wing near the root contribute less to the roll damping, while bits of wing near the tips contribute more (because of leverage).

So the trick is, we want the roots to stall first. If we set the roots at a higher angle of incidence than the tips, when the wing as a whole attains its maximum coefficient of lift, the roots will be stalled and the tips will be unstalled, and there will still be a positive amount of roll damping. At the stall, the airplane will drop its nose straight ahead, rather than dropping one wingtip. This is a very desirable handling characteristic.

This design trick (decreasing incidence along the span from root to tip) is called *washout*. The opposite notion (having less incidence at the roots than the tips) is called *washin*. Nobody would design a plane with washin.⁷

To help compensate for propeller drag effects (as discussed in [section 9.5](#)), sometimes one wing is given more washout than the other. This is called simply asymmetric incidence.⁸

Finally: Deploying the flaps has the effect of increasing the washout. That's because the flaps are only installed on the inboard sections of the wings. When they are deployed, they increase the incidence of that section, as discussed below.

5.5 The Effect of Flaps


Flaps are important. They are used during landing ([section 12.7.1](#)), takeoff ([section 13.2](#)) and other low-speed maneuvers ([section 12.10](#) and [section 17.2.8](#)).

So, the question is, what do flaps do? Well, that question has no less than five different good answers:

1. Extending the flaps lowers the stalling speed.
2. Extending the flaps increases the wing section's angle of incidence.
3. Extending the flaps effectively increases the washout, since on most planes the inboard sections have flaps while the outboard sections do not.
4. Extending the flaps increases drag. This is helpful during landing, but unhelpful during climb and cruise.
5. Extending the flaps perturbs the trim speed. This is an undesirable side effect. See below, and see also [section 12.10](#).

5.5.1 Effect on Stalling Speed

Extending the flaps gives the airfoil a shape that is more resistant to stalling. It can fly at a higher angle of attack without stalling. This is shown in [figure 5.9](#), especially in the left-hand panel thereof. More importantly, it can produce a higher coefficient of lift, perhaps as high as 2.5, whereas the same wing without flaps would stall before its coefficient of lift got higher than 1.3 or thereabouts. This higher coefficient corresponds to a lower stalling speed (assuming other things like weight are held constant).



[Figure 5.9](#): Flaps Affect Coeff. of Lift and Incidence

5.5.2 Effect on Incidence

Extending the flaps increases the incidence of the wing as a whole. You have effectively rotated the whole wing by a few degrees. Its leading edge is in the same place, but its trailing edge is lower, relative to the rest of the plane. This is shown in the right-hand panel in [figure 5.9](#). You need to account for this change in incidence, so you can judge angle of attack by looking out the window, as discussed in [section 2.4](#).

I always wince a little when I hear someone say ``when we extend the flaps it increases the lift''. Well, I hope not. I hope that lift equals weight throughout the flap-extension process. Of course lift *would* increase if you kept the same pitch attitude while increasing the incidence, but proper technique involves lowering the nose while the flaps are extending, to maintain lift equal to weight. It's also true that at some *later* time, we will reduce airspeed, and at that time we will need more *coefficient of lift* for the same amount of lift. A correct way to say it is this: extending the flaps *permits* a higher coefficient of lift. The best way to say it is quite simple: extending the flaps lowers the stalling speed.

5.5.3 Effect on Washout

Extending the flaps raises the incidence of the wing-roots relative to the rest of the wing. That is, it increases the washout.

It turns out that this increase in incidence is in some sense larger than the increase in the stalling angle of attack. This has important and rather surprising consequences. Consider the typical situation where flaps are installed only on part of each wing. When flaps are extended, the affected part of the wing is flying at a higher angle of incidence, and therefore a higher angle of attack, compared to the unflapped part of the wing. Therefore the flapped section will stall sooner!

I know this sounds paradoxical, but it is 100% true: even though the flapped section has a shape that is intrinsically more stall resistant, it will stall before the unflapped section does.

This stalling behavior is actually quite useful. As discussed in [section 5.4](#), to get good low-speed handling, we want the wing-roots to stall first. That ensures we don't run out of roll damping before we run out of vertical damping. Therefore designers typically install flaps only on the inboard part of the wings.

To see how the flapped section can be producing more lift (even though it may be operating near, or even beyond, its critical angle of attack), please refer⁹ to [figure 5.10](#). In this figure, the horizontal axis is not angle of attack, but rather angle of attack *minus* incidence. This quantity has a simple physical interpretation: it is the angle at which the relative wind hits the fuselage. It has the nice property that it doesn't depend on which part of the wing is being considered, and it doesn't change when the incidence is changing due to flap extension.




Figure 5.10: Flaps — More Stalled Yet More lift

Before the flaps are extended, the wingtip and root have the same shape and the same performance, as shown by the blue curve. When the flaps are extended, the *back* part of the airfoil rotates down. This means that the flapped section has been rotated to a higher angle of incidence. To measure the incidence, you can look at where each coefficient-of-lift curve crosses through zero. You will see that the magenta curve has been shifted to the left. Another way to see this is at the top of the figure: the zero-lift-direction (marked ZLD) of the flapped section now points more nose-up. The max-lift-direction (marked MLD) has rotated in the same direction by an even larger amount.

In the situation shown in the figure, the flapped section is flying beyond its critical angle of attack (magenta curve) while the wingtip is flying below its critical angle of attack. This corresponds to a fairly low airspeed, such as might be used for a short-field approach.

In this situation, even though the flapped section is stalled, its coefficient of lift is still quite high, indeed

higher than what the unflapped section could produce at *any* angle of attack.

5.5.4 Effect on Drag

Deploying the first notch of flaps (on most airplanes) adds relatively little drag. Deploying the last notch adds much more.

5.5.5 Effect on Trim

On most airplanes, extending the flaps tends to lower the trim speed, just as if you had dialed in some nose-up trim. You will need to dial in some nose-down trim to compensate. (On most Mooneys, extending the flaps actually raises the trim speed, and you will need to dial in some nose-up trim to compensate.)

The main contribution to the nose-up trim change is that the tail flies in the wake of the wing. The extended flaps give a more downgoing angle to the downwash, which then hits the tail. On aircraft with a high T-tail, such as a Seminole, the tail is much less affected by the downwash from the wings, and there is typically very little trim change with flap extension.

Another contribution comes from the drag. On a low-wing airplane, the extended flaps tend to drag the bottom of the plane backward, forcing the nose down. This partially cancels the previously-mentioned effect of the downwash on the tail.

Conversely, on a high-wing airplane, the drag of the flaps tends to drag the top of the airplane backwards, forcing the nose up. This adds to the previously-mentioned downwash effect. Therefore we expect high-wing aircraft to have more trim change with flap extension.

If you set up for level flight at 90 knots and gradually¹⁰ extend the flaps (leaving the power and trim controls alone), you can expect to see the following reductions in trim speed:

	Cherokee	C-152 (2200 RPM)
First notch	5 knots	10 knots
Second notch	10 knots	25 knots
Third notch	minor	5 knots
Total	15 knots	40 knots

In a C-152, extending the flaps *with the engine at low power* causes much less trim change, so in everyday operations you will not become familiar with the large changes shown in the table. However, when you start a go-around, you will have full power and full flaps, and therefore a dangerously low trim speed (something like 45 KIAS). Watch out for nasty pitch-up on go-around! The Skyhawk (C-172) and Skylane (C-182) behave about as badly the C-152. See [section 12.10](#).

5.6 Summary

- The stall occurs at the critical angle of attack, which is the point where a further increase in angle of attack does not create a further increase in coefficient of lift.
- Lift does not go to zero at the stall. In fact, the coefficient of lift reaches its *maximum* at the stall.
- Vertical damping goes to zero at the stall.
- Roll damping goes to zero at about the same point and for similar reasons. However, a well-designed airplane will maintain a little bit of roll damping even after it has lost vertical damping.
- The airplane is very ill-behaved near the stall because of the loss of vertical damping and roll damping.
- It is possible to fly at an angle of attack above the critical angle of attack, but it is not possible to maintain steady flight at an airspeed below the stalling airspeed.
- Flaps affect the airplane in five different ways: stalling speed, drag, incidence, washout, and trim.

1

See [chapter 10.1](#) for a general discussion of equilibrium, stability, damping, and related concepts.

2

This chapter concentrates on the airplane's initial reaction, taking into account just the wing. In the longer term, the airplane reacts to an increased load by pitching down and speeding up, but this occurs after and because of the effects discussed here, and because the tail gets into the act, as discussed in [chapter 6](#).

3

See [chapter 10](#) for a discussion of damping in general.

4

For more on the relationship of lift and weight, see [chapter 4](#).

5

The induced drag will be about the same as in unstalled flight at the same airspeed, but the form drag will be much increased. See [section 4.3](#) for a discussion of types of drag.

6

For simplicity, will consider pure rolling motion. More complicated motions such as Dutch roll can make a (negative) contribution to the damping budget. See [section 10.6.1](#).

7

You can experience washin by flying upside down. A plane that has washout in normal flight will effectively have washin during inverted flight. For this reason, high-performance aerobatic aircraft are often built with essentially zero washout.

8

Because of an ambiguously-worded passage in the FAA Flight Training Handbook, some people seem to have gotten the impression that the term washin was a fancy term for asymmetric incidence. It is not; no engineer (or well-informed pilot) would use the term that way. You should stick with the definitions given here.

9

In books such as [reference 5](#), you will see curves that resemble [figure 5.10](#), in that the coefficient-of-lift curve intercepts the x -axis somewhere to the left of the origin. [Figure 5.10](#) chooses the x -axis so that intercept is equal and opposite to the incidence, but in those other books they choose the x -axis differently. Their intercept is not related to the incidence except possibly by coincidence. See [section 2.14](#) for a discussion of the choices involved.

10

In the C-152, V_{FE} , the max speed for operating with flaps fully extended, is 85 knots. You can briefly pull on the yoke to get the speed below 85 before extending the first notch of flaps. After that, you won't need to pull anymore, because of the trim-speed change which is the point of this demonstration.

6 Angle of Attack Stability, Trim, and Spiral Dives

Maintain thine airspeed,
lest the ground arise and smite thee.
— Aviation proverb.

This chapter discusses how you should use the trim wheel, how the airplane responds to changes (or attempted changes) in angle of attack, and how you should recover from a spiral dive.

6.1 The Basic Stability Principle

To control pitch attitude, conventional pilot technique is to push or pull on the yoke until the airplane is doing what you want, and then to use the trim wheel to ``trim off'' the yoke forces — thereby telling the airplane to remember that the current aircraft behavior is what you prefer.


But let's look into this a little more closely. What aspect of the behavior is the trim wheel supposed to ``remember''?

- the preferred rate of climb?
- the preferred pitch attitude?
- the preferred airspeed?
- the preferred angle of attack?

The last answer is far and away the best: the airplane is trimmed for a definite angle of attack. As we shall see, knowing this has important safety implications. Trim for angle of attack!

As discussed in [section 2.12](#), the airspeed indicator is the closest thing you have to an angle-of-attack indicator in typical light aircraft; therefore at standard weight (and load factor), trimming for airspeed is almost as sensible as trimming for angle of attack.

Angle of attack stability is crucial to well-behaved flight. It can be achieved without any complicated moving parts; even a balsa-wood toy glider maintains a definite angle of attack. To see how it works, let's start by considering the forces on the teeter-totter shown in [figure 6.1](#).




[Figure 6.1](#): Balance Insensitive to Rainfall

In the top panel of the figure, we have an ordinary playground teeter-totter with two buckets of water on it. Each bucket contains a four-inch depth of water. The left bucket has half as much horizontal area, so it contains half as much volume as the right bucket. Since the smaller bucket is twice as far away from the pivot, the torque from the small bucket is just equal (and opposite) to the torque from the big bucket; all the torques cancel.¹

Now let's consider what happens if an inch of rain falls on our teeter-totter. The new situation is shown in the

bottom panel of [figure 6.1](#). In both buckets, the depth of water increases by one inch, and in both buckets this represents a 25% increase. The system remains in equilibrium.²


We now contrast this with the slightly different teeter-totter arrangement shown in [figure 6.2](#). The initial situation is shown in the top panel. This time, the small-area bucket (the one on the left) is filled to a depth of only one inch. The other bucket is filled to a depth of four inches. In order to get things in balance, the large bucket must be moved much closer to the pivot — four times closer than it was previously, and all-in-all eight times closer than the small bucket.



[Figure 6.2](#): Balance Sensitive to Rainfall

Let's consider what happens if an inch of rain falls on this new arrangement. Once again, the depth of water increases in both buckets by one inch. This still represents a 25% increase for the right-hand bucket, but it now represents a 100% increase in the left-hand bucket. The same additional depth has a disproportionate effect. The system is no longer in equilibrium; it will tilt down to the left.

You may be wondering what all this has to do with airplanes. Well, this sort of reasoning is exactly what is needed to explain the angle-of-attack stability of an airplane. The situation is shown in [figure 6.3](#).



[Figure 6.3](#): Aircraft Sensitive to Angle of Attack

In the top panel, the airplane is just cruising along in still air. The wing is flying at a normal cruise angle of attack (four degrees), while the tail is flying at a much lower angle of attack (only one degree). This is in analogy with the two buckets, one having four inches of water and the other having only one inch.³

The torques are in balance because the even though the tail is "loading" (producing much less lift than it is capable of) it is much, much farther away from the pivot point. You can check the balance mathematically: the tail has one-quarter as much coefficient of lift and one-half as much area, but it has eight times as much

lever arm — so all the torques cancel.

The bottom panel of [figure 6.3](#) shows what happens if the airplane flies into an updraft. Because of the updraft, the relative wind is no longer coming from straight ahead, but is coming from a point one degree below the forward horizon. In the first instant after the airplane enters the updraft, the pitch attitude will not have changed (it won't have had time to change) so at least for a moment both the tail and wing will be flying at an angle of attack one degree higher than previously: two degrees and five degrees, respectively. This represents a 100% increase for the tail but only a 25% increase for the wing. This creates a pitching moment. The aircraft will pitch nose-down into the updraft. The pitch-axis torque budget will return to equilibrium only when the original angle of attack has been restored.

The same logic applies to any other situation where the airplane finds itself flying at an angle of attack different from its trimmed angle of attack. Any increase or decrease in angle of attack will have a disproportionate effect on the tail. The airplane will pitch up or down until it restores its trimmed angle of attack.

Angle of attack stability results from this simple principle: ``the thing in back flies at a lower angle of attack than the thing in front''. Aircraft designers have a special word for any situation where two airfoils have different angles of incidence, namely *decalage*,<sup>4</sup> from the French word for ``shift" or ``offset".⁵ The more wing/tail decalage you have, the more vigorously the airplane will oppose any attempted deviation from its preferred angle of attack.

* Other Flying Objects Are Not Similar

This property of being trimmed for a particular angle of attack is truly remarkable. It is not shared by other so-called ``aerodynamic" objects such as darts, arrows or bombs. They can't be trimmed for any angle of attack other than zero. If you drop a bomb from a great height, it will (to an excellent approximation) wind up pointing straight down and going straight down, with a velocity essentially as large as could possibly be obtained from an object of that size and weight. In contrast, an ordinary airframe in ordinary gliding flight goes horizontally 10 (or more) feet for every foot of descent. Its airspeed is tenfold less than the terminal velocity that would be expected for an object of that size and weight, and its vertical speed is at least a hundredfold less than terminal velocity.

If you reduce the amount of drag on the bomb, it will fall faster. If you reduce the amount of drag on the airframe, it will be able to descend *slower*.

Don't let anybody tell you the tail on an airplane works ``just like" the feathers on an arrow.

6.1.1 Center of Mass Too Far Aft

Let's consider what happens to an airplane that has insufficient decalage. It is all too easy to create such a situation, by violating the aft limit of the airplane's weigh-and-balance envelope. Suppose you are hauling a bunch of husky skydivers. Suppose initially the loading is within the weight-and-balance envelope, but one by one all the jumpers wander to the very back of the cabin. As more and more weight accumulates in the back of the plane, the center of mass (center of gravity) moves aft, and you have to dial in more and more nose-down trim. The tail has to fly at a higher and higher angle of attack to support the added weight back there. Eventually you reach the point where the wing and the tail are flying at the same angle of attack — no decalage. At this point the airplane will not necessarily immediately fall out of the sky, but you'd better be careful.

The airplane will no longer have any angle of attack stability. It won't maintain its trimmed airspeed. (There are lots of things that could disturb the angle of attack, such as (a) an updraft, as depicted in [figure 6.4](#), in analogy to the previous subsection, or (b) a speed change, which would cause a loss of lift — which in turn would cause an angle of attack change as discussed in [section 5.2](#).) If you think you've got the airplane trimmed for 100 knots and 4° angle of attack, it will be equally happy to fly at 200 knots and 1° angle of attack, or 50 knots and stalling angle of attack!




Figure 6.4: Center of Mass Too Far Aft

In such a situation, you will need to keep very close watch on the angle of attack. You will need to constantly intervene to prevent the airspeed from wandering off to a dangerously high or dangerously low value — above V_{NE} or below V_S — leading to in-flight structural failure or a nasty stall. This is in marked contrast to a normal airplane with a normal amount of angle-of-attack stability which will maintain a definite angle of attack (and therefore a more-or-less constant airspeed) all by itself.

Not only is our aft-loaded airplane much more likely to stall than a normal airplane, the resulting stall will be the worst stall you've ever seen. In a normal stall, only the wing stalls; the tail keeps flying normally. The nose then drops, and the stall recovery begins automatically. Pushing on the yoke helps things along. But in our aft-loaded plane, notice that the tail is flying at just as high an angle of attack as the wing. It is perfectly possible that the tail will stall first. When this happens, the nose will pitch up! This guarantees the wings will stall shortly after the tail does. Now you've got an airplane with both the wing *and* the tailplane stalled. Pushing forward on the yoke will only make the tailplane *more* stalled. This is not a good situation.

At this point, the jumpers won't have to be asked twice to leave the plane. After they've left, you may be able to recover from the stall.

The stall is not the only thing you need to worry about with an aft-loaded airplane. You could just as easily get an airspeed excursion to a very high airspeed. That in turn could lead to structural failure.

The moral of the story: don't mess with the weight-and-balance envelope. The airplane's manufacturer did extensive analysis and testing so they could put the largest possible weight-and-balance envelope in the Pilot's Operating Handbook.

6.1.2 Center of Mass in the Middle

Now let's take another look at what happens when the center of mass is in the middle of the allowed envelope. Suppose you get another group of passengers (since the skydivers from the previous scenario are unwilling to fly with you anymore, and have taken up basket weaving instead).

Once again, suppose the center of mass starts out in the middle of the weight-and-balance envelope, as was depicted back in [figure 6.3](#). Now suppose a few of the passengers move somewhat toward the front of the cabin. The center of mass will move forward. The tail will have less weight to support. If you don't do anything, the nose will drop and the airspeed will increase. Your first impulse will be to maintain altitude and airspeed by pulling back on the yoke. If the passengers promptly returned to their original positions, you would promptly be able to release the yoke pressure. But let's imagine that they *stay* forward. Rather than hold a steady back pressure on the yoke, you will dial in some nose-up trim to relieve the pressure.

As the center of mass moves farther and farther forward, you will need to dial in more and more nose-up trim to maintain the desired angle of attack. At some point the center of mass will move ahead of the center of lift of the main wing. The tail will then need to provide a negative amount of lift in order for the torques to be in balance, as shown in [figure 6.5](#). There is nothing wrong with this; indeed most aircraft operate with negative tail lift most of the time.




Figure 6.5: Moderately Forward CM, Tail Download

The wing will have to generate enough lift to support the entire weight of the airplane, plus a little bit extra to overcome the downward force on the tail. (We have to have the forces in balance as well as the torques.) This in turn implies the wing will generate a little more induced drag, but the loss in performance is so small that you ordinarily won't notice it.

You will have lots and lots of decalage, so the airplane will have plenty of angle of attack stability. You can check this in the figure.

Some people are under the misimpression that the tail *must* fly at a negative angle of attack for the airplane to be stable. That's just not true. The real rule is just that the thing in back needs to fly at a *lower* angle of attack than the thing in front. If the angle is so much lower that it becomes negative, that is just fine, but it is not required.

The amount of stability you have depends on the angle of attack of the tail relative to the wing, not relative to zero.

6.1.3 Center of Mass, Lift, and Area

An amusing consequence of the decalage rule involves the *center of area* and *center of lift* of the airplane. To find the center of area non-mathematically, make a top-view picture of the airplane (on reasonably rigid paper). Cut away the background, leaving just the airplane itself, and see where it balances. The balance-point will be precisely the center of area.

The mathematical rule involved is a generalization of the rule you use to calculate the location of the center of mass. Various examples of the rule include:

- To locate the center of mass: total up the product of mass times distance, summing over all elements of mass in the airplane. Divide by total mass; the result is the distance from the datum to the center of mass.
- To locate the center of area: total up the product of area times distance, summing over all elements of area. Divide by total area; the result is the distance from the datum to the center of area.
- To locate the center of lift: total up the product of lift times distance, summing over everything in the airplane that produces lift. Divide by total lift; the result is the distance from the datum to the center of lift.

All distances in these calculations are measured from some arbitrarily chosen reference point, called the *datum*. (The choice of datum doesn't matter, as long as you use the same datum for all measurements.)

6.1.4 Pitch-Axis Equilibrium

In steady flight the airplane must be in equilibrium. All torques must cancel, as discussed in [section 19.7](#). There are various ways pitch-axis torques can arise; an extreme example is shown in [figure 6.6](#). The engine is

mounted high up on a pylon. (Seaplanes commonly do this.) In particular, the thrust is created some distance above where the drag is created. This means we have two forces and a lever arm — i.e. a torque.




Figure 6.6: Thrust Not Aligned With Drag Makes Torque




Figure 6.7: Weight Not Aligned With Lift Makes Torque

The obvious way to cancel this torque is to have the center of lift (of the whole airplane) slightly offset from the center of mass (of the whole airplane). This causes a pitching moment — a torque around the pitch axis — as shown in [figure 6.7](#).

The amount of torque produced by the thrust/drag misalignment will depend on the throttle setting. Specifically, when you open the throttle such a seaplane will tend to pitch down and accelerate; you will need to pull back on the yoke and/or dial in lots of nose-up trim to compensate. This is a rather undesirable handling characteristic. Airplane designers try to minimize the thrust/drag lever arm. Indeed, given a choice, it is better to put the thrust slightly below the drag, in which case opening the throttle causes the airplane to pitch up slightly and reduce its trim speed.

In all cases, the lift/weight lever arm ([figure 6.7](#)) is always very, very short compared to the thrust/drag lever arm ([figure 6.6](#)), since weight and lift are huge compared to thrust and drag.

There are other miscellaneous contributions to the pitch-axis torque budget. For one thing, any airfoil (even a barn door) produces a certain amount of torque — not just pure lift. The amount of torque grows with angle of attack, but some airfoils have the obnoxious property that the amount of torque is not strictly proportional to the amount of lift. Changing the airfoil (e.g. by extending flaps) changes the amount of torque.

The horizontal tail has a huge amount of leverage, and its coefficient of lift is adjustable over a very wide range. This means that by moving the yoke and/or trim, the pilot can move the center of lift (of the whole airplane) over a wide range. This in turn produces lots of torque to overcome the various nonidealities just mentioned.

I reiterate: the center of lift of the whole airplane is always very, very nearly aligned with the center of mass of the whole airplane. Otherwise the aircraft would not be in equilibrium.

On the other hand, because of the decalage rule, the center of *area* will always be behind the center of lift (and hence behind the center of mass). This is because the tail is ``loafing''. It is not doing its share of the lifting. The tail is a long way behind the center of mass, so it has a whole lot of leverage. It has a lot of area, out of proportion to the lift it is producing. This means the center of area will be aft of the center of lift.

To reiterate: the center of mass is significantly ahead of the center of *area*, not the center of lift.

Another misconception that is more nearly true is the notion that the center of mass *of the whole airplane* has to be ahead of the center of lift *of the wing alone*. This condition will occur if the tail is producing a negative amount of lift. As we have seen, this is possible, but not necessary.

Here's an explicit example. I've actually done the following experiment:

- I took a Cessna 172 Skyhawk and put a couple of large pilots in the front seats, with no luggage and no other passengers. That meant the center of mass was right at the front of the envelope, so the tail had to produce considerable negative lift in order to maintain equilibrium. There was lots and lots of angle of attack stability.

- I took the same Skyhawk and put a small pilot in the front seat, a large mad scientist in the back seat, and 120 pounds of luggage in the rear cargo area. That put the center of mass right at the rear of the envelope, so the tail had to produce considerable positive lift in order to maintain equilibrium. The airplane still had plenty of stability. (As far as the pilot could tell, it was just as stable as it ever was.)


The easiest way to determine whether the tail lift is positive or negative is to observe the direction of motion of the tip vortices. This can be done using yarn contraptions that dangle from the tips of the horizontal tail.

6.1.5 Canards Operate on the Same Principle

Some airplanes have the main wing in the back. They get their stability from a much smaller wing (called a *canard*) in the front. Anybody who believes that "the thing in back always has to fly at a negative angle of attack" will have a hard time understanding how this works. The thing in back is the main wing! It had better be flying at a normal, positive angle of attack.

In fact, you can build a whole sequence of planes, gradually transforming a canard configuration into a normal configuration by making the rear wing smaller and the forward wing larger. If you do it right, all of them will have positive lift from the tail, and all of them will be stable — all for the same reason.

According to the decalage rule, the thing in front must be flying at a higher angle of attack. The canard configuration is analyzed in [figure 6.8](#).



[Figure 6.8](#): Canards Operate on the Same Principle

In the top panel, the airplane is buzzing along in still air. The main wing (in the back) is operating at a normal cruise angle of attack, 4° . The canard is operating at 10° angle of attack. This gives us 6° of decalage, which should be plenty. All the forces and torques are in balance.

Then, as shown in the lower panel, the airplane flies into an updraft. The updraft affects the canard and the main wing equally, increasing both angles of attack by one degree. This represents a 25% change for the main wing, but it represents only a 10% change for the canard. The airplane will pitch nose-down, as it should. The system will return to equilibrium only when it returns to the original (trimmed) angle of attack.

In a canard-type airplane, the center of mass is clearly always ahead of the main wing, but this is not what creates stability. The center of mass has to be ahead of the center of area (including the area of the canard). The only way this can happen is if the canard produces a huge amount of lift, out of proportion to its area. The next time you see such an aircraft parked on the ramp, take a look. You will see that the canard is installed at a tremendously large angle of incidence.⁶

Since the canard must fly at a higher angle of attack than the main wing, stalling of the canard is always a concern during maneuvers that involve a high angle of attack — e.g. landing. This explains why canard airplanes tend to have high landing speeds (which in turn requires rather long runways) — to land at a nice slow speed you have to fly the main wing at the highest possible coefficient of lift. Now for stability the canard

has to fly at an even higher coefficient than that, which calls for some compromises and/or some very tricky designs.

6.1.6 Beyond Decalage

Decalage is the main issue but not the only issue affecting the airplane's angle of attack stability. The following points are mentioned only briefly, because they are of more interest to airplane designers than to pilots.

1. In maneuvers where the airplane is rotating around its pitch axis, pitch damping must be taken into account, as discussed in [section 6.1.8](#).
2. The tail flies in the propwash (to a greater extent than the wing). This reduces stability, because it reduces the steepness of the tail's lift versus angle of attack curve. Remember that stability depends on the torque due to the tail increasing more steeply than the torque due to the forward wing when an overall angle of attack change occurs. Alas, the propwash hits the tail at the same angle regardless of what the relative wind is doing, so stability is reduced.
3. The tail flies in the downwash of the wings. This reduces stability, again because it reduces the steepness of the tail's lift versus angle of attack curve. The air flowing off the back of the wing tends to flow straight off the trailing edge, regardless of the angle at which it approached the wing. Also, when the airplane's overall angle of attack changes, the aft wing can move in or out of the forward wing's wake. This changes the lift curve of the stabilizer in ways that are hard for designers to predict. Further, any change in the downwash pattern can move the angle of attack to a new equilibrium point. Therefore, on most aircraft, extending the flaps perturbs the trim speed, as discussed in [section 5.5](#).
4. In a steep turn, the trimmed angle of attack will decrease slightly, because rotations are not commutative, as discussed in [section 19.6.5](#).
5. A propeller disk in front of the airplane reduces stability, because of the way the airflow through the disk changes with angle of attack. (Pusher props increase stability.)
6. A cambered wing reduces stability. Conversely, you can make an airfoil that doesn't need a tail to be stable, if you give it enough reverse camber; flying-wing aircraft use this trick.
7. The aspect ratio of an airfoil affects the steepness of its lift versus angle of attack curve. You get more stability if you have a short fat wing and a long skinny tail.
8. Sweepback affects the lift versus angle of attack curve.
9. Ground effect changes everything. This is important because you want the airplane to be well behaved during takeoff and landing, not just during cruise.
10. As discussed in [section 6.1.7](#), designers can use springs and/or bobweights to pull the airplane slightly away from its purely aerodynamic trim point.

To reiterate: decalage is the primary means for creating angle of attack stability. The other effects mentioned in this subsection determine how much decalage will be needed.

6.1.7 Springs and Bobweights


In addition to the purely aerodynamic contributions discussed at the beginning of [section 6.1](#), some airplanes have non-aerodynamic contributions. Imagine an aircraft that is not quite in trim from a purely aerodynamic point of view, so that you must apply pressure to the yoke. Now imagine that you relieve this pressure using a spring connected to the yoke. The airplane is now in trim in an overall sense. It is trimmed approximately, but alas not exactly, for a definite angle of attack. This is because at a higher airspeed, the aerodynamic force on the yoke is larger. This force overpowers the spring, changing the angle of attack.

Designers can also use weights (called *bobweights*) to pull the airplane slightly off its aerodynamics trim point. That makes the angle of attack depend on load factor as well as airspeed. Designers generally try to design

an airplane to use aerodynamic trim alone, but sometimes adding springs and/or bobweights are the expedient way to create an acceptable ``control feel''.

6.1.8 Pitch Damping

Let's consider what happens during a maneuver where the aircraft is rotating around the pitch axis. This includes loops, phugoids (as discussed in [section 6.1.12](#)), and steep turns. Note that for any bank angle steeper than 45 degrees, a turn involves more pitch-axis rotation than yaw-axis rotation.



[Figure 6.9](#): Pitch Damping

[Figure 6.9](#) shows what the relative wind does during an upward-pitching maneuver. The angle of attack of the tail is increased relative to the angle of attack of the wing. The aerodynamic effect is similar to the effect you would get by applying some nose-down trim. That is, the airplane wants to fly at a lower angle of attack than it would in the corresponding situation without the pitching motion. This effect is the main contribution to pitch damping. It is helpful because it makes phugoid oscillations die out after a few cycles. On the other hand, it makes graveyard spirals slightly more dangerous.

6.1.9 Center of Mass Too Far Forward

Let's finish our discussion of the pitch axis torque budget by considering what happens if the center of mass is too far forward. The airplane in [figure 6.10](#) has too much weight in the forward cargo area. In order for the torques to be in balance, the tail must be flying at a tremendous negative angle of attack.



[Figure 6.10](#): Center of Mass Too Far Forward

Stability is not the problem here. The aircraft has vast amounts of decalage and will be exceedingly stable. If the situation is not too extreme, the aircraft will be flyable until the time comes to raise the nose for the landing flare. When you pull back on the yoke, the tail will stall. This is the mirror image of the usual stall: the tail stalls because its angle of attack becomes too negative. The more stalled it gets, the less (negative) lift it produces. The nose of the airplane will snap down like a mousetrap.

This can definitely happen any time you exceed the forward limit of the weight-and-balance envelope; please don't get the idea that you are OK unless you actually put anvils in the forward baggage locker.

Some aircraft have very tight restrictions on the center of mass. Beechcraft Sundowners and V-tailed Bonanzas are notable examples; a Sundowner with just two pilots and full fuel is well beyond the forward limit of the center-of-mass envelope. The correct solution to this problem is to use ballast. For the Sundowner, 50 pounds of ballast at the back of the luggage compartment typically suffices.

I once knew some people who liked to fly a Sundowner but didn't like to bother with the ballast; they complained that the airplane was tricky to handle in the flare and they wondered why it had to go into the shop for nose gear repairs three times in a six-month period. The airplane was destroyed in a crash so they don't have this problem anymore.

Ballast may seem low-tech, but it does the job. I recommend using jugs of water as your ballast. That way if you need ballast on the outbound leg but need full load-carrying capacity on the return leg (satisfying the balance requirement with cargo and rear-seat passengers) you can dump out the water and keep the jugs; with other forms of ballast you'd need to worry about how to replace or retrieve it.

Here's a dirty trick that might save your neck in an emergency. If you need to land an airplane that is too nose-heavy, use soft-field technique. That is, carry some engine power during the flare. The propwash over

the tail will give you a little more control authority and delay the tailplane stall. On the other hand, if you are smart enough to anticipate this situation, you ought to be smart enough to load the airplane properly so that the situation doesn't arise.

6.1.10 Other Failure Modes

There are lots of ways to violate the weight-and-balance envelope. As discussed above,

- If the center of mass is too far aft, the airplane will lose its angle of attack stability. The airplane will not maintain its trimmed airspeed.
- If the center of mass is too far forward, you run the risk of a negative-alpha tail stall.

Additional things that you need to worry about include:

- Compliance with your airplane's official weight-and-balance limits is required by FAR 91.9(a). The regulators take violations pretty seriously, as indeed they should.
- There is always a limit to how far you can deflect the yoke. You may run out of control travel before you reach the stall or the zero-stability point. This is not entirely good news, because it means you lose control of the pitch axis sooner than you otherwise would.
- Of course, if you put too many anvils in the baggage compartment you need to worry about structural failure of the compartment floor — in addition to whatever stability and control problems you have.
- As the center of mass moves forward, the phugoid oscillations (as discussed in the [section 6.1.12](#)) tend to become more pronounced.
- Et cetera, et cetera.

6.1.11 Practical Considerations

Airline crews are required to check the weight and balance in detail for *every* flight. In practice, general aviation pilots often pre-calculate typical cases. For instance, I know that in one of the planes I commonly fly, two pilots (of any reasonable size) and full fuel is well within the envelope, so I know I don't need to check the details.

If I am flying an unfamiliar airplane, or an unusual mission (e.g. taking three linebackers as passengers in a Skyhawk) then I will check the weight and balance very carefully.

I have a computer program that makes it quick and easy.

6.1.12 Phugoid Oscillations

As we have seen, it is a good thing for the airplane to have plenty of stability of angle of attack, and this is relatively easy to arrange.

In fact, the airplane's desire to return to its trimmed angle of attack is so strong that it generally returns too quickly, and overshoots. To say it in slightly more technical terms, airplanes essentially never have as much *damping* of the pitch axis as you would like.

You can do the experiment yourself easily enough: Trim the airplane for straight and level flight at some reasonable airspeed. Pull back on the yoke until the airplane slows down about ten knots, and then let go. The airplane will not just return to its trimmed condition (pitch attitude, airspeed, and angle of attack) but will pitch down and speed up *too much*. Of course, the airplane will shortly discover this, and will pitch up and slow down again — but will overshoot in the other direction. This is shown in [figure 6.11](#). This phenomenon is called *phugoid oscillations* (pronounced fyoo'goid). After a few cycles, the oscillations will die out⁷ and the airplane will return to the previous trimmed condition.




Figure 6.11: Phugoid Oscillations

As the center of mass moves forward, you get more and more stability, but less and less damping of the pitch axis — therefore worse phugoids.

Fortunately, the phugoid oscillation is so slow that you can easily arrest the oscillation. If at point (1) in the figure you push the nose down to level pitch attitude, the airplane will be on altitude, on airspeed, and level — and the phugoid will be over. Similarly, if at point (3) you pull the nose up to level pitch attitude, the phugoid will be over instantly. If starting at point (2) you hold level pitch attitude, the airplane will take a while to accelerate to its trim speed; you will need to maintain back pressure on the yoke until it does. Similarly, starting at point (4) you can push on the yoke until the airplane decelerates to its trim speed.

You may find this recovery procedure counterintuitive at first, so it's good to practice it a few times. See also [section 10.6.2](#) for a general discussion of how to recognize oscillations and how to respond.

You can expect a phugoid whenever the airplane's airspeed or pitch attitude is disturbed from the trimmed equilibrium condition. Rough handling of the controls will do it for sure. Even if you leave the controls alone, a series of updrafts and downdrafts can easily initiate a phugoid (if you are not paying enough attention to the pitch attitude). This will result in much larger altitude and airspeed excursions than would have occurred if level pitch attitude had been maintained.

On July 19, 1989, the #2 engine of a DC-10 disintegrated, disabling the hydraulic systems that operate the flippers.⁸ The pilots managed to fly the beast to the Sioux City airport, controlling it with just the #1 and #3 throttles — the only controls still available. Every power change provoked a few cycles of phugoid oscillation. The pilots had never been taught about phugoids; they had to figure it out on the fly (so to speak). The captain of this flight, Al Haynes, has given a number of lectures recounting the experience. A videotape exists, too — highly recommended.

If you think the word "phugoid" looks strange, you're right. The origins of the word are highly amusing. Apparently Lanchester (who was the first to analyze these oscillations) wanted to coin a fancy name, based on Greek roots. He started with the English word "flight", which is, unfortunately, a homonym. From there, he stumbled onto the Greek word for *fleeing* instead of *flying*. The same root "fugh" has come down to us in the words "fugitive" and "centrifuge". So a term that was meant to translate as "aeronautical oscillation" actually comes out as "fugitive oscillation". Oh well.

6.2 Spiral Dive

6.2.1 Which Way Is Up?

People think they know which way is up, but they don't. The semicircular canals in your inner ear will tell you which way is up for a few seconds, but after that, you don't know. Not without looking.


If you can see the horizon, that tells you which way is up. If you can see the ground below you, that tells you which way you are turning. If you have a horizon gyro and a rate-of-turn gyro, and you are skilled at interpreting them, that's fine. But suppose you are flying in clouds, or over an unlighted area on a dark, overcast evening. If you look away from the instruments, you have no idea which way is up, or which way you are turning.

Sooner or later you will get into a bank, and then the bank angle will increase rather rapidly (due to the overbanking tendency, [section 9.4](#)).

6.2.2 Overview

The result is called a *spiral dive*. It has a well-deserved nickname: *graveyard spiral*. Except for running into something, a spiral dive is almost the only way you can inadvertently destroy an airplane.⁹


This will be a good application and illustration of what we have just learned about angle of attack stability. This subsection gives a quick overview of the situation; a more detailed discussion is presented in [subsection 6.2.3](#).



[Figure 6.12](#): Forces in a Steep Turn

Imagine that you are initially trimmed for straight and level flight at, say, 100 knots. Then you inadvertently enter a steeply banked turn. [Figure 6.12](#) shows the forces acting on the plane in level flight and in the turn. Let's imagine that the plane weighs exactly one ton. In level flight the downward force of gravity is exactly canceled by the lift produced by the wings, so the wings must be producing one ton of lift.

In the turn, though, the wings must produce enough force not only to support the weight of the airplane (vertically), but also to change the airplane's direction of motion (horizontally). The total force can be quite large: In a 60° turn, two tons of force is required. In a 75° turn, almost four tons of force is required, as shown in [figure 6.13](#).




[Figure 6.13](#): Forces in a Very Steep Turn

In order to produce 4 tons of lift, the airplane must fly at roughly 200 knots — twice the wings-level trim speed.

Now let's imagine that after spiralling for a while, you discover what is going on. The first thing you should do is to roll back to wings-level attitude. That solves your most urgent problem, but does not get you completely out of danger.

So let's think about the new situation. The airplane is still going roughly 200 knots. (It is going to slow down, but it hasn't yet done so.) It is still trimmed for cruise angle of attack. Therefore, the wings are still producing 4 tons of lift. You've got a real problem. Previously, you had 4 tons of lift pointing mostly horizontally, pulling you around the turn. Now you've got 4 tons of lift, pointed vertically — pulling you into a loop! This situation is illustrated in [figure 6.14](#).



[Figure 6.14](#): Rolling Out of a Spiral Dive

Note that a properly-executed aerobatic loop-de-loop involves only 4 Gs of force at the bottom. There is no doubt that an airplane rolling out of a 75° -banked spiral has enough energy to flip right over onto its back —

unless you do something. This is the phugoid from hell.


Therefore, when rolling out of a steeply-banked turn, plan on *pushing* on the yoke, to prevent a dangerously nose-high attitude.

6.2.3 General Discussion

Spiral dives are really important. Now that we've learned the ``lay of the land'', let's go through the scenario again in a little more detail.

The first step in the scenario is to have one wing down. There are lots of ways this could happen. If the airplane is not in good lateral trim (perhaps because you've burned more fuel from one tank than another, or because the passengers and cargo are not symmetrically distributed) one wing might drop as soon as you let go of the controls. Even if the trim is perfect, turbulence certainly can make one wing go down.

Having a wing down produces a whole series of consequences.¹⁰ The earliest step in this series is shown in [figure 6.15](#). The airplane has just entered a bank. The airspeed is the same as it was before the bank began, simply because it has not yet had time to change.




[Figure 6.15](#): Earliest Consequence of Bank

We see that the vertical component of lift is insufficient to balance the weight of the airplane. (Compare this figure with [figure 6.12](#) or [figure 6.13](#).) The unbalanced force will cause the airplane to drop straight down. This is the same as the ``albatross effect'' discussed in [section 5.2](#). Then, as soon as any appreciable downward velocity develops, the airplane will pitch down and speed up — because the airplane wants to maintain its trimmed angle of attack¹¹ as discussed in [section 6.1](#).


The combined effect of vertical damping and angle of attack stability will cause the airplane to speed up until the lift vector is long enough that its vertical component balances the weight of the airplane, as depicted in [figure 6.12](#). The load factor is defined to be the ratio of the lift the wing is actually producing, relative to the lift required for unaccelerated flight.

To say it yet another way, the load factor specifies how many *G* you pull in a steady turn. It grows explosively at large bank angles, as shown in [figure 6.16](#).



[Figure 6.16](#): Load Factor versus Bank Angle

The trim speed increases almost as dramatically, as shown in [figure 6.17](#). In a 60° bank, the airplane will want to maintain a speed that is roughly 141% of its wings-level trim speed. In a 75° bank, the trim speed is roughly 200% of the wings-level trim speed. In every airplane I know of, if you start out at cruise and then double the airspeed, you will be well beyond V_{NE} (never-exceed airspeed). This creates the risk of immediate structural failure, especially if you do something foolish like pull back on the yoke.



[Figure 6.17](#): Trim Speed versus Bank Angle

The trim speed grows in proportion to the square root of the load factor. There is a simple reason for this. Recall (from e.g. [section 4.4](#)) the key formula:

$$\text{lift} = \frac{1}{2} \rho v^2 \times \text{coefficient of lift} \times \text{wing area} \quad (6.1)$$

When you enter a spiral dive, the wing area of the airplane doesn't change, the density of the air (ρ) doesn't change, and the coefficient of lift [12](#) doesn't change much, either.

Consider the following scenario: imagine you are not proficient at instrument flying, but you find yourself flying through clouds or flying on a dark night over the desert. You will very soon lose track of which way is up.

At some point you perceive that something is wrong, because you are being pushed into your seat by unusual G loads. Four G s will definitely get your attention. You should also be able to hear the unusual wind noises, as the airplane accelerates to roughly 200% of its normal cruise speed. You will *not* have any sensation that you are turning. Even if you suspect you are in a turn, you will not be able to tell which direction you are turning, without referring to outside references or gyroscopic instruments.

Because of the overbanking tendency, the bank angle will continue to increase. The airspeed, descent rate, and load factor will increase accordingly. There will be no significant slip angle.

[6.2.4](#) Recovering From a Spiral Dive

If you find yourself in an unusual turning, descending situation, the first thing to do is decide whether you are in a spiral dive or in a spin. In a spiral dive, the airspeed will be high and increasing; in a spin the airspeed will be low. Also, the rate of rotation in a spiral is much less; the high speed means the airplane has lots of momentum and can't turn on a dime.

Here is the correct procedure for recovering from a spiral dive. [13](#)

- Smoothly roll the wings level.
- Simultaneously use your other hand to retard the throttle. This is not an essential step, but it might prevent the engine RPM from going beyond redline.
- Do not pull on the yoke at all. When you finish rolling out of the turn, the airplane will have 15 degrees or so of nose down pitch attitude, but it will immediately pitch up all by itself, at the rate of roughly 15 degrees per second. You should just *wait* a second or so until the airplane returns to level flight attitude, and then *push* on the yoke to maintain a reasonable attitude.
- As the airspeed returns to normal, the amount of pushing you need goes to zero. You can re-open the throttle and fly away.

If you have good outside references, by all means use them to re-establish wings-level attitude and then to re-establish a reasonable pitch attitude.

If you don't have good outside references, you should *not* rely on the attitude indicator (artificial horizon). The attitude indicator contains a gyro mounted on ordinary mortal gimbals, which can only accommodate a limited range of pitch and bank angles. A steep spiral can easily cause the gyro to tumble, whereupon it will need several minutes of relatively straight and level flying before it can re-erect itself. Military aircraft have non-tumbling attitude indicators, but you're not likely to find such things in a rented Skyhawk. Therefore, you should roll the wings level by reference to the rate-of-turn gyro. [14](#) Being a rate gyro (as opposed to a free gyro) it has no gimbals, and therefore can't possibly suffer from gimbal lock.

Remember: to recover from an unusual attitude, use the rate-of-turn gyro to level the wings. [15](#) This is a good

example of the sort of information you have to get from books. Presumably during training you will never do anything bad enough to tumble the attitude indicator.

Controlling the pitch attitude without relying on the artificial horizon requires thoughtful use of the airspeed indicator. At the point where the wings have just been returned to level, the airspeed will be something like twice what it ought to be. It will decrease slowly at first, then faster and faster. Your job is to keep the airspeed from unwinding too quickly. Pick some rate like 5 knots per second, and push on the yoke enough to keep the airspeed needle from moving faster than that.

Don't worry that pushing on the yoke will cause the airplane to fly into the ground. The airplane will climb and it will pitch up all by itself; your job is to keep it from pitching up too much. Remember¹⁶ the law of the roller coaster: 9 feet per knot, per hundred knots ([section 1.2.1](#)). As you slow down from the high-speed dive, most of that airspeed energy will be converted back to altitude. Wait until the airspeed returns to a reasonable value before you worry about returning to your exact intended altitude.

You can also use the altimeter to help manage the pitch attitude. As soon as the altimeter needle starts moving upward, you should push on the yoke to keep the needle from moving too quickly.

Unless you know you are proficient on instruments, you should not rely too heavily on the vertical speed indicator. It has weird built-in delays that can be hard to interpret.

Very, very few pilots have been taught how to handle a spiral dive correctly. In fact, the official FAA **Flight Training Handbook** ([reference 21](#)) currently¹⁷ says not once but four times (on pages 65, 187, and 188) that the procedure for recovering from a spiral dive calls for pulling *back* on the yoke. It says that you shouldn't pull back too much, but nowhere does it say that you might need to push forward. The FAA **Instrument Flying Handbook** also discusses spiral dives without giving the slightest hint that forward pressure might be necessary. The vast majority of other pilot training books suggest the same wrong procedure.

In some aircraft, including many trainers, retarding the throttle produces a nose-down pitch change which helps with the recovery, just like a small push on the yoke. Although this helps, it is definitely not sufficient in all cases. What's worse, there are some aircraft (as mentioned in [section 6.1.4](#)) in which retarding the throttle produces a nose-up pitch change.

In a not-very-steep spiral, it hardly matters what recovery procedure you use, but the more serious the spiral, the more crucial it is to use the correct procedure.

Let's look again at what happens if you use the wrong procedure. You are buzzing along in the clouds, and you get into a spiral dive. You smoothly roll the wings level (so far so good). The next thing you know the plane pitches up into a ridiculous nose-high attitude. If we are talking about a really high-speed spiral dive, the airplane will loop right over on its back. If the spiral was more moderate, you will ``only" do a tail slide or hammerhead or something.

This is just about the last thing you need. You were in a spiral dive, which was bad enough — but now you are in some horrendous unusual attitude, stalled and/or upside down, still in the clouds.

If you use the correct procedure, recovering from the spiral dive is straightforward. If you use the wrong procedure, the ensuing unusual attitude could be very hard to recover from.

If you use the widely-taught ``standard" procedure and pull back on the yoke, it can only make things worse. Pulling back will increase the angle of attack, and therefore the coefficient of lift. This might make things *much* worse, for several reasons:

- It will make you more likely to wind up in an unusual attitude.
- Since the wings were already developing 4 Gs, you don't need to increase that very much before you snap the wing spar.
- Even if you don't break the airplane, you might break the pilot. It varies a lot from person to person, but 6 Gs, especially suddenly and unexpectedly, is enough to drain the blood from your head and collapse some of the blood vessels in your brain. Even if the G load is removed, it will take a while for you to regain consciousness. Furthermore, even after you regain consciousness you will not be as smart as you used to be. Your thought processes will be severely impaired for a couple of minutes if not longer. Since you were already in an emergency situation *before* you blacked out, this is really the

last thing you need.

The correct recovery procedure is rather counter-intuitive. Because the airplane is descending and because it is going too fast, your instincts will tempt you to raise the nose. The problem is that the airplane's instincts tell it to do the same thing — and it will pitch up *too much* unless you intervene.

6.2.5 Try It Yourself

You don't need to take my word for what happens — you can go out and do some experiments yourself. You probably want to take an instructor along, but it is not absolutely necessary if you are careful. Experiment with shallow banks before messing with really steep banks.

Start by trimming the airplane for level flight at a low-cruise airspeed, say 100 knots, and clearing the area. Roll the airplane into a 45° bank and let it descend and accelerate. Leave the throttle alone, leave the trim alone, and don't push or pull on the yoke. Apply enough aileron to keep the bank from getting steeper than 45° . Wait a few seconds for the airspeed and descent rate to stabilize, then roll the wings level and watch what happens.

After you know what happens in a 45° bank, try it again at 50° , and work your way up to 55° . Don't even think about exceeding 60° without having an aerobatic-qualified instructor on board. The margin between an "interesting" spiral and a genuine emergency becomes small, as you can see in [figure 6.16](#).

When you roll out of the 55° or 60° banked spiral, the nose will be pointed about 15 degrees below the horizon. If you count off one second, the pitch attitude will be level. After two seconds, it will be 15 degrees nose up. After three seconds, it will be 30 degrees nose up, which is an awful lot. You will be quite happy to grab the controls at that point and push the aircraft back to a reasonable pitch attitude.

Repeating the experiment under the hood is also edificational.

Normally when you demonstrate a steep turn, the airspeed does not increase — in fact it decreases. That is because you retrim and/or pull back on the yoke, causing the angle of attack to increase. In contrast, our explanation of the spiral dive assumed that it was an *inadvertent* spiral dive, so the angle of attack stays nearly the same and may even decrease¹⁸ slightly. One thing is certain: the wings have to create enough lift to support the effective weight of the airplane (real weight times load factor). If the coefficient of lift stays the same, the speed has to increase; if the speed stays the same, the coefficient of lift has to increase.

6.3 Summary

- The airplane has considerable angle of attack stability. The airplane is trimmed for a definite angle of attack.
- Speaking of pitch stability is less accurate than speaking of angle of attack stability.
- The biggest contribution to angle of attack stability is *decalage*. The thing in back flies at a lower angle of attack than the thing in front. The thing in back may, but need not, fly at a negative angle of attack. A canard obviously requires the thing in back (the main wing!) to have positive angle of attack.
- Angle of attack stability is reduced as the center of mass moves aft.
- The trim wheel is used to choose what the trimmed angle of attack will be. The yoke is a convenient extension of the trim wheel, for temporary changes in angle of attack.
- The trimmed angle of attack corresponds to a definite airspeed (assuming constant load).
- Configuration changes (flap extension, power changes etc.) have side-effects on the trim speed.
- In a turn, the load factor increases. In order to maintain its trimmed angle of attack, the airplane pitches down and speeds up. In a steep turn (assuming you don't change the angle of attack using the yoke and/or trim), the speed-up is substantial.
- When you roll out of a steep turn, the airplane tends to pitch up all by itself. You may need to push on the yoke to maintain control.

1

For a general discussion of the concepts of force, torque, and moment, see [section 19.7](#).

2

For a general discussion of the principles and terminology of equilibrium stability, and damping, see [chapter 10](#).

3

We find it convenient to use the center of mass of the airplane as our pivot point, measuring all lever-arms from that point, so the force of gravity contributes nothing to the pitch-axis torque calculations. Of course, the answers come out the same no matter what pivot is chosen. See also [section 6.1.4](#) for a discussion of sundry additional pitching moments.

4

...rhymes with ``day-garage'', with the accent on the last syllable.

5

The word can equally well refer to a difference in angle of attack between the two wings of a biplane.

6

One famous exception: the Wright brothers' original airplane (``Flyer I'') on display in the Smithsonian does not have enough decalage to produce positive angle of attack stability. It must have required a goodly amount of skill and constant attention during flight.

7

See [section 6.1.8](#) for a discussion of why it dies out.

8

... not to mention the ailerons, rudder, flaps, et cetera.

9

Other disasters such as in-flight fires are vanishingly improbable by comparison.

10

Among other things, airplane will turn, pitch down, speed up, and experience a load of more than one G. Also, once the airplane is substantially banked (more than 30 or 40 degrees) the overbanking tendency will cause the bank to get steeper and steeper.

11

The angle of attack remains constant to a good approximation, but because of the pitch damping discussed in [section 6.1.8](#) and other nonidealities discussed in [section 6.1.6](#), it won't be *exactly* constant.

12

Remember, coefficient of lift is determined by angle of attack — and the airplane is trimmed for definite angle of attack. We are assuming an inadvertent spiral dive, so you presumably haven't changed the angle of attack by pushing or pulling on the yoke, or by messing with the trim.

13

Recovery from a spin is discussed in [section 18.7](#).

14

That is, the turn needle or turn coordinator, whichever you happen to have.

15

When the wings are level, you will observe zero rate of turn on the rate gyro. Remember that this instrument indicates rate of turn, not bank angle *per se*.

16

See [section 1.2.1](#).

17

Mid-1997. The FAA is working on a revision of the Flight Training Handbook. Let's hope these sections are the first to be changed.

18

... due to pitch damping, as discussed in [section 6.1.8](#).

7 More About Energy and Power

- Don't worry about ridiculous lies.
- Worry about plausible lies.
- Don't worry about things that confuse you.
- Worry about things that you know for sure, because not all of them are true.

7.1 Introduction

There is an age-old conundrum in the pilot community: Some people suggest that the yoke controls altitude while the throttle controls speed (just like in a car). Other people suggest just the reverse, namely, that the yoke controls airspeed while the throttle controls altitude.

So, which is correct?

Answer: neither one is correct. Both suggestions are based on wishful thinking. You might *wish* for an airplane where one control changes altitude and nothing else, while another control changes speed and nothing else, but that is not how real airplanes work.

The truth is simple enough:

- The yoke (in conjunction with trim) controls angle of attack, and hence determines airspeed. Airspeed is linked to altitude via the law of the roller-coaster and via the power curve.
- The throttle controls power. Power can be used to overcome drag, to accelerate, and/or to climb.

This is the right way to think about the issue.

I like to say ``the yoke is the main speed control, but it is not *just* the speed control". That is, if you want to change speed you simply must move the yoke and/or trim.¹ However, moving the yoke and/or trim has multiple effects: there is not just a speed change but also a short-term change in altitude because of the law of the roller-coaster, plus a long-term change in altitude because of the power curve.

Your piloting performance is sometimes judged on how well you maintain your assigned altitude and airspeed. Since you do not have a simple up/down control or a simple fast/slow control, even seemingly simple maneuvers require using *combinations* of controls. Let's look at some examples.

7.2 Making Changes in Airspeed

Once upon a time, a friend of mine bought a fancy new airplane. Although he already had lots of experience piloting complex aircraft, this was a step up in performance, so he thought it would be wise to take a week-long course at an internationally-famous training center. He also got lots of other instruction in the new plane. Even so, after dozens of hours of experience in the new plane, he still didn't feel ``in command". He kept getting into unpleasant high-workload situations. Among other things, he complained that it took forever to get the thing slowed down.


When I discussed this with him, it didn't take long to discover a couple of easily-fixable problems. For starters, he had been told to control airspeed using the throttle. He had the firm impression that to reduce speed somewhat, he should just close the throttle somewhat. I pointed out that such an idea couldn't possibly be right, for two reasons:

- Being high and fast is very different from being low and fast, so any rule of the form ``if you are fast do such-and-such" must be dangerously wrong.
- Even when the right procedure calls for closing the throttle, you don't ``just" close the throttle, for reasons that we now discuss.

This discussion assumes you want to change airspeed while maintaining straight-line flight. This includes straight and level flight, and it also includes the important case of final approach, where you are descending on a straight line, following a nice stable glideslope.

7.2.1 Front Side of the Power Curve

[Figure 7.1](#) shows the obvious but not-recommended procedure for decelerating on the front side of the power curve.




[Figure 7.1](#): Slowly Decelerating on the Front Side of the Power Curve

Since flying at lower speed requires less power, if you just reduce the power the right amount (as shown by the solid red curve), the airspeed will eventually dribble down. The power required is shown by the dashed green curve; it gradually decreases as the airspeed decreases.

The problem with this technique is that the airspeed keeps decreasing for a very long time. You will need to retrim over and over and over.

Slowing down means shedding kinetic energy. The area between the two curves² shows exactly how much kinetic energy you have shed.

[Figure 7.2](#) shows a much cleverer procedure. The idea is to make a *temporary* reduction in power that is big enough so that the airplane decelerates in a reasonable time. Then you can re-open the throttle to maintain the desired final outcome. If you do a little extra work with the throttle, you will do a lot less work with the other controls — and you will get a nicer result (faster deceleration).



[Figure 7.2](#): Cleverly Decelerating on the Front Side of the Power Curve

Remember, deceleration requires shedding kinetic energy, as indicated by the area between the curves. The area in this figure is the same as the area in the previous figure; we have just ``collected'' the area so that we can shed the energy in a reasonably short time. That means you don't have to spend the rest of your life re-trimming as the airspeed gradually changes.

7.2.2 Back Side of the Power Curve

Imagine you are on final approach. On long final you are maintaining a speed near V_Y (a normal approach speed in many aircraft), using 1700 RPM of engine power. Then, suddenly, the tower controller asks that you land and hold short of a crossing runway. You decide to convert the normal approach to a short-field approach. This requires decelerating from V_Y to a somewhat slower speed. The procedure is shown in [figure 7.3](#).

You need to shed some kinetic energy, as shown by the shaded area in the figure. Since that always takes time, you should immediately retard the throttle. You are now getting rid of mechanical energy (via drag) faster than it is being replaced (via the engine). You want to pay for this energy deficit by cashing in airspeed, not altitude, so you must pull back on the yoke and then roll in some nose-up trim to get rid of the force on the yoke. When the airspeed reaches short-field approach speed, you re-open the throttle. Returning to 1700 RPM will not suffice; you will need *more* power to complete the approach at this low speed than it would have at the higher speed.




Figure 7.3: Decelerating on the Back Side of the Power Curve

This is an interesting contrast with the previous situation (e.g. [figure 7.2](#)). The required power *increases* as the airspeed decreases. Therefore you do not even have the option of making the speed-change with only one power-change. It requires two (opposite and unequal) power changes.

7.2.3 Right versus Wrong Procedures

Another view of what is happening is shown in [figure 7.4](#). The dotted line shows the descent rate needed to remain on the glideslope as a function of airspeed. You started out at point A, using 1700 RPM. You are now at point B, using more than 1700 RPM to remain on the glideslope.




Figure 7.4: Energy Management on Approach

This combination of controls (close the throttle a little, pull the nose up, then open the throttle more than a little) is the only way to decelerate without an altitude excursion when you are in the mushing regime.

The analysis given above — thinking about the energy change in terms of the area between the two curves — is simple, practical, and absolutely correct.

In contrast if you tried to analyze this maneuver in terms of an up/down control versus a fast/slow control, it would be very confusing. Let's try it anyway.

- Suppose you think of the yoke as purely the fast/slow control and the throttle as purely the up/down control. At the moment you decide to convert to the short-field approach, your only problem appears to be excess airspeed. Therefore you pull back on the yoke. Poof! You decelerate sure enough, but you go above the glideslope in the process. You notice this, and reduce the throttle. You gradually descend back onto the glideslope. You re-open the throttle to 1700. That doesn't quite suffice, so you slowly drop *below* the glideslope. You notice this before too long and add power. Eventually find the right combination of settings. Summary: you get the job done, but it is rather sloppy. You have unnecessary altitude excursions and airspeed excursions, and you do some unnecessary work.
- In contrast, now suppose you think of the yoke as the up/down control and the throttle as the fast/slow control. At the moment you decide to decelerate, you close the throttle a little. Contrary to your wishes, the airplane does not decelerate; in fact it probably accelerates a little.³ You pull on the throttle a little more. Still no deceleration. Now the airplane is starting to descend below the glideslope. You notice this and pull back on the yoke. Now things seem (but only seem) better, since you are now back on the glideslope at a reduced airspeed.

At this point you are in real danger. You are losing energy rapidly, because you are operating on a draggy part of the power curve with a reduced throttle setting. The energy deficit must be paid by cashing in altitude and/or airspeed. Unfortunately, most pilots, especially beginners, pay more attention to altitude than to airspeed. As you lose energy you will keep pulling back on the yoke to maintain altitude. This allows you to stay on the glideslope in the short run — but at a terrible cost. You might very soon cash in *all* of your airspeed.

Let's hope that you notice the decreasing airspeed before you stall. Using the (fallacious) idea that the throttle controls airspeed, you shove open the throttle. This does not immediately cause the airplane to accelerate; in fact it probably causes a slight deceleration (which is definitely not what you need right now). It also causes you to start climbing above the glideslope. You notice this and shove

forward on the yoke.

You might eventually stumble onto the right combination of yoke and throttle, but the process won't be pretty.

Conclusion: trying to pretend that the airplane has a pure up/down control or a pure fast/slow control is a losing proposition.

The yoke works by moving certain control surfaces at the back of the airplane. Fifty years ago, Langewiesche ([reference 1](#)) named those surfaces the *flippers*. He wisely refused to call them ``elevators" lest you think that their primary effect was to ``elevate" the airplane.

The flippers primarily control airspeed,⁴ not elevation.

Of all the oversimplified wishful-thinking ideas, the notion that the yoke is the up/down control is the most deadly. You may think that neither you nor anybody else would be dumb enough to keep pulling back until the stall occurs — but the accident statistics indicate otherwise. The stall/spin accident is the #1 most-common type of fatal accident, year in and year out.

Stall/spin accidents occur during departure as well as approach. Once again, during departure the airplane is normally at or near V_Y , so the notion that the yoke is the up/down control is guaranteed to be wrong — dangerously wrong.

The problem is compounded because during approach and departure the airplane is at low altitude. At a higher altitude you would have more time to figure out the problem, and you would be able to regain vital speed by cashing in some altitude.

7.3 You Can Get Away With A Lot During Cruise

You may be wondering how such a dangerous notion could be come so widespread. The answer is simple: the notion that yoke is the up/down control appears to work, *most* of the time.

Nearly all of your pilot time is spent in normal cruising flight. Now suppose at some point you find yourself 100 feet below your desired cruising altitude. What do you do? You pull back on the yoke. This is what everybody does. It works. There's nothing wrong with it.

Here is the detailed analysis: You start out with a shortage of altitude which implies a shortage of mechanical energy. In the short term you can't change the mechanical energy, but you can convert airspeed into altitude using the law of the roller-coaster.

At this point you have returned to the desired altitude. You are still low on energy, but since the new airspeed is closer to V_Y , you are on a less-draggy part of the power curve and you will eventually make up the deficit. As the airspeed rebuilds, you gradually release your tug on the yoke. You don't need to touch the throttle during this maneuver.

There is an important assumption in this analysis that often goes unstated: Most pilots are very aware of their precise altitude, but (alas) not nearly so aware of their precise airspeed. Similarly: most flight instructors, air traffic controllers, and checkride examiners will complain immediately if you deviate from your assigned altitude, but they hardly ever seem to notice or care about airspeed excursions. This is not 100% logical, but it is a fact of life.

In this scenario, we corrected for an altitude excursion by means of an airspeed excursion. Under the circumstances, it was a perfectly reasonable thing to do.

For comparison, here's a scheme for correcting the same 100-foot altitude excursion *without* an airspeed excursion. You notice that you have an energy shortage, so you open the throttle a little. The airplane will enter a nice climb, with negligible change in airspeed. When you reach the assigned altitude, you return the throttle to its previous setting and the maneuver is complete. You leave the yoke and trim alone.⁵

This scheme might seem like the ideal way to perform the correction maneuver, but it is very rarely used in practice. There are a couple of reasons for this.

- Commonly, the purpose of the flight is to get somewhere as quickly as possible. Therefore, in cruising flight, the throttle is already as far open as it should go. When a shortage of mechanical energy

develops, increasing the engine output is not an option. The only option is to choose a less-draggy speed (closer to V_Y) while the energy rebuilds.

- If you make a temporary reduction in speed by pulling on the yoke, when you let go the airplane will return to its previously-trimmed speed. It's simple. In contrast, there is no corresponding idea of "throttle trim". If you move the throttle temporarily, it is not particularly easy to move it back to exactly the right place afterward. What's worse, you also need to worry about the mixture control and (possibly) the engine RPM control. Making a temporary change in power might require moving *three* controls (or *six* controls in a twin), and it would be an obnoxious task to get them all back to their proper positions afterward.

I repeat that the aerodynamically logical way to fly the airplane precisely is to trim for the airspeed you want and then manage the altitude with the throttle. When in doubt, do it this way. If you were a 100% logical Vulcan you might do it this way all the time. However, during cruise, it is more convenient to leave the throttle alone, use the yoke as if it were the up/down control, and accept modest airspeed excursions.

It is OK to use the yoke as if it were the up/down control provided:

- you are on the front side of the power curve, and
- you are willing to accept airspeed excursions.

The second proviso is just as important as the first. Suppose you decide to descend to a substantially lower altitude. You could do this by shoving forward on the yoke and/or dialing in lots of nose-down trim, but if you're not careful you could exceed the maximum normal-operations speed. As always, when in doubt, trim for the airspeed you want and then manage the energy situation by controlling engine power and/or controlling drag.

Again: It is OK to use the yoke as the up/down control provided:

- you are on the front side of the power curve, and
- you are willing to accept airspeed excursions.

7.4 Let "George" Do It

Sometimes I get a student who says "The yoke *has* to be the up/down control. I know because the autopilot controls altitude just by moving the yoke".

All I can say is that autopilots are not exempt from the laws of physics — the power curve and the law of the roller-coaster. The same rule applies: "George" (the autopilot) can control altitude using just the yoke provided you are on the front side of the power curve *and* you are willing to accept airspeed excursions.

This point is so important that I will analyze the short-field approach scenario one more time — using the autopilot.

Refer back to [figure 7.4](#). You start out at point A. The autopilot is using the yoke as if it were the up/down control, trying to keep you rigorously on the glideslope. When you decide to decelerate, you retard the throttle, whereupon the autopilot pulls back on the yoke, keeping the airplane on the glideslope by cashing in some airspeed. When the airspeed reaches short-field approach speed, you re-open the throttle.

You are now at point B. Things appear OK, but there are two things that could go wrong.

- Suppose a momentary updraft carries the airplane above the glideslope. The poor dumb autopilot will push forward on the yoke. This will convert the excess altitude to airspeed. The airplane will return to the glideslope, but since its new airspeed is closer to V_Y , it will tend to climb. The more it climbs, the more the autopilot will push forward on the yoke. This unstable feedback process will continue until the airplane reaches point C — the other point where the airplane can stay on the glideslope with the chosen amount of power. Since this point is on the front side of the power curve, "George" can get away with controlling the altitude using just the yoke. This is not, of course, a good short-field approach speed.

This airspeed excursion from point B to point C will probably leave you unable to complete the approach. You can go around and try again. This may be disappointing, but there is something much worse that could have happened starting from point B.

- Using the same logic, let's see what happens supposing at point B a *down*draft carries the airplane

below the glideslope. The poor dumb autopilot will pull back on the yoke. This will convert some airspeed to altitude. The airplane will return to the glideslope, but since its new airspeed is farther than ever from V_Y , you will tend to descend. The more you descend, the more the autopilot will pull back on the yoke.

This is crazy! The autopilot is performing the "flare" maneuver while you are still way out on final approach, cashing in all your airspeed in the vain attempt to maintain altitude! At this throttle setting, there is simply not enough energy to carry the airplane along the glideslope at any speed below point B , and pulling back on the yoke temporarily disguises and permanently worsens the problem.

We hope that the autopilot runs out of pull-back authority before it causes the wings to stall. In either case, the airplane is going to descend below the glideslope.

The only way out of this mess is to notice that you have an energy shortage. The sooner you open the throttle the better off you'll be.

If you want to prevent problems of this sort, don't try to control altitude using the yoke unless you're on the front side of the power curve *and* you're willing to accept airspeed excursions. The easiest way to control both airspeed and altitude is to trim for the right airspeed, leave the yoke alone, and control altitude with the throttle.

7.5 Max Performance using the Power Curve

7.5.1 Best Rate of Climb

The purpose of this section is to get a deeper understanding of the power curve, and to see how it applies to maximum-performance climbs and descents. If you aren't interested in such details, you can skip to the next section.

Let's start by comparing [figure 7.5](#) to [figure 7.6](#). As shown in [figure 7.5](#), the highest point on the power curve represents the best rate of climb. The corresponding airspeed is denoted V_Y .




Figure 7.5: Best Rate of Climb




Figure 7.6: Minimum Sink, Best Endurance Glide

Similarly, as shown in [figure 7.6](#), the highest point on the power curve is the point that causes the minimum sink rate. This gives the maximum *time* aloft. Again, the corresponding speed is denoted V_Y .⁶

Not coincidentally, V_Y marks the boundary between the "front side" and the "back side" of the power curve. As discussed in [section 1.2.5](#) and [section 1.3.3](#), you have to know whether you are above or below this special airspeed to know whether speed changes will give you a long-term climb or a long-term descent.

The value of V_Y will change by a few knots depending on configuration (engine power setting, flap setting, et cetera).

7.5.2 Zero Power Available

Another point on the power curve that is sometimes important is shown in [figure 7.7](#). If there is a point in the

mushing regime where the power curve crosses zero, I will call the corresponding V_Z , the airspeed where there is zero rate of climb.




Figure 7.7: Zero Rate of Climb

In many airplanes, when the engine is developing full power there is no such point; the airplane can climb even at the critical angle of attack, as shown in [figure 7.5](#). On the opposite side of the same coin, with zero engine power, there is once again no such point as V_Z ; the airplane cannot maintain level flight at any airspeed. Therefore the concept of V_Z is only useful in certain circumstances; these include: your airplane's engine could be small by design, or you could be operating at an altitude close to the airplane's absolute ceiling, or you could be having some mechanical problems.

Imagine you are flying at V_Z , with the throttle already wide open, and you want to climb and maintain a higher altitude. Your only option is to *dive*. The dive will give you an airspeed higher than V_Z , closer to V_Y , and by maintaining this new airspeed you will be able to climb.

You can see that flying at V_Z just above the treetops would be a very bad situation. You would not be able to climb at V_Z and you would not be able to accelerate without diving. You are stuck. The only solution is to make sure you never get into such a situation.

Remember that V_Z is defined to be in the mushing regime; the corresponding zero-climb point on the front side of the power curve is completely benign. On the front side you can always climb; all you need to do is pull back on the yoke.

7.5.3 Best Angle of Climb

We see that the power curve is rather flat on top. That means that if you fly a couple of knots faster than V_Y , your rate of climb will hardly be affected at all. You will reach your destination a percent or two sooner, so this sort of "cruise climb" is generally a sensible thing to do.

A more interesting situation arises when you *don't* want to get where you are headed any sooner than necessary — such as when you are trying to climb over an obstacle. In this case it makes sense to climb at an airspeed a few knots below V_Y . The more you slow down, the more time you will have to accumulate altitude before reaching the obstacle. But don't get carried away; the power curve tells you that if you slow down enough, you will degrade the climb performance to the point where further reductions in airspeed don't pay.




Figure 7.8: Best Angle of Climb

As indicated in [figure 7.8](#), the optimum obstacle clearance strategy (the best angle of climb) is achieved at the point where the tangent to the power curve goes through the origin. That means that small changes in airspeed are causing exactly proportionate changes in rate of climb, hence no change in angle of climb. The airspeed where this occurs is denoted V_X . Larger changes away from V_X can only degrade the angle of climb.

We now consider the situation during a descent. This could happen because the aircraft is above its absolute ceiling (which is the normal situation for gliders) or because the engine is operating at a reduced power level.

If all you want is maximum time aloft, you should fly at V_Y as discussed above.

If, however, you want to clear an obstacle and/or glide to a particular place fairly far away, you care about distance (not just time aloft). Once again we observe that the power curve is rather flat on top. That means that if you glide a couple of knots faster than V_Y , your time aloft will hardly be affected at all, but you will get to your destination sooner. This gives you a better chance of getting there before you run out of altitude.




Figure 7.9: Maximum Distance Glide

We can use the tangent trick again. The best distance (i.e. best angle) is achieved at the point where the tangent to the power-off power curve goes through the origin. That means that small changes in airspeed are causing exactly proportionate changes in descent rate.

Once again, the airspeed where you get the best angle can be denoted V_X even though in this case it is a descent angle not a climb angle. In the particular case where you have exactly zero engine power, the best angle occurs right at the point where the aircraft achieves its best lift-to-drag ratio. The airspeed where this occurs is denoted $V_{L/D}$.

7.5.4 Power Depends on Altitude via True Airspeed

Let's compare high-altitude flight with low-altitude flight at the same angle of attack. Assume the weight of the airplane remains the same. Then we can make a wonderful chain of deductions. At the higher altitude:

- the lift is the same (since lift equals weight)
- the lift-to-drag ratio is the same (since it depends on angle of attack)
- the drag is the same (calculated from the previous two items)
- the thrust is the same (since thrust equals drag)
- the indicated airspeed is the same (to produce the same lift at the same angle of attack)
- the true airspeed is greater (because density is lower)
- the power required is greater (since power equals drag times TAS)

The last step is tricky. Whereas most of the aerodynamic quantities of interest to pilots are based on IAS, the power-per-thrust relationship depends on TAS, not IAS.

This means that any aircraft requires more power to maintain a given IAS at altitude. This applies to propellers, jets, and rockets equally.

Another way of getting the same result is to observe that the drag *force* is the same, so getting from point A to point B requires the same amount of energy — since energy is just force times distance. On the other hand, at altitude the airplane gets from A to B more quickly, because of the increased TAS. This requires more power — since power is energy per unit time.

This has no direct effect on V_Y or V_S , or on the general shape of the power curve; it just shifts the curve downward by a scale factor. At high altitudes, this shift will have a huge effect on cruise speed and rate of climb.

7.5.5 Other Power and Altitude Effects

There are lots of additional insights to be gained from thinking about the power curve and its tangents.

- During a descent, the airspeed for best (flattest) angle of descent is necessarily greater than the airspeed for best (slowest) rate of descent. This might not have been obvious from the numbers in your POH.
- During a normal climb, the airspeed for best angle of climb is necessarily less than the airspeed for


best rate of climb.

- As the rate of climb decreases (due to increased power required, reduced engine power output, or whatever) V_X increases and gets closer to V_Y . An interesting case occurs when the airplane is at its absolute ceiling. Then there is barely enough power to maintain level flight at V_Y . The tangent (through the origin) is horizontal, so V_X must be equal to V_Y .

Previously ([section 7.5.4](#)) we considered how much power was *required* as a function of altitude, without any mention of how much engine power you were actually using. Now we consider the effects of engine power.

You might choose to change the engine power, or you might be forced to do so. As altitude increases, sooner or later its power output will decrease. Any power-change will cause small distortions in the shape of the power curve. Let's try to understand why.

Recall that $V_Y(100)$ denotes the airspeed for best rate of climb when the engine is producing 100% of its rated power, while $V_Y(0)$ denotes the airspeed for minimum sink (best endurance) with zero engine power.



[Figure 7.10](#): Power Curve Affected by Engine & Propeller Efficiency

It would be nice if engine efficiency and propeller efficiency were independent of airspeed, but this is only approximately true. Designers often sacrifice a little climb performance in order to get better cruise performance. This means that the effect of engine power is to raise some parts of the power curve more than others, as shown in [figure 7.10](#).

In particular, points to the right of $V_Y(0)$ are raised a little more than points to the left thereof. As a consequence, $V_Y(100)$ must sit somewhere to the right of $V_Y(0)$. At intermediate power settings, V_Y is somewhere between $V_Y(0)$ and $V_Y(100)$. The shift is usually not large.

One reason why efficiency depends on airspeed is *propeller slip*. The propeller is not a solid disk that throws the air straight backwards; there is a certain amount of leakage between the blades and around the edge of the disk. Actually, propellers are typically about 80% efficient at cruise, which is surprisingly good.

At a given engine RPM, propeller slip depends in complicated ways on the indicated airspeed (which determines the drag on the airplane, hence the load on the propeller) and on true airspeed (which determines the angle at which the blades meet the oncoming air).

Taking the efficiency to be independent of airspeed is a reasonable approximation for constant-speed props² but not as good for fixed-pitch props.

Propwash over the wings changes the stall speed, which moves sideways the leftmost points on the power curve. Other propwash effects fiddle with the curve in various minor ways.

[7.5.6](#) Wind Effects

V_Y is not affected by wind (because it only involves altitude and time, not distance). On the other hand, if you are gliding into a headwind toward a distant objective, you want to glide a little bit faster than in the no-wind case, because you want to give the wind less time to push you away from your objective.




Figure 7.11: Best Penetration of 30 Knot Headwind

Once again we can use the tangent construction, as shown in [figure 7.11](#). If there is a 30-knot headwind, the tangent should go through a point 30 knots to the right of the origin. Because of the shape of the curve, the point of tangency does not move 30 knots, but only about 7 knots. Glider pilots call this point the *penetration* speed. As a rule of thumb, when gliding into a moderate headwind, increase the glide speed by about a quarter of the windspeed.

When gliding with a tailwind, you can go farther by gliding a little more slowly than in the no-wind case, but not by very much. Even with an infinite tailwind, it would never pay to glide more slowly than V_Y .




Figure 7.12: Best Penetration of 500 fpm Downdraft

If you are gliding through a downdraft, you want to fly a little faster so you can get out of it as soon as possible. The construction in [figure 7.12](#) can be used to analyze the situation. Given a 500 fpm downdraft, the tangent should pass through a point 500 fpm above the origin. By the same token, if you are flying through an updraft, you want to stay in it as long as possible, so you can reduce the glide speed. The tangent should pass through the appropriate point below the origin.

7.5.7 Weight Effects

A Cherokee Six is a rather popular airplane. It has very good load-carrying ability; more than half of the legal max gross weight is useful load. Even allowing for a bantamweight pilot and a modest amount of fuel, you can imagine flying it at half of max gross weight.

For reasons discussed in [section 2.12.4](#), at reduced weights every point on the power-off power curve is rescaled to a lower speed. In particular, if the weight is reduced by a factor of 0.5, the stalling speed, the best-endurance speed, the maneuvering speed, etc. are reduced by a factor of 0.707 (a 29% reduction). The vertical speeds are reduced by the same factor. This is shown by the lower two curves (the power-off curves) in [figure 7.13](#).




Figure 7.13: Power Curves at Reduced Weight

Start with the standard-weight, power-off curve, then shrink it. For each point, the new airspeed is 71% of the old airspeed, and the new vertical speed is 71% of the old vertical speed. This produces the half-weight, power-off curve.

Now, when we apply power, the full-weight curve moves up by about 1000 feet per minute, thereby turning a

500 fpm descent into a 500 fpm ascent at V_Y . When we apply power to the half-weight airplane, the same amount of energy is devoted to lifting half as much mass, so the curve shifts by twice as much — 2000 fpm. The best rate of climb is more than 1500 fpm.

The cruise speed increases a little at the lighter weight, but not by very much. This is because almost all of the drag at high speeds is parasite drag, which depends on the shape of the airplane, not on its weight or angle of attack.

7.6 Variations in the Power Curve

As mentioned in [section 1.2.5](#), the general shape of the power curve is more-or-less the same for all airplanes, but there are some variations.

7.6.1 Power Curve Depends on Aspect Ratio

Consider a typical airplane in which the stalling speed is 60 KCAS and V_Y is 75 KCAS. We know that V_Y depends on a balance between induced drag and parasite drag, so let's consider what happens if we rearrange things a little bit.

In particular, imagine replacing the wings. The new span will be twice as large, and the new chord will be half as large. This leaves the wing area unchanged, but increases the *aspect ratio* (the ratio of span to chord) by a factor of four.

In the modified airplane, the stalling speed will be very nearly the same, since this depends mainly on wing area. Also the parasite drag will be more-or-less unchanged.

However, the amount of induced drag at any particular airspeed will be less, since the long wing doesn't need to produce such strong wake vortices, as discussed in [section 3.12.3](#).

Therefore V_Y will no longer be 75 KCAS. We can fly slower (thereby reducing parasite drag) without incurring a proportionate increase in induced drag.

The same thing happens if you do something that increases the parasite drag, such as towing a banner. The airspeed V_Y representing the optimal tradeoff between induced and parasite drag shifts to a lower value.

In the extreme case of a high aspect ratio and lots of drag, V_Y might be only a few knots above the stall. You could reasonably take off, fly around all day, and land, without ever operating on the back side of the power curve.

At the other extreme, consider an airplane with a short wingspan, lots of chord, and not very much drag. A typical fighter jet is a good example. For such a plane, V_Y is very much higher than the stall speed. Takeoff, landing, and many other maneuvers must be conducted quite far back on the back side of the power curve.

7.6.2 Sketching the Curve

If you know a few points on the power curve, you can sketch the whole curve. As mentioned in [section 1.2.5](#), the general shape of the curve is the same for all airplanes, so you just need to shift and rescale the curve to fit your particular airplane's performance numbers.

Some of the numbers are easy to obtain, while others are not. For instance:

1. The power-off and power-on stalling speed can be obtained from the POH, and can be easily measured. The corresponding rates of descent generally cannot be obtained from the POH, and would be very hard to measure.
2. The POH gives the airspeed for best rate of climb, and the resulting vertical speed.
3. The POH gives the airspeed for best angle of glide, and resulting angle, from which you can infer the vertical speed.
4. The cruise airspeed can be obtained from the POH. At cruise power setting, the rate of climb at this speed is zero by definition. But what is the power-off rate of descent at this airspeed? You cannot find that in the typical POH, so you may want to measure it experimentally.

You need an estimate of the cruise-airspeed power-off descent rate in order to plan your descent as you

approach your destination, or when ATC asks you to cross a surprisingly-nearby fix at a surprisingly-low altitude.

On the other hand, the rate of descent at stalling angle of attack doesn't usually matter, because if you cared about rate of descent you'd be flying at some other airspeed.

I don't know all the details of the power curve for the airplanes I fly, and unless you are an airplane designer or test pilot, you probably don't need to know the details either. Accurately measuring the entire power curve is (a) unnecessary, (b) *much* harder than you might think, and (c) beyond the scope of this book.

7.6.3 Some Theory

The following mathematical formula may be of additional help in sketching and understanding the power curve. Using the basic lift/drag model introduced in [section 4.4](#), we expect that

$$\begin{aligned} \text{dissipation at } V &= 0.75 \frac{V^3}{V_Y} + 0.25 V_Y^3 \\ & \quad (7.1) \end{aligned}$$

Using this formula, you can get an estimate of the shape of the front side of the power curve using only one measurement, at least for planes where V_Y is not too close to the stall. Here's the idea: measure the sink rate at V_Y , and attribute three quarters of it to induced drag and one quarter of it to parasite drag. Then, as the airspeed increases, the power dissipated by induced drag will go down like the reciprocal of the airspeed while the power dissipated by parasite drag will go up like the cube of the airspeed. As you can see from [figure 4.15](#), this won't be exact, but it will be close.

There is a discussion of coefficients, forces, and powers in [section 4.4](#). See also [section 4.5](#).

7.6.4 Power Requirements versus Speed

Suppose we want an airplane with a reasonably high cruise speed. How much power does it take?

In particular, let's suppose our airplane can stay airborne at an airspeed of $V_Y = 75$ KIAS, using 100 horsepower (at a particular altitude). Now let's suppose we want the cruise speed to be double that speed, namely 150 KIAS (at the same altitude). Then we expect (based on the formula given above) to need 240 horsepower during cruise.

If we want to double the cruise speed again, to 300 KIAS, we need to increase the power to over 1600 horsepower! We see that in the high-speed regime, doubling the power causes an eightfold increase in the parasite drag power. (The total increase in dissipation is somewhat less than eightfold, because the induced drag component isn't increasing.)

Note that when you increase the airspeed from 75 to 150 KIAS, the power goes up by a factor of 2.4 but the gas mileage gets worse by only 20%. That's because mileage depends on fuel per unit distance, not fuel per unit time, and you would get to the destination in half the time. Similarly, when you increase the speed from 150 to 300 KIAS, the power goes up by a factor of 6.8, but the gas mileage gets worse by only a factor of 3.4.

Of course, you can reduce the power requirement (and fuel requirement) by redesigning the airplane to reduce the coefficient of parasite drag, but big improvements are usually not very easy to achieve.

7.6.5 Power Requirements versus Altitude

The previous section considered different speeds at the same altitude; now we consider different altitudes at the same indicated airspeed.

The angle of attack is the same, the lift force is the same (just equal to weight), and the drag force is the

same — all independent of altitude, if we keep the indicated airspeed the same.

However, this does *not* mean that the required power remains the same. The drag power is equal to the drag force times the airspeed (true airspeed, not indicated airspeed). This means that for any given cruise IAS, the power required grows as a function of altitude, in the same proportion as the TAS/IAS ratio.

This puts a limit on how high you can fly, even if you have a turbocharged engine whose output is independent of altitude.

To look at the same fact another way, let's consider speed at constant power (rather than power at constant speed). Assuming constant engine power, if your cruise altitude is higher your indicated cruising speed must be lower (closer to V_Y), so that you can operate at a point on the power curve appropriate to the reduced power available.

At any altitude where you have plenty of power, the IAS is large compared to V_Y and drops only slowly as the available power declines (because power depends roughly on the *cube* of the IAS). In this regime the TAS is increasing even as the IAS is dropping. In contrast, as the altitude approaches the absolute ceiling, the cruise IAS is near V_Y , small changes in IAS don't make much difference, and any decrease in available power causes the IAS to drop toward V_Y so fast that the TAS drops, too.

7.7 Energy Management Stunts

7.7.1 High-Speed Steep Descent

Here is an anecdote that illustrates a peculiar technique for getting rid of energy in a hurry. I tried this once, back when I was a private pilot with about 100 hours' experience. I was approaching a tower-controlled airport and had requested landing clearance. Unfortunately, the tower controller was tied up for a while, talking on his land-line. Eventually he said, ``cleared to land, if you can make it from there''. The problem was that I was 2000 feet above the runway, and less than two miles from the touchdown zone. That makes a ten degree glide slope — pretty darn steep.

The wisest thing would have been to foresee and avoid the whole situation; that is, I should not have allowed myself to get so close at such a high altitude. Failing that, the next-wisest thing would have been to request approval for a 360° turn, so I could lose altitude smoothly.

However at that point in my pilot career I had more aerodynamic knowledge than wisdom, so I used another (rather unprofessional) method for getting rid of the excess energy. As a flight instructor I really don't recommend this technique, but as a physicist I have to admit that it works (if properly carried out) and it illustrates a couple of interesting points about energy management.

Anyway, here's the story: I accepted the clearance, immediately extended full flaps, reduced the power to idle, and dived at the ``top of the white'' — the maximum allowable flaps-extended airspeed.

The situation is illustrated in [figure 7.14](#), which compares my steep, high-speed glide with a normal power-off glide. To give an indication of speed, the figure shows a stopwatch symbol every 15 seconds along each path.



[Figure 7.14](#): High-Speed Dive versus Normal Glide


The high-speed dive was different from the normal approach in several ways:

- At each instant, the airplane was at a lower altitude than it would have been if I had flown a normal-speed approach. (Compare the altitude of corresponding stopwatches in [figure 7.14](#).) This is because my chosen airspeed was on a very draggy part of the power curve. I was relying on this to solve my energy problem.
- At each instant, the airplane was closer to the airport than it would have been if I had flown a normal-speed approach. (Compare the horizontal position of corresponding stopwatches in the figure.) This is an unavoidable consequence of the higher airspeed. It was unhelpful, because it

meant I had less time to get rid of the excess energy.

- Effect (1) was bigger than effect (2). That is, the drag increase was disproportionately larger than the airspeed increase. This is true anywhere on the front side of the power curve, at speeds greater than $V_{L/D}$. (As discussed in [section 7.5](#), at speeds near $V_{L/D}$, a small increase or decrease in airspeed leaves the direction of flight unchanged; you just move a little faster or slower along the same glidepath.)
- Even though I was cashing in altitude energy at a prodigious rate, I had to remember that at each instant I had more *airspeed energy* than I would have had normally. I needed a plan to deal with this surplus at some point.

If my high-speed glidepath had taken me directly to the runway threshold, I would have arrived at the threshold with far too much kinetic energy, and would have had a hard time landing the plane. Fortunately, I could tell early in the maneuver that my glidepath led to a point about an eighth of a mile short of the runway. About a quarter of a mile from the runway, as my glide intercepted the normal glidepath, I smoothly pulled back on the yoke. This flare-like maneuver brought the airspeed down to normal. I retrimmed appropriately. I was then able to follow a steep but non-ridiculous power-off approach path the rest of the way to the runway, at a normal airspeed, followed by a normal flare and landing.



[Figure 7.15](#): High-Speed Diving Approach Details

This strategy — diving at a very high airspeed toward a point short of the runway, so there will be enough time and distance left to get rid of the excess airspeed, is diagrammed in [figure 7.15](#). I reiterate that this stunt is not normal pilot technique. Still, it is a good energy-management illustration, and sometimes it is helpful during forced-landing practice.

Remember that the flaps are not the only way of getting rid of unwanted energy. Depending on what type of airplane you fly, you may prefer a high-speed *slip* and/or extending the landing gear early. You can also make a circle or two.

[7.7.2](#) Low-Speed Steep Descent

Looking at [figure 7.9](#), you may suspect that you can increase the angle of descent by flying at speeds well below $V_{L/D}$. In principle, this is possible — but such a procedure is even more unwise and unprofessional than the high-speed procedure discussed in the previous section.

The main problem is that by the time you achieve a significant increase in descent angle, your airspeed will be much too close to the stall. A slight gust, windshear, or imperfection in pilot technique could cause a stall. Remember, stalling on approach is the #1 way to cause a fatal accident.

A secondary problem with such a procedure is that it probably involves such a nose-high pitch attitude that you can't see where you are going. A third problem is that you might not have enough energy to flare; if you try to raise the nose too quickly it will just cause an accelerated stall.

It is possible to construct scenarios (such as landing on a very short runway with an obstructed approach) where a steep descent on the back side of the power curve is the only way to get the job done. However, before attempting such a task, you should make sure you have the appropriate specialized training and practice. In most cases it is wiser to just choose a different place to land.

[7.7.3](#) Skimming in Ground Effect

Here is trick for *saving* a little bit of energy. I hope you never get into a situation where you need to use this trick — but it might save your bacon if the situation arises.

Suppose no engine power is available, and the aircraft is too low and/or too far from the desired landing place. Using our energy-management logic, we see that the only real way to stretch the glide is to find a low-drag mode of operation. The solution is sort of the reverse of a soft-field takeoff ([section 13.4](#)) — you should make use of ground effect.

Specifically, the procedure is to maintain best-glide speed⁸ right down into ground effect, even if this means that you enter ground effect over the swamp a tenth of a mile short of the intended landing place. Once you are in ground effect, start pulling back on the yoke. Because there is very little induced drag in ground effect (as discussed in connection with soft-field takeoffs in [section 13.4](#)), the airplane can fly at very low airspeeds with remarkably little drag. You can then fly all the way to the landing area in ground effect. It is like a prolonged flare; you keep pulling back gradually to cash in airspeed and pay for drag. This technique will not solve all the world's problems, but it is guaranteed to work better than trying to stretch the glide by pulling back before entering ground effect.

Conversely: if you are approaching a short runway and have a few knots of excess airspeed on short final, you should pull back on the yoke and get rid of the excess airspeed *before* entering ground effect. If you think you can't get rid of it on short final, remember it will only be harder to get rid of in ground effect. A timely go-around might be wise.

If you want to practice skimming in ground effect, find a long, long, long runway to practice on, and be careful not to run off the far end.

7.8 Summary

Most pilots are very aware of their precise altitude, but (alas) not nearly so aware of their precise airspeed or angle of attack.

The airplane is trimmed for a definite angle of attack, and hence (more or less) a definite airspeed. The yoke is part of the angle-of-attack control system. Pulling back on the yoke will always make you slow down.

If you are on the front side of the power curve *and* if you don't mind airspeed excursions, you can use the yoke as a convenient, sneaky way to control altitude. This is because airspeed is linked to altitude via the law of the roller-coaster and via the power curve.

Warning: just because this works OK 99% of the time, don't get the idea that it works all of the time. Bad habits are easy to learn and hard to unlearn. Do not get the idea that pulling back on the yoke always makes the airplane go up. On the back side of the power curve, it doesn't work — and might kill you. In critical situations (including approach and departure), you simply must control the airspeed using the yoke and trim.

The throttle controls power. Power is energy per unit time. To overcome drag requires power. To accelerate requires power. To climb requires power.

In flight, if you open the throttle a normal airplane will not accelerate — it will climb.

Whereas opening the throttle causes energy to enter the mechanical system, you can also encourage energy to leave the mechanical system by extending the flaps, the spoilers, the landing gear, etc., and/or by choosing a draggier place to sit on the power curve.

If you want to fly precisely, you need to look at the altitude *and* the airspeed, size up the energy situation, and then decide what to do with the yoke *and* the throttle.

¹

Once again, this assumes the airplane is in flight (not resting on its wheels) so that the trim mechanism is effective. This also neglects the small nonidealities discussed in [section 6.1.4](#).

²

Be careful to call these the ``power versus time'' curves. If you shorten this to ``power curve'', people will think you mean the power-versus-airspeed curve.

³

...for reasons discussed in [section 6.1.4](#).

⁴

...or (more precisely) angle of attack, as discussed in [chapter 2](#).

⁵

...except for perhaps using the yoke to prevent slight phugoid oscillations at the beginning and end of the climb, where the pitch attitude changes.

⁶

The term V_y is defined to be the airspeed for best rate of climb. It applies even when the rate of

climb is negative.

7

...which incorporate a governor that adjusts the pitch of the propeller.

8

Actually you might want to fly a tiny bit faster than best-glide speed, so you enter ground effect sooner.

8 Yaw-Axis Torque Budget

In aircraft (unlike cars, bikes, or sailboats) you have separate control over which way it is *pointing* relative to which way it is *going*.

8.1 Overview

This chapter discusses the motion of the airplane around the yaw axis.¹ The conventional definitions of the three principal axes are shown in [figure 19.9](#) in [section 19.6.1](#). For a discussion of the terminology and general principles of forces and moments, you can refer to [section 19.7](#). (The motion around the other axes is discussed in the following chapters.)

For a precise definition of terms such as *yaw angle*, *heading*, and *slip angle*, please refer to [section 19.6.3](#).

It is important to keep the slip angle to a minimum, that is, to make sure the aircraft is *pointing* the same direction as it is *going* through the air. This is important for several reasons:

- Efficiency: If your objective is to turn to the left, it doesn't make sense to let the maneuver begin with a big inadvertent yaw to the right. Also, all that slipping creates lots of unnecessary drag.
- Comfort: Passengers really hate being sloshed from side to side. In many trainers, the pilots are seated so close to the pivot point (the center of mass) that they don't appreciate the effect yawing motion can have at other locations.
- Safety: Whereas if you stall in coordinated flight the nose will just drop straight ahead, if you manage to stall in sufficiently uncoordinated flight, you will get a spin (see [chapter 18](#)) or a snap roll (see [chapter 11](#)), which is much harder to recover from.

Maintaining zero slip angle while maneuvering requires coordinated use of the ailerons and rudder, so pilots speak of "zero slip angle" and "good coordination" almost interchangeably.

This section considers, one by one, the various phenomena that affect the airplane's motion around the yaw axis. There are surprisingly many such phenomena, including the helical propwash, yaw-axis inertia, adverse yaw, P-factor, and gyroscopic precession — plus the stability and damping created by the vertical fin and rudder.

8.2 Yaw Stability

An airplane is not always *pointing* in the same direction as it is *going*. This is a new concept for most people, since for ordinary objects such as cars, bicycles and sailboats, the direction they are pointing is (more or less) always the same as the direction they are going, and certainly there is no separate control of the two directions.

As an extreme example of the new concept, take a Frisbee and draw on it the picture of an airplane. When you throw the Frisbee, the picture of the airplane will turn around and around and around. The direction it is pointing has no connection with the direction it is going.

In a sailboat or airplane, you can change the heading with the rudder. In the airplane (unlike the sailboat) the resulting sideways forces are too small to be of much use in changing the direction of motion (but see [section 8.10](#)).

To change the direction of flight of a Frisbee or airplane, the proper procedure is to put it into a bank and *lift* it around the turn.

Unlike a Frisbee, an airplane is not free to turn completely around the yaw axis. If the slip angle (i.e. the difference between the direction it is pointing and the direction it is going) gets to be more than a few degrees, performance is compromised. The vertical fin is designed to keep the slip angle from getting too large.

[Figure 8.1](#) shows a situation where the airplane's heading has been disturbed out of its usual alignment with the airflow. There are lots of ways this could happen, including a gust of wind, a momentary uncoordinated deflection of the controls, or whatever.




Figure 8.1: Response to a Yaw Angle

In this situation, the relative wind is striking the vertical fin and rudder at an angle. Like any other airfoil, the fin/rudder produces lift in proportion to its angle of attack, so it will produce a force (and therefore a torque) that tends to re-align the airplane with the wind. We say that the airplane has lots of yaw-axis stability.

The colloquial name for yaw-axis stability is ``weathervaning tendency''. That is, the airplane tends to align itself with the relative wind, just as a weathervane does. [Section 8.11](#) discusses weathervaning during taxi.

8.3 Yaw Damping

Pure yawing motions are reasonably well damped. The process is analogous to the process that produces damping of pure vertical motions and pure rolling motions (see [chapter 5](#)). When the tail is swinging to the right with an appreciable velocity, it sees a relative wind coming from ahead *and to the right*. The resulting angle of attack produces a leftward force that damps the rightward motion.

A leftward force in proportion to a rightward velocity is exactly what constitutes damping.

8.4 Helical Propwash

One of the very first things that people find out about when they start learning to fly is that it takes right² rudder (sometimes a lot of right rudder) to keep the airplane going straight at the beginning of the takeoff roll. The physics of the situation is portrayed³ in [figure 8.2](#).




Figure 8.2: Helical Propwash

It would be nice if the propeller would just take the air and throw it straight backwards, but it doesn't. The propeller airfoil necessarily has some drag, so it drags the air in the direction of rotation to some extent. Therefore the slipstream follows a helical (corkscrew-like) trajectory, rotating as it flows back over the aircraft.

The next thing to notice is that on practically all aircraft, the vertical fin and rudder stick up, not down, projecting well above the centerline of the slipstream. That means the helical propwash will strike the left side of the tail, knocking it to the right, which makes the nose go to the left, which means you need right rudder to compensate.

You don't notice the effect of the helical propwash in cruise, because the aircraft designers have anticipated the situation. The vertical fin and rudder have been installed at a slight angle, so they are aligned with the actual airflow, not with the axis of the aircraft.

In a high-air-speed, low-power situation (such as a power-off descent) the built-in compensation is more than you need, so you need to apply explicit left rudder (or dial in left-rudder trim) to undo the compensation and get the tail lined up with the actual airflow.

Conversely, in a high-power, low-air-speed situation (such as initial takeoff roll, or slow flight) the helix is extra-tightly wound, so you have to apply explicit right rudder.

Helical propwash sometimes contributes to left/right asymmetry in multi-engine aircraft, as discussed in [section 17.2.12](#).

8.5 P-Factor

The term *P-factor* is defined to mean ``asymmetric disk loading''. It is an extremely significant effect for helicopters. When the helicopter is in forward flight, the blade on one side has a much higher airspeed than

the other. If you tried to fly the blades at constant angle of attack, the advancing blade would produce quite a bit more lift than the retreating blade.

8.5.1 Blade Speed

For airplanes, the same effect can occur, although it is usually small. For the effect to occur at all, you need to have an angle between the propeller axis and the relative wind. To be specific, imagine that the aircraft is in a nose-high attitude, but its direction of motion is horizontal (i.e. the relative wind is horizontal). Then the downgoing blade will be going down and a little bit forward, while the upgoing blade will be going up and a little bit backward. The downgoing blade will effectively have a slightly higher airspeed. Since this blade is on the right-hand side of the airplane (once again assuming a typical American engine) it will tend to torque the airplane around to the left and you'll need right rudder to compensate.




Figure 8.3: P-Factor

The situation is depicted in [figure 8.3](#). The airplane is in level flight, with a 10 degree nose-up attitude. The motion of the blade through the air is shown in magenta. It consists of the rotational motion (shown in green) plus the forward motion of the whole airplane (shown in red). The motion of the downgoing blade is shown with solid lines, while the motion of the upgoing blade is shown with dotted lines. You can see that the speed of the downgoing blade is larger than the speed of the upgoing blade.

This is the main contribution to P-factor: the advancing blade sees more relative wind, while the retreating blade sees less relative wind.

8.5.2 Blade Angle

There is a widespread misconception that P-factor arises because the angle of the right (downgoing) propeller blade is larger than the angle of the left (upgoing) propeller blade. Many books erroneously call attention the angle of the blade relative to the ground. The blade doesn't care about the ground; the only thing that matters is the angle of attack, i.e. the angle between the blade and its own motion through the air.

The correct analysis is shown in [figure 8.4](#). As a point of reference, the left panel shows level pitch attitude in normal level flight, where no P-factor occurs. Meanwhile, the right panel shows the airplane in a 10 degree nose-up attitude (still in level flight). Since we want to discuss angle of attack, I have attached a "reference line" pointer to each of the blades, just like the reference line used in [section 2.2](#). The angle of attack of the propeller blade is just the angle between the reference line and the blade's motion through the air.⁴




Figure 8.4: P-Factor and Angle of Attack

When the propeller disk is inclined to the direction of flight (so that P-factor really is occurring) the downgoing

blade has a slightly greater angle of attack (compared to the upgoing blade) as shown in [figure 8.4](#). This occurs because the vector representing the airplane's motion has "better leverage" when it meets the resultant, because the resultant is shorter and because it is more nearly perpendicular to the airplane's motion.

This angle-of-attack effect is of course zero when propeller axis is aligned with the direction of flight.⁵ The effect is never very large, because

1. At low speeds, the airplane's forward velocity (as represented by the horizontal red arrow in [figure 8.4](#)) is so small that it can't have much effect on anything.
2. At high speeds, the airplane has a low angle of attack, so the angle between the propeller disk and relative wind is necessarily small (except for helicopters, tilt-rotors, and such).
3. At very high speeds, when you are going fast enough to over-run the geometric pitch of the propeller (so that the resultant coincides with the reference line in [figure 8.4](#)), you might think that a small difference in angle of attack would be a 100% effect. I suppose that's true, but in this case the total thrust is practically zero, and 100% of nothing is nothing.

This angle-of-attack effect is in addition to (and usually smaller than) the airspeed effect discussed previously. Both are small compared to the helical propwash effect.

Remember, we don't care whether the downgoing blade makes a bigger angle *to the vertical* than does the upgoing blade. The blade doesn't care which way is up — all it cares about is where the relative wind is coming from. Imagine a tailwheel-type airplane stationary in the run-up area on a windless day. You can incline the propeller disk as much as you want relative to vertical, but there will be no P-factor unless there is *wind* blowing through the propeller disk at an angle.

[8.5.3](#) Initial Takeoff Roll

There are quite a lot of myths surrounding P-factor. For some reason, P-factor gets blamed for the fact that typical aircraft require right rudder on initial takeoff roll. This is impossible for several reasons.

- Nearly everybody these days learns to fly in nose-wheel type aircraft, which means the propeller disk is vertical during the initial the takeoff roll. Since there is no angle between the relative wind and the propeller axis, P-factor obviously cannot occur.
- Now let's suppose, just for sake of argument, that you are flying a taildragger, in which the propeller disk is actually non-vertical during the initial takeoff roll. Common experience is that the most right rudder is required at the very beginning of the takeoff, before much forward speed has been achieved. The **Flight Training Handbook** says this is because P-factor is worst at low airspeeds. But wait a minute — real P-factor is proportional to airspeed. In the initial moments of the takeoff roll, there is no relative wind, so there can't possibly be any P-factor. Of course, if you are taking off into a headwind, there could be a little bit of P-factor — but does that mean if you take off with a slight tailwind there will be a negative amount of P-factor, requiring left rudder? Don't bet on it.

The real reason that you need right rudder on initial takeoff roll is because of the helical propwash, as discussed in [section 8.4](#). P-factor exists in some circumstances, but it cannot possibly explain the behavior we observe during initial takeoff roll.

[8.5.4](#) Observing P-Factor

It is not easy to observe P-factor. It is hard to separate P-factor from other yaw contributions such as helical propwash ([section 8.4](#)) and twisted lift ([section 8.8.4](#)).

An important preliminary experiment is to observe what happens during the takeoff roll in a multi-engine aircraft. (To be specific, let's consider the case where the engines rotate clockwise as seen from the rear.) In some airplanes, the propwash hits the tail, and you must apply right rudder to compensate, just like in single-engine planes.

In other airplanes, most of the helical propwash misses the vertical tail in normal flight. This causes no problems and no compensation is required. See [section 17.2.12](#) for details on this. This is the perfect way to illustrate that there is no P-factor when the propeller disk is not inclined.

If you really want to observe nonzero P-factor, you can proceed as follows: Take a twin-engine (or four-engine) aircraft with non-counter-rotating propellers. Attach a slip string. Establish coordinated cruising flight, with the same amount of power on both sides. Confirm that the ball and string are centered. Pull up

into a nonturning climb at very low airspeed (i.e. very high angle of attack), maintaining cruise power. Maintain coordinated flight as indicated by the slip string. Observe the rolling tendency due to propeller drag. Shift weight (e.g. fuel) from left to right to get rid of the rolling tendency, so you can fly straight *without* deflecting the ailerons, i.e. without incurring any twisted lift.

You will observe the inclinometer ball will be slightly off-center. This can be attributed to P-factor. To be explicit:

- Each engine is producing the same *amount* of thrust.
- But the result is asymmetric thrust, because the *location* of the thrust vectors is shifted (because of asymmetric disk loading).
- We have gone to some trouble to minimize other contributions, such as twisted lift. Also note that even if there were some helical propwash hitting the tail, that wouldn't be a good explanation for the off-center inclinometer. You should be able to compensate for helical propwash using the rudder alone, without banking.
- To summarize: we have symmetric power settings but asymmetric thrust. A detailed analysis of the effects of asymmetric thrust can be found in [section 17.2.5](#).

The effect of P-factor is not very large. You can easily compensate using a little bit of right rudder and right bank. Indeed, in typical situations you can just ignore it entirely.

You can never use rudder deflection as an indication of P-factor, because any situation that exhibits P-factor will also exhibit a large amount of helical propwash.

The single-engine version of the previous experiment goes like this: Put a slip string on each wing, far enough out on the wing that it is not unduly disturbed by the propwash, yet close enough in that you can see it. In a high-wing aircraft, you'll have to put the string on the bottom of the wing. Put strings on both sides in symmetric locations, so you can tell for sure what string position corresponds to symmetric airflow. Then confirm that in normal nonturning cruising flight, you have symmetric airflow (as indicated by the strings) and zero inclination (as indicated by the inclinometer ball). Finally, set up a situation in which the largest possible P-factor occurs: flaps retracted, minimum airspeed, and full power.

Once again, the indication of P-factor in this situation would be to have the ball be off-center when the strings were centered. I have tried this experiment, but the P-factor was too small to observe.

Here's another possible experiment. Take your favorite aerobatic airplane and paint the starboard rudder pedal green and the port rudder pedal red, just so we can keep straight which is which. Now go to a safe altitude and set up for inverted slow flight. In this high-power, low-speed situation, do you need to push the port (red) pedal or the starboard (green) pedal? If P-factor is more important, the answer will be port, because that is now the downgoing, advancing blade. If helical propwash is more important, the answer is starboard, because the relationship between the propeller, rudder, and rudder pedals is unchanged by the inversion.

[8.6 Gyroscopic Precession](#)

A spinning object will respond to a torque around one axis with a motion around another axis. This remarkable and counterintuitive phenomenon — *gyroscopic precession* — is discussed in more detail in [section 19.9](#).

Gyroscopic precession is often quite noticeable at the point where a taildragger raises the tail, early in the takeoff roll.⁶ If the airplane were an ordinary non-spinning object, you could raise the tail using the flippers alone. The flippers do not actually dictate the *motion* of the fuselage; they just produce a *force* and a torque around the pitch axis. For a gyroscope, a torque around the pitch axis produces a motion around the yaw axis. If you try to raise the tail of a real airplane using flippers alone, it will yaw to the left because of precession.

To get a gyroscope to actually start *moving* around the pitch axis, you need to apply a torque around the yaw axis. This is what the vertical fin and rudder are for. See [section 19.9](#).

Of course, an airplane has some plain old mass in addition to its gyroscopic properties. In order to lift this ordinary mass you need to use the flippers. Therefore, the tail-raising maneuver requires both flippers and

rudder — flippers to change the pitch of the ordinary mass, and rudder to change the pitch of the gyroscope.

8.7 Canted Engine

Often the engine is mounted in such a way that direction of the thrust vector is a little to one side of the axis of the airplane. This is done in order to compensate for various nonidealities such as helical propwash. It contributes to the yaw-axis torque budget in the obvious way.

8.8 Rudder Usage During Rolls

Turning the airplane properly requires coordinated use of ailerons and rudder. Getting it exactly right is a bit tricky.

Remember that in an airplane, the direction you are moving is not necessarily the same as the direction you are pointing. There are several crucial things that happen during a turn:

- 1) You use the wings to change the direction your center of mass is going. I call this the *CM-turn*.
- 2a) You use the rudder to change your heading (i.e. to overcome yaw-axis inertia, i.e. to provide yaw-axis acceleration).
- 2b) You use the rudder to overcome steady adverse yaw due to twisted lift, as discussed below.
- 2c) You use the rudder to overcome transitory adverse yaw due to differential drag.

Item 1 is relatively straightforward: you put the airplane into a bank. The horizontal component of lift will change the direction of motion.

Item 2a is important because if the airplane didn't have any vertical tail, banking would cause it to just *slip* off in the new direction without changing its heading. It is much nicer to yaw the plane to align its axis with the new direction of motion, so you apply the rudder, thereby creating a yaw rate that matches the CM-turn rate.⁷

Now we come to item 2b. We must consider adverse yaw. As discussed in [section 8.8.4](#), during a steady roll, the aerodynamic forces produced by the two wings are equal in magnitude, but one force vector is twisted slightly forward while the other one is twisted slightly rearward. This causes a yawing moment in exactly the wrong direction: if you are rolling to the right it tries to make the airplane yaw to the left. To compensate you must deflect the rudder whenever the ailerons are deflected.

Finally, we come to item 2c. Suppose you are flying an airplane where there is a lot of mass out on the wings. Whenever you are starting or ending a roll maneuver, you need to accelerate one wing upward and the other wing downward. As discussed in [section 8.8.3](#), this briefly requires extra lift on one wing and reduced lift on the other wing. This unequal lift produces unequal induced drag. This drag causes additional adverse yaw.

For any given rate of roll, you need to use lots more rudder at low airspeeds, for reasons discussed in [section 8.8.6](#).

Procedures for maintaining coordination during turns are summarized in [section 8.8.7](#); the intervening sections describe in a little more detail what is the problem we are trying to solve.

8.8.1 Analysis of a Roll

To make the discussion more concrete, let's consider a roll starting from straight-and-level flight and rolling to the right. As we can see from [figure 8.5](#), there are multiple timescales in the problem.




Figure 8.5: Timeline for Roll Maneuver

- $[t_0, t_1]$ the time it takes you to move the ailerons;
- $[t_1, t_2]$ the time it takes for the roll rate to reach the value commanded by this aileron deflection;
- $[t_2, t_3]$ the time it takes the yaw rate to reach the value corresponding to the rate of CM-turn;
- $[t_3, t_4]$ the time you hold the ailerons deflected.

Let's analyze what happens if you move the ailerons fairly abruptly. Although generally I recommend flying with a smooth, gentle touch, (1) there will be times when you want to roll the airplane on short notice, so let's learn how to do it; and (2) the abrupt case makes it easier to understand what is going on.⁸

In some airplanes, such as a Piper Cub, the roll rate will reach its final very quickly (within a small fraction of a second), because the airplane has very little inertia about the roll axis. Practically all the mass (pilot, passenger, fuel, and engine) is arranged in a straight line right on top of the roll axis, so they don't contribute much moment of inertia. In other airplanes, such as a Cessna 310, the roll rate responds much more slowly, because lots of mass (engines and tip tanks) is situated far from the roll axis.

Before the roll rate is established (i.e. during the time $[t_1, t_2]$) the plane will experience transitory adverse yaw due to differential induced drag. The nose will swing a little toward the outside of the turn. The effect is usually rather small, since

1. these differential drag forces are typically small (during slow flight) or very, very small (during cruise), compared to the differential lift forces that cause the roll, and
2. these forces must act against the yaw-axis inertia, which is at least as large as the roll-axis inertia.

The rest of the discussion applies no matter how slowly or abruptly you moved the ailerons.

After the time t_2 , a steady roll rate exists. Even though the ailerons are deflected, there is no difference in lift from one wing to the other, for reasons discussed in [section 8.8.4](#). Since there is no difference in lift, there will be no difference in induced drag, hence no transitory adverse yaw.

However, one wingtip is diving, so its force vector is twisted slightly forward. The other wingtip is rising, so its force vector is twisted slightly rearward. Even though each force has practically the same *magnitude* as it would in non-rolling flight, the twist means there is a slight component of force in just the right direction to produce a steady adverse yawing moment.

In addition, because the airplane is rolling, a bank is developing. This bank causes a CM-turn; that is, the airplane is changing its direction of motion. In order to keep it *pointing* in the same direction it is *moving*, you need to deflect the rudder during the roll, as discussed in [section 8.8.5](#).

At time t_6 , the ailerons are neutralized, but the rolling motion has not yet stopped. (Again, there is a delay due to roll-axis inertia.) At this point there are several things going on:

1. There is a difference in lift between the two wings, as needed to damp out the roll. This creates a negative amount of transitory adverse yaw. This requires a left-rudder contribution to compensate.
2. However, the airplane is still rolling, and a still-increasing rate of yaw is needed to coordinate with the still-increasing rate of CM-turn. This requires a right-rudder contribution.
3. Similarly, because the airplane is still rolling, the twisted lift requires a right rudder contribution.

In practical situations, the first item (transitory adverse yaw) is usually smaller than the other two. During the interval $[t_5, t_7]$ the roll rate is decreasing, so you need less and less rudder deflection.

Analogous statements would apply if you started from a left turn and used right aileron and right rudder to roll out of the turn. Similarly, it is easy to do a similar analysis for rolling into a left turn and/or rolling out of a right turn.

8.8.2 Designers' Tricks

Imagine an airplane without a vertical fin. It would behave more like a Frisbee than a sailboat — if you gave it a yaw rate, inertia would make it just keep on yawing until some torque acted to stop it. Even if it wasn't yawing, there would be no reason to expect the yaw angle (i.e. heading) to be anywhere close to the desired value.

In a real airplane, of course, the vertical fin and rudder supply the forces required to keep the yaw angle and yaw rate under control. An overview of how you use the rudder during turns can be found in [section 8.8](#).

Aircraft manufacturers know about how turns are affected by twisted lift and yaw-axis inertia. They generally try to provide the needed yaw-axis torque automatically, using various tricks. One trick is to interconnect the rudder and ailerons with a spring. That means you automatically get a certain amount of rudder deflection in proportion to the aileron deflection. They choose the proportionality factor so that you can more or less fly "with your feet on the floor" at cruise airspeeds. Of course, vastly more rudder is needed at lower airspeeds; fortunately you can easily overpower the interconnect spring by pushing on the controls in the obvious way.


Here's another trick, which you may have noticed on many airplanes: when one aileron goes down a little, the other one goes up a lot. (This is called *differential aileron deflection*.) The designers were trying to arrange for the upward-deflected aileron to generate a lot of parasite drag. If they do it just right, the drag force is just enough to overcome twisted lift and yaw-axis inertia during a steady roll. The so-called *Frise aileron* uses a similar trick. It has lip on the bottom, well ahead of the hinge. The lip sticks down into the airstream when the main part of the aileron is deflected up. Again, the purpose of the lip is to generate drag on the wing with the upward-deflected aileron.

In addition to overcoming yaw-axis inertia (during a steady roll), the designers also want to overcome transitory adverse yaw (when ailerons have been deflected but the roll hasn't yet started). Fortunately, transitory adverse yaw is rather small, and by adjusting the amount of differential deflection, and the amount of the Frise effect, pretty good cancellation can be achieved.

The bad news is that this compensation only works at one airspeed. The designers arrange it so you can fly with your feet on the floor during cruise. This is a mixed blessing, because it can lull you into complacency. At lower airspeeds, where it is most important, you still need to use lots of rudder to keep things coordinated. Don't forget!

8.8.3 Transitory Adverse Yaw

Suppose you wish to roll into a right turn. You will deflect the ailerons to the right, as shown in [figure 8.6](#). During the brief time after the ailerons are deflected and before the steady roll is established, this will increase the lift created by the left wing, and decrease the lift created by the right wing. Unfortunately, there is no way to produce lift without producing drag, so the left wing will be dragged backwards while the right wing lunges forward. This is the exact opposite of what we wanted; the airplane yaws to the left even though we wanted it to turn to the right. Being a good pilot, you have anticipated this, so you apply right rudder as well as right aileron, to make sure the nose swings the right way.



[Figure 8.6](#): Transitory Adverse Yaw


Even if you don't get the footwork exactly right, the nose will eventually swing around and point more-or-less the right way, because of the airplane's inherent yaw stability (as discussed in [section 8.2](#)).

Once a steady roll rate is established (no acceleration around the roll axis), the two wings are producing the same amount of lift, so this type of adverse yaw will no longer exist.⁹

Now let's consider what happens if you wish to roll out of the turn. The airplane is banked to the right and already turning to the right. You will deflect the ailerons to the left. This will cause extra drag on the right wing, and reduced drag on the left wing. The airplane will yaw to the right, continuing and exaggerating the turn that you were trying to stop. Anticipating this, you apply left rudder along with the left aileron

8.8.4 Steady Adverse Yaw -- Twisted Lift

Now let's consider what happens during a steady turn. As illustrated in [figure 8.7](#), the airplane as a whole is moving forward, but the left wingtip is moving forward and *up* while the right wingtip is moving forward and *down* (because of the rolling motion).



[Figure 8.7](#): Steady Adverse Yaw -- Twisted Lift

Let's see what the *local* angle of attack is at the wingtip. We use the trusty formula
$$\text{angle of attack} + \text{angle of climb} = \text{pitch} + \text{incidence} \quad (8.1)$$

In the figure, the right wingtip has a negative angle of climb, since it is going forward and down. But the deflected aileron gives it a lower incidence, effectively twisting that section of airfoil nose-down. By the same token, the left wingtip has a positive angle of climb (due to the rolling motion) and an increased incidence (due to the aileron).

In a steady roll, the incidences just cancel the climb angles, so that the left wing and the right wing end up flying at the same angle of attack. If they didn't cancel, you wouldn't have a *steady* roll.

The cancellation means there is no torque around the roll axis, but the yaw axis is a different story. As you can see in the figure, the force vector for the downgoing wing is twisted forward, while the force vector for the upgoing wing is twisted rearward. This pair of fore-and-aft force components creates a torque around the yaw axis. You need to deflect the rudder to compensate.

Some people try to argue that these force-components should be called ``drag" forces since they are directed fore and aft, in the same direction as the overall relative wind. However, it is much better to think of them as components of the local lift, since the twisted lift remains perpendicular to the *local* relative wind. The strongest argument is this: a drag force should dissipate energy in proportion to force times airspeed, but it is clear that the twisted lift forces do not dissipate energy.¹⁰

8.8.5 Yaw-Axis Inertia

(In this section, we will assume that you are flying at such a low airspeed that the designers' tricks discussed in [section 8.8.2](#) are not sufficient to produce automatically coordinated turns.)

Whenever the airplane is in a bank, it will make a CM-turn. A pure CM-turn, however, is not what you want. A pure CM-turn means that even though the airplane is *moving* in a new direction, the *heading* hasn't changed. The airplane has a nonzero slip angle. The uncoordinated airflow acting on the tail will eventually set up a yawing motion that matches the CM-turn rate, converting it from a pure CM-turn to a more-or-less¹¹ coordinated turn. If the yaw axis damping is weak, as it usually is, the nose will slosh back and forth several times as it tries to catch up with the CM-turn.

At any particular CM-turn rate, once the yaw rate is established, no further yaw-axis torque is required. Like a toy top, once the airplane starts rotating around the yaw axis it will be happy to continue rotating.¹² The only time you need a yaw-axis torque is when the yaw rate is changing.

So, we see that during a steady roll,

- The ailerons are deflected.
- The bank angle is increasing and, correspondingly, the rate of CM-turn is increasing.
- To match the CM-turn rate, the yaw rate must increase.
- To increase the yaw rate, the rudder should be deflected.

Conclusion: the rudder should be deflected when the ailerons are deflected.

8.8.6 Amount of Rudder Required

As we have seen, there are three reasons you need to apply the rudder during roll maneuvers: yaw-axis inertia, differential induced drag, and twisted lift.

The amount of rudder deflection you need depends on the shape of your airplane, and also depends on airspeed.

*** Twisted Lift**

Example 1: Consider an airplane with long wings and with most of the mass concentrated near the middle of the airplane. A typical glider is an excellent example, but almost any ordinary-shaped airplane will do. In this case there will be very little roll-axis inertia, and accordingly very little transitory adverse yaw. There will also be rather little yaw-axis inertia. Therefore in such a plane, the dominant effect will be steady adverse yaw due to twisted lift.

Example 2: Suppose you are flying along in *any* airplane on a sunny summer day. You encounter a situation where your right wing is in an updraft, while your left wing is in a downdraft. You deflect the ailerons in order to maintain zero bank, zero roll rate, and constant heading. This combination of non-horizontal relative wind and deflected ailerons creates twisted lift, the same as shown in [figure 8.7](#) (except that the roll rate is zero in this case). Therefore this is a perfect example of steady adverse yaw, and you must deflect the rudder to compensate. (This rudder requirement could not be explained by differential drag or yaw axis inertia. This is pure twisted lift.)

The yawing moment due to twisted lift is essentially independent of airspeed. It just depends on the deflection-angle of the ailerons. Meanwhile, though, the force produced by the rudder is proportional to airspeed squared. Therefore you need lots more rudder deflection (per unit aileron deflection) when the airspeed is low.

*** Differential Induced Drag**

Example 3: Consider an aircraft where there is a lot of mass located far away from the roll axis. A twin with heavy engines mounted way out on the wings, plus tip-tanks full of fuel, is a good example. Such a plane will have lots of roll-axis inertia, and therefore lots of transitory adverse yaw. You will still have to worry about yaw-axis inertia and twisted lift, but *in addition* to those effects you will need to apply *extra* rudder deflection when ailerons are first deflected, before the steady roll develops.

The amount of rudder required depends dramatically on airspeed. In addition to the rudder-force issue discussed above, the amount of transitory yawing moment itself increases when the airspeed decreases. The key to understanding this is to realize that whereas the coefficient of lift is more or less proportional to the angle of attack (for moderate angles of attack), the coefficient of induced drag is more or less proportional to the *square* of the angle of attack.

The left side of [figure 8.8](#) shows the same situation as in [figure 8.6](#), along with the coefficient of drag curve. On this curve I have indicated the different angles of attack for the two wingtips, and the correspondingly different amounts of drag. We see that the coefficient of drag curve is relatively flat on the bottom, so at relatively small angles of attack (high airspeeds), a difference in angle of attack doesn't cause too much difference in drag.




Figure 8.8: Slow Flight Means More Transitory Adverse Yaw

In contrast, the right side of [figure 8.8](#) shows the same aircraft in *slow flight*. Both wings are operating at a higher angle of attack. Because the coefficient of drag curve is steeper in this regime, the same difference in angle of attack (i.e. the same aileron deflection) creates more difference in drag (i.e. more transitory adverse yaw).

* Yaw-Axis Inertia

Example 4: Consider a long, thin, single-engine biplane carrying lots of cargo. Since it has a rather short wingspan, there will be rather little twisted lift, i.e. rather little steady adverse yaw. Similarly, since all the mass is close to the roll axis, there will be very little roll-axis inertia, i.e. very little transitory adverse yaw. There will, however, be lots of yaw-axis inertia.

Example 5: Let's return to the case where your right wing is in an updraft, while your left wing is in a downdraft. This time, however, you don't deflect the ailerons; you just accept the resulting roll rate. During the steady roll, you will need to deflect the rudder to supply the yaw-axis momentum to match the ever-increasing CM-turn rate. (This rudder requirement could not be explained by twisted lift or differential drag. This is pure yaw-axis inertia. Also note that no designers' tricks could maintain coordination in this situation, since the ailerons are not deflected.)

Once again, the amount of rudder required increases markedly at low airspeeds. There are three main contributions; the first two essentially cancel each other:

1. The roll rate depends on the deflection-angle of the ailerons, times airspeed. That means at low airspeeds the roll rate is less, which reduces the amount of rudder required.
2. The amount of turn that results from a given bank angle increases at low airspeeds, as discussed in [section 16.6](#). This increases the amount of rudder required.
3. Finally, as always, rudder effectiveness depends on airspeed squared, increasing the amount of rudder deflection required at low airspeeds.

8.8.7 Summary: Coordinated Turning Procedures

A proper turn consist of two ingredients: a CM-turn and a heading change. In an idealized ``basic" airplane, you would use the ailerons to bank the airplane and lift the CM around the corner, and you would use the rudder to change the heading and combat adverse yaw. In a typical modern airplane *at cruise airspeeds*, deflecting the ailerons alone creates a fair approximation of the proper torques around both axes. In all airplanes *at low airspeeds*, proper rudder usage is vitally important.

The basic rule is simple:

- if you are rolling to the right, you must apply right rudder;
- if you are rolling to the left, you must apply left rudder.

The *amount* of rudder will depend inversely on the airspeed.

Another version of the rule substitutes the word "aileron" for "roll":

- right aileron requires right rudder;
- left aileron requires left rudder;

In a steady roll, the two versions are more or less equivalent; at the beginning and end of a roll (when the roll rate does not match the aileron deflection) the truth lies somewhere in between. Split the difference.

The foregoing rule allows you to anticipate the need for rudder deflection. As discussed in [section 11.5](#), you have many ways of knowing when you've got it right:

1. The acid test involves looking out the window. You should perceive that the rate of heading change is proportional to the amount of bank.
2. You can also look to the side and perceive that the wings flap straight up and down, not slicing fore and aft as you roll.
3. You can see that the inclinometer ball remains almost centered.
4. Yet more information comes from the seat of your pants.

By the way: If you think about it for a moment, you can see that in inverted flight (negative angle of attack) you will have a *negative* amount of adverse yaw — if you deflect the yoke to the left you will need to push on the right rudder pedal, and vice versa — just the opposite of what you would do in noninverted flight. When you are actually in the plane, hanging upside down, this is not as confusing as it seems on paper. A little thought and a little practice will make it fairly self-evident which wing you should lower to make a CM-turn and which rudder pedal you should push to change the heading.


As mentioned at the beginning of this chapter, there are lots of reasons why you should use the rudder properly during turns. Alas, the learning process is complicated by the fact that in many cases the airplane will "cover up" small mistakes for you. In particular, whenever the airplane is in a slip, the vertical fin will automatically try to return the plane to zero slip angle. This is the yaw-axis stability discussed in [section 8.2](#). The plane will (under most conditions) eventually establish an approximately correct rate of heading change anyway. The goal of correct rudder usage is to establish the correct yaw-axis motion *without* a slip developing even temporarily.

The dependence of adverse yaw on airspeed can lead to trouble. Pilots spend almost all of their time buzzing around at cruise airspeeds, where ignoring the rudder is OK or nearly so. Sometimes this leads to complacency. The problem arises on approach and/or departure, where airspeeds are much lower. Proper coordination becomes more challenging, exactly at the place where it is most important (since the margins for error are also smaller). If you mishandle the ailerons at low speed and low altitude, you could well cause a spin or a snap roll, with no chance for recovery.

[Section 11.5](#) describes a few useful tricks for perceiving exactly how much rudder is needed to achieve perfect coordination.

8.9 Long-Tail Slip

Now let's see what happens while the airplane is in an established turn. In particular, let's consider an airplane with a fairly long fuselage, flying in a fairly tight turn. As shown in [figure 8.9](#), there is no way that the airflow can be lined up with the front part of the fuselage and the back part of the fuselage at the same time. The fuselage is straight, and the path through the air is curved. You can't have a straight line be tangent to a circle at two different points. You have to choose.



[Figure 8.9](#): Airplane in a Tight Turn — Rudder Neutral

Because of the air hitting the sides of the fuselage, and other effects, the fin/rudder might not completely determine the slip angle, but it will be the main determining factor. For sure, the airflow at the front of the fuselage — and over the wing — will have a significant slip component.


You can see from the geometry of the situation that the amount of long-tail slip is proportional to the length of the airplane and inversely proportional to the turning radius. The latter depends on the *square* of the airspeed, as well as the bank angle.

Diagram illustrating a vehicle on a banked curve without friction. The vehicle is shown on a curved path, and the text "No Slip" indicates that the vehicle is not sliding. The diagram shows the forces acting on the vehicle: the normal force (perpendicular to the surface) and the weight (vertical). The vehicle's velocity is indicated by a blue arrow.

I emphasize that even though you are holding inside rudder (bottom rudder) during the turn, this is definitely not a skidding turn (unless you get carried away and use *too much* inside rudder). This rudder usage is completely unrelated to the uncoordinated ``boat turn" discussed in [section 8.10](#).

The best way to check the alignment is with a *slip string*¹³ — a piece of yarn exposed to the airflow where the pilot can see it. If (as is usually the case) you don't have a slip string, you can refer to the inclinometer ball.¹⁴

My friend Larry has a sailboat. It doesn't have ailerons. You steer it with the rudder.¹⁵ This changes the direction the boat is pointing. As shown in [figure 8.11](#), this causes the water to flow crosswise past the hull, creating a sideways force that eventually changes the direction the boat is going.



[Figure 8.11](#): Boat Making a Boat Turn




Figure 8.12: Airplane Making a Boat Turn

All the same words can be applied to an airplane. Keeping the wings level, you press the right rudder pedal. This causes the airplane to yaw to starboard. As shown in [figure 8.12](#), air will then hit the fuselage on the port side, creating a sideways force¹⁶ that will gradually shove the airplane around in a right-hand turn. (There will also be a lot of drag, but that is not our concern at the moment.) The force of the wind on the rudder (needed to yaw the plane) is smaller than, and in the opposite direction to, the resulting force of the wind on the fuselage.

In powered flight, the horizontal component of thrust will make an additional contribution to the boat turn.

To reiterate: the airplane will turn to the right if you hold the right rudder pedal down — even if the wings are not banked. Of course, turning the airplane properly (using the wings) is ten times more effective and more efficient than a boat turn

8.11 Weathervaning During Taxi

When the airplane is on the ground, it feels the force of the ground *and* the force of the wind.

Since the tail is far, far behind the wheels, a crosswind will create a torque around the yaw axis. It will tend to blow the tail downwind, forcing the nose to turn upwind, just like a weathervane.

Now, suppose you are moving (as opposed to parked). The weathervaning tendency causes the nose to turn into the wind. The wheels are still on the ground, making lots of friction, so the airplane will roll in the direction determined by the wheels, i.e. the direction it is heading. Therefore the airplane will travel toward the *upwind* side of the runway. This may seem ironic or even paradoxical, but it's true — the crosswind causes the airplane to move upwind.¹⁷ You have to deflect the rudder to downwind to compensate.

8.12 Asymmetric Thrust

In a multi-engine airplane, if the engine on one side has failed, or for any reason is developing less thrust than its counterpart on the other side, this will produce a torque (possibly a very large torque) around the yaw axis. This is discussed in [section 17.2.4](#).

8.13 Yaw-Axis Torque Budget — Summary

We have finally come to the end of this section, having covered the most important causes and effects of torques and motions around the yaw axis. There are quite a number of such processes:

- The helical propwash effect is important, especially in high-power / low-airspeed situations.
- Gyroscopic precession means that deflecting the flippers will cause a yawing motion (and deflecting the rudder will cause a pitching motion).
- Adverse yaw means that deflecting the ailerons will cause a yawing moment.
- The long-tail slip effect means that an inadvertent turn will be a slipping turn. This effect is very significant in gliders. It is much less noticeable in typical powered aircraft, but it has important implications for roll stability, as discussed in [section 9.3](#).
- P-factor exists in principle but is usually insignificant.
- Actual motion around the yaw axis will create a yawing moment that tends to damp the motion.
- Yawing the plane changes the direction it is *pointing* which does not automatically change the

direction it is *going*; the ``boat turn" effect exists but is feeble and inefficient.

- The pilot can deflect the rudder to oppose the unwanted yawing effects and create the desired ones.

Some of these ideas will be revisited when we discuss ``Dutch roll" in [section 10.6.1](#).

Perceiving coordination and maintaining coordinated flight is important. Further discussion of this topic appears in [chapter 11](#), along with a discussion of how and why to perform intentional slips.

[1](#) For a discussion of the terminology of forces and moments, you can refer to [section 19.7](#).

[2](#) All the examples in this section assume a typical American engine that rotates clockwise as seen from behind.

[3](#) The figure exaggerates the curvature of the streamlines.

[4](#) You can also think of the blade's angle of attack as the angle between the reference and the blade's relative wind. Relative wind and direction of motion are the same concept, just reversed 180 degrees. Be careful though, because there are various different relative winds, including the instantaneous wind relative to the moving blade and the average wind relative to the overall airplane.

[5](#) Interestingly, it goes to zero again when the axis is perpendicular to the direction of flight, as in a helicopter.

[6](#) but if you pay attention you can notice it in many other situations

[7](#) Yaw-axis acceleration (which may be a somewhat unfamiliar subject) is discussed in more detail in [section 8.8.1](#).

[8](#) This analysis ignores the overbanking tendency and various other small effects.

[9](#) Although in a steady turn you may need some rudder deflection because of the long-tail slip effect, as discussed in [section 8.9](#), and in a steady roll you will need some rudder deflection because of twisted lift and roll-axis inertia, as discussed in the following sections.

[10](#) The real drag vector gets twisted, too, but the consequences are too small to worry about.

[11](#) It won't be exactly coordinated because of the long-tail slip effect, as discussed in [section 8.9](#).

[12](#) But you will generally need some rudder deflection to compensate for the long-tail slip effect.

[13](#) This is more commonly called a yaw string, but in fact it measures the slip angle, not the yaw angle. This is discussed in more detail in [section 17.2.3](#).

[14](#) However, the inclinometer ball and the slip string actually measure quite different things. The distinction is noticeable whenever the rudder is deflected, as discussed in [section 17.2.4](#).

[15](#) Boat lovers' note: there are some ocean liners that do use roll-control devices rather like ailerons, although they are primarily for passengers' comfort, not for steering. Also, to be sure, there are some boats that can be steered by banking them. On my sailboard, for instance, you have to bank it the wrong way (i.e. to the outside of the turn) by shifting your weight. On some light racing yachts you can steer them pretty well just by shifting the weight of the crew around. Many speedboats bank into the turns. But we're getting off the subject. The point is that Larry's boat (like lots of others) leans to leeward whether you are turning left, turning right, or going straight. The reason it doesn't tilt any more than it does is because there are a couple thousand pounds of lead in the keel. You can't bank it by shifting your weight, and it wouldn't turn much if you did. You steer it with the rudder.

[16](#) Technically, this force is classified as a lift force, since it is perpendicular to the relative wind — even though it is produced by the fuselage (not the wings), and even though it is horizontal. See the official definitions in [chapter 4](#).

[17](#) In those rare cases where there is inadequate friction on the wheels (such as a seaplane, or an

airplane taxiing on a slick icy surface) it is quite possible for the wind to blow the airplane downwind. This of course has nothing to do with torque; it's just a plain force.

9 Roll-Axis Torque Budget

Non-pilots have the idea that pilots keep the airplane in the air by firm intervention and lightning reflexes. In fact, though, pilots are judged more on smoothness than quickness.

This chapter considers the various forces that could impart a rolling moment to the airplane.¹

9.1 Dihedral


Back in [section 8.2](#), we discussed how an uncoordinated relative wind² will affect the yaw axis; now let's see how it will affect the roll axis.

The first thing that people think of in this connection is *dihedral*. The word comes from the Greek word for ``two planes'' and just means that the two wings are not coplanar, as shown in [figure 9.1](#).



[Figure 9.1](#): Definition of Dihedral

In the presence of dihedral, any uncoordinated (side-to-side) airflow will hit the bottom of one wing and the top of the other wing, as shown in [figure 9.2](#). This means one wing will be forced up and the other forced down. If you work out all the angles between the total relative wind and the wings, you find that indeed the angle of attack is increased on the upwind wing and reduced on the downwind wing. The difference in lift produces a rolling moment. Any process whereby uncoordinated airflow produces a rolling moment it is called a *slip-roll coupling*; dihedral is a good example of this.



[Figure 9.2](#): Dihedral — Slip Produces Rolling Moment


The rolling moment (the torque about the roll axis) will be proportional to the dihedral angle, and proportional to the amount of slip.

9.2 Other Forms of Slip-Roll Coupling

Dihedral is only one of several reasons why an airplane might have a slip-roll coupling. A high-wing airplane has a certain amount of slip-roll coupling because of *interference effects*. That is, when the airplane is in a slip, the fuselage interferes with the airflow over the wing. As shown in [figure 9.3](#) and [figure 9.4](#), the stream lines have to bend a little in order to flow around the fuselage. This creates an updraft at the root of the upwind wing, and a downdraft at the root of the downwind wing. This creates a rolling moment that tends to raise the upwind wing.




[Figure 9.3](#): Redirection — High-Wing Airplane in a Slip (1)




[Figure 9.4](#): Redirection — High-Wing Airplane in a Slip (2)

As shown in [figure 9.5](#) and [figure 9.6](#), on a low-wing aircraft the effect is reversed. There is an updraft at the root of the downwind wing, and a downdraft at the root of the upwind wing. This contributes a *negative* amount of slip-roll coupling.




[Figure 9.5](#): Redirection — Low-Wing Airplane in a Slip (1)



[Figure 9.6](#): Redirection — Low-Wing Airplane in a Slip (2)

A related interference effect is shown in [figure 9.7](#). A fuselage moving sideways through the air is a very non-streamlined object. Downstream of such an object you expect to find a large, messy wake. The air in the wake is less capable of producing lift when it flows over the wing.



[Figure 9.7](#): Turbulence — High-Wing Airplane in a Slip



[Figure 9.8](#): Turbulence — Low-Wing Airplane in a Slip

Because the air that the wing really cares about is coming from ahead and *below*,³ this type of interference is more pronounced in a high-wing airplane — the fuselage is in a stronger position to disturb the relevant airflow. This can be seen by comparing [figure 9.7](#) with [figure 9.8](#).

The magnitude of the effect of the wake is very difficult to predict. It will depend not only on the general shape of the fuselage, but also on the details of the surface finish.⁴ It will also depend very nonlinearly on the airspeed and slip angle.

In general, interference effects mean that (to achieve an adequate amount of slip-roll coupling) low-wing

airplanes typically need more dihedral than high-wing airplanes. You can check this by looking at typical airplanes at your local airport.

A third effect is illustrated in [figure 9.9](#). If you put a swept-wing airplane into a slip, the more-forward wing produces more lift. That wing (the left wing in the figure) presents effectively more span to the airstream. It is a common mistake to think that the increased span explains the increased lift. The mistake is to overlook the fact that when it presents more span, it necessarily presents less chord. Lift, other things being equal, is proportional to wing area, and it is well known that area is not changed by a rotation. The correct explanation has more to do with the direction of airflow. Air flowing spanwise along an airfoil doesn't produce lift. The key idea is that the chordwise component of the airflow is bigger for the left wing.



[Figure 9.9](#): Swept Wing Airplane in a Slip

A fourth effect is shown in [figure 9.10](#). In practically all aircraft, the rudder sticks up above the roll axis. When the aircraft is in a slip, the rudder produces a substantial force. This force times this lever arm produces a torque about the roll axis.

Anything else that sticks up above the roll axis and produces sideways drag or sideways lift contributes the same way. This includes the wings of a high-wing airplane, although the effect is small since spanwise flow along a wing doesn't create much force — just a little bit of sideways drag. This is another reason why high-wing airplanes can get by with less dihedral (for the same amount of slip-roll coupling).



[Figure 9.10](#): Tall Rudder in a Slip

All four effects just mentioned are in the same direction, and can be combined: You can have a high-wing, swept-wing airplane with lots of dihedral and a really high tail — in which case you would probably have more slip-roll coupling than you need.⁵

The propwash contributes a negative amount of slip-roll coupling when the engine is producing power. If you yaw the nose to the right, the uncoordinated component of the wind will blow more of the propwash to the right wing. The extra lift on the right wing will roll you to the left.

Slip-roll coupling is the reason why you can make a relatively normal turn with the rudder (inelegant though it is). If you gently press on the right rudder, you will cause a skid that will eventually produce a bank to the right. Of course the skid itself will also cause a boat turn to the right. If you hold a constant rudder deflection, the boat-turn force will only be proportional to the rudder deflection, whereas the bank (and the associated non-boat turn) will keep getting larger and larger because of the slip-roll coupling.

[9.3](#) Roll-Axis Stability

We are now all set to understand how the airplane responds if, for some reason, one wing goes a bit lower than the other.

The airplane will start to turn. If the turn were perfectly coordinated, the airplane would be happy to keep turning around and around and around. Fortunately, as we recall from our discussion of the long-tail slip effect ([section 8.9](#)), "an inadvertent turn will be a slipping turn". This tiny amount of slip, acting through the slip-roll coupling, will tend to roll the airplane back to wings-level straight-ahead flight.

This process gives the airplane a slight amount of *roll-axis stability*.

Airplane designers always make sure the airplane has a certain amount of slip-roll coupling, for exactly this reason.


The roll-axis stability is rather weak, because the two necessary ingredients are individually weak: The slip-roll coupling is usually moderately weak, and the long-tail slip effect is so weak that (except for glider pilots) most pilots never notice it unless it is pointed out.

Common experience indicates that roll-axis stability is indeed rather weak. If you are cruising along in turbulent air and take your hands off the controls for a couple of moments, you do not expect the nose to pitch up or down 30 degrees, and you do not expect it to yaw left or right 30 degrees, but you would not be at all surprised to have a 30 degree bank develop.

Even in the best of conditions, the stability generated by the long-tail slip with slip-roll coupling can only overcome a small amount of uncommanded bank. For larger bank angles, the overbanking tendency ([section 9.4](#)) takes over and makes the roll axis unstable.


* Dihedral in the Absence of Slip

Before going on, let's take another look at what happens in a coordinated turn. Sometimes it is argued that when the airplane is in a bank, the lowered wing has a bigger footprint (a bigger projection on the ground) than the raised wing, as shown in on the left in [figure 9.11](#). A similar argument was used back in [section 9.2](#) to explain why swept wings produce a slip-roll coupling. There is one slight difference: the swept-wing effect is real (because it involves the direction of the air) whereas the supposed effect of dihedral in a coordinated bank is completely imaginary. The wing doesn't know or care where the ground is. It cares only where the air is coming from. In a coordinated turn, the air is coming from straight ahead, so dihedral has no effect.



[Figure 9.11](#): Dihedral Has No Effect in the Absence of Slip

Other myths about dihedral involve the angle of the lift vectors of the two wings. The correct answer is the same: in the absence of slip, dihedral has no effect. As long as the air is coming from straight ahead, the lift vectors are symmetrically disposed, as shown in [figure 9.12](#).



[Figure 9.12](#): Symmetric Lift Vectors

In a coordinated turn, the aircraft is happy to continue turning forever; it will definitely not have any tendency to return to wings-level flight. Indeed, it will have the opposite tendency, called the overbanking tendency, which we now discuss.

9.4 Differential Wingtip Speed; Overbanking

[Figure 9.13](#) shows the aircraft in a coordinated turn. The outside wingtip follows a path of length $2\pi R$ (big R) while the inside wingtip has the proverbial "inside track" — its path is only $2\pi r$ (little r). Since the outside wingtip travels farther in the same amount of time, it must be moving faster.




Figure 9.13: Overbanking Tendency

The same fact is depicted a second time in the figure — the relative wind is depicted to be stronger on the outside wingtip. Since the lift generated by an airfoil depends on the square of the airspeed, the outside wing would produce more lift (other things being equal). This means that the aircraft in a turn (especially a properly coordinated turn) will tend to bank into the turn more and more. The steeper the turn, the more pronounced this *overbanking tendency* becomes. The next thing you know, you are in a spiral dive (as discussed in [section 6.2](#)).

In order to combat this tendency, you need to deflect the ailerons against the turn.

The strength of this effect depends on the ratio of the wingspan to the radius of turn. If you have stubby wings, high airspeed, and shallow bank angle, you'll never notice the effect. On the other hand, in a glider you might have long wings, low airspeeds and steep turns — in which case you might need nearly full outside aileron just to maintain a steady bank angle.

It is interesting to combine this with what we learned about long-tail slip effect ([section 8.9](#)) — in the slow, steeply banked turn in the glider, you would be holding nearly full inside rudder (to prevent the long-tail slip) and nearly full outside aileron (to counteract the overbanking tendency). If you are not expecting this, it will appear very strange. You are holding completely crossed controls, yet the turn is perfectly coordinated. You can confirm this by observing that the slip string is perfectly centered.

You don't want to have to figure this out while sitting in the glider, trying to make a steep turn. Sometimes it pays to read the book before you go flying.

9.5 Rolling Moment due to Propeller Drag

The engine makes a contribution to the roll axis torque budget.

As we remarked earlier, the propeller does not throw the air straight back; there is some rotational drag on the propeller blades. According to Newton's law of action and reaction, you can see that if the prop throws the air down on the right, it tends to make the airplane roll to the left.

To put it more crudely: take a model airplane (where the propeller rotates to the right) and hold it by the propeller. If you start the engine, the airplane will rotate to the left.

As shown in [figure 9.14](#), some of the rotating air hits the top of the right wing and the bottom of the left wing.⁶ This tends to reduce the amount of roll — but it can never reduce it to zero or cause a roll to the right. Similarly, any air intercepted and "straightened out" by the tail reduces the rolling moment somewhat. Using Newton's law again, we see that if any air escapes while still rotating down to the right, the airplane will roll to the left.



Figure 9.14: Rolling Moment Produced by Propeller

The only way to restore equilibrium is to take a corresponding amount of air and throw it down on the left. Airplane designers have long since learned about this propeller drag rolling moment, and they take steps to compensate for it. For instance, they set the left wing at a slightly higher angle of incidence than the right wing. This is called, unsurprisingly, *asymmetric incidence*. It is especially useful to apply this trick to the part

of the wing that flies in the propwash, so that the effect increases as engine power increases. On a Piper Cherokee, the roll-axis trim is rather easy to adjust — in the flap extension mechanism for each flap there is a turnbuckle that allows the flap to be raised or lowered until the roll-axis trim is just right.

If the roll-axis trim is just right in cruise, it will be nowhere near right during a soft-field takeoff. In that case, the propeller drag will be worse because of the high power, and the fancy rigging of the airfoils will be less effective because of the low airspeed. The result: you will have to deflect the yoke to the right, using the ailerons to counter the prop drag rolling moment.

9.6 Engine Inertia

Newton's second law asserts that force equals mass times acceleration. There is a rotational version of this law, asserting that the rotational force (i.e. torque) equals the rotational inertia times the rotational acceleration. That means whenever the engine RPMs are increasing or decreasing, a torque is produced.


There is also a rotational version of Newton's third law, asserting that if you impart a clockwise rotational momentum to one thing, you must impart a counter-clockwise rotational momentum to something else.

Consider an airplane which has the engine aligned in the usual way, but where the propeller-drag effects (discussed in [section 9.5](#)) are negligible. The easiest way to arrange this is to have a single engine driving two counter-rotating propellers. The Wright brothers used this trick in their first airplane.

While (and only while) the engine speed is changing, the airplane will tend to roll. It will roll to the left if the engine is accelerating, and it will roll to the right if the engine is decelerating.

In steady flight in this airplane, the the engine's rotational inertia has no effect. The fact that the engine / dual propeller system is producing power does not imply that it is producing any net torque.

To clarify the distinction, compare the two situations shown in [figure 9.15](#) and [figure 9.16](#). We have an ordinary single-engine airplane. We have removed the propeller and bolted a huge brake drum onto the propeller shaft. In [figure 9.15](#), the brake shoes are attached to the floor of the hangar. When we run the engine, the brake will produce a huge torque that will make the airplane want to roll to its left. This is completely analogous to the propeller drag effect discussed in [section 9.5](#). There is an instrument called a *prony brake* that measures the torque-producing capability of an engine in precisely this way.



[Figure 9.15](#): Prony Brake Attached to the Floor



[Figure 9.16](#): Prony Brake Attached to the Airplane Itself

In [figure 9.16](#), the brake shoes are attached not to the floor, but to the airplane itself. Even if the engine is producing torque (straining against its engine mounts) all the torques flow in a closed circuit and cancel. The airplane as a whole exhibits *no* rolling tendency.

Newton's law is quite explicit about this: if you want to give the airplane some left-rolling momentum, you have to give *something else* some right-rolling momentum. This ``something else" could be the air (as in [figure 9.14](#)) or perhaps the hangar floor (as in [figure 9.15](#)).

The angular motion of the internal engine parts can only affect the rolling moment if you *change* their rotational speed.

Engine rotational inertia should not be confused with propeller drag. In a direct-drive propeller installation, the propeller-drag torque does act on the fuselage via the engine mounts, but that is a coincidence, not a law of physics. In a gear-drive installation, most of the propeller-drag torque acts on the fuselage via the gearbox.

Finally, consider the case where the engine rotates one way and the propeller rotates the other way (which is easy to arrange using a gearbox). In a steady slow-flight situation, I guarantee you will need to deflect ailerons to compensate for propeller drag; engine inertia *per se* will have no effect on pilot technique.

9.7 Climbing and Descending Turns

In a level turn both wingtips are moving horizontally. In a climbing turn, both wingtips will be climbing, but they will not make equal angles to the horizon. This is because the climb angle depends on the ratio of the vertical speed to the forward speed. As a result of the different climb angles, we get different angles of attack for the two wingtips. The geometry of the situation is shown in [figure 18.6](#) (in the chapter on spins). Another way to think about this is to recognize that it involves rotating around a non-vertical axis, as discussed in [section 19.6.4](#).

Let's do an example. Consider an airplane with a 35-foot wingspan making a standard-rate turn (3 degrees per second) at 100 K_{TAS} while climbing or descending at 500 fpm. We can calculate the resulting angle of attack at the wingtips, as shown in [table 9.1](#).

		Airspeed (K _{TAS})	Vertical speed (fpm)	Angle of climb	Angle of attack
Climbing turn	inside wingtip	99.46	500	2.844°	4.485°
	outside wingtip	100.54	500	2.814°	4.515°
	difference	2.2%			0.7%
Descending turn	inside wingtip	99.46	-500	-2.844°	4.515°
	outside wingtip	100.54	-500	-2.814°	4.485°
	difference	2.2%			-0.7%

[Table 9.1](#): Climbing and Descending Turns

We see that the change in angle of attack typically has less effect than the change in airspeed. In a climbing turn, the angle effect contributes to the overbanking tendency, while in an ordinary descending turn, it somewhat reduces it.

In a spin (which has a higher vertical speed, lower airspeed, and vastly higher rate of turn) the angle effect is extremely significant, as discussed in [section 18.6.1](#). Far from reducing the rolling moment, increasing the angle of attack on the inside wing (which is stalled) only makes the situation worse.

9.8 Roll-Axis Torque Budget — Summary

There are several effects that can give rise to a rolling moment. The most important ones are:

- A slip tends to cause a rolling moment, for several reasons, including: dihedral produces slip-roll coupling; the fuselage shadowing one wing (especially on a high-wing airplane) produces slip-roll coupling; swept wings produce slip-roll coupling, and a tall rudder that sticks up above the roll axis produces slip-roll coupling.
- This slip-roll coupling combines with the long-tail slip effect (discussed in [section 8.9](#)) to give the airplane a small amount of stability around the roll axis.
- At medium or large bank angles the overbanking tendency makes the roll axis unstable; the bank angle will tend to get larger and larger. This produces a spiral dive.
- There exists some medium-small bank angle where the two just-mentioned effects cancel. (That is, the overbanking tendency cancels the stability due to slip-roll coupling plus long-tail slip.) At this bank angle, the airplane will happily continue turning, at constant bank angle, without any help from the

pilot.

- The airplane tends to roll left in high-power / low-airspeed situations, because of propeller drag.
- If you suddenly change the speed of rotation of the engine, the rest of the airplane will be subjected a brief rolling impulse. (Similarly, if you change the *direction* of the axis of rotation, gyroscopic precession will cause yawing and/or pitching moments.) If the speed and direction of rotational motion is unchanging, engine torque will have no noticeable⁷ effects.

Some of these ideas will be revisited when we discuss Dutch roll in [section 10.6.1](#).

1

For a discussion of the terminology and general principles of forces and moments, you can refer to [section 19.7](#).

2

i.e. an airflow pattern that is flowing left-to-right or right-to-left over the fuselage.

3

See the discussion of upwash in [section 3.1](#), including [figure 3.2](#).

4

See the discussion of dimples on golf balls in [section 18.3](#) and in [reference 17](#).

5

Excessive slip-roll coupling will cause the airplane to suffer from Dutch roll, as discussed in [section 10.6.1](#).

6

The figure greatly exaggerates how tightly the flow pattern is wound.

7

There will be all sorts of ``closed-circuit'' torques, but they will not affect the handling of the airplane.

10 Equilibrium, Stability, and Damping


Three of the most useless things in aviation are:

- The airspace above you.
- The fuel not on board.
- The runway not in front of the wheels.

Several parts of this book make use of the concepts of equilibrium, stability, and damping. This section defines the concepts a little more precisely and clarifies the relationships between them.


10.1 Equilibrium

The word *equilibrium* is quite ancient. The word has the same stem as the name of the constellation ``Libra" — the scale. The type of scale in question is the two-pan balance shown in [figure 10.1](#), which has been in use for at least 7000 years. The compound word ``equilibrium" translates literally as ``equal balance" and means just that: everything in balance, no unbalanced forces.



[Figure 10.1](#): Equilibrium — Forces in Balance

The *wheel* is more modern than the balance; it has been in use for ``only" some 5500 years. It provides some more sophisticated illustrations of equilibrium and related concepts. If we make the wheel lopsided by attaching a weight to the rim at one point, the system will be out of equilibrium unless the weight is exactly at the top or bottom. If we prepare the system in an unbalanced state and then let go, it will immediately start rotating.



[Figure 10.2](#): Equilibrium and Stability

To reiterate, and as indicated in [figure 10.2](#), there are three ways to make the wheel be in equilibrium: (1) position the weight at the bottom, (2) remove the weight entirely, or (3) position the weight at the top.

10.2 Stability

Stability has to do with how the system responds if we move it a little ways from its equilibrium position. There are three possibilities:

- *Positive stability* means that if the system is displaced a little ways from its equilibrium position, it will generate a force tending to push it back towards equilibrium. The wheel with the weight positioned at the bottom is an example of positive stability.

- *Neutral stability* (also called *zero stability*) means that if the system was in equilibrium and you displace it slightly, it remains in equilibrium. No force is generated. The perfectly balanced wheel is an example of this.
- *Negative stability* means that if the system is displaced a little ways from its equilibrium position, it will generate a force that tends to push it farther from equilibrium. The wheel with the weight at the top is an example of negative stability.

It usually doesn't make much sense to talk about stability except for systems that are in equilibrium or nearly so.

For a multi-dimensional system, we get to ask about the stability of each axis separately. For example, consider an egg resting on a horizontal table. An ideal egg has zero stability against motion in one direction: it is free to roll around its axis of symmetry. On the other hand, it has positive stability against motion in the end-over-end direction; if you rock the egg slightly by pushing its nose down, it will tend to return to its original state.

10.3 Damping

A system exhibits *damping* if motion of the system produces a force that opposes the motion.

A bicycle wheel provides a good demonstration of a system with very little damping. Assuming the bearings are good and the wheel is not touching anything, when you spin the wheel it will keep going for more than a minute. Air friction produces very small forces that eventually cause the wheel to slow down.


A bicycle wheel that is rubbing against something is much more heavily damped. When it is in motion, rubbing friction can create large forces that oppose the motion and bring the motion to a stop.

A dynamical system can exhibit negative amounts of damping, but this is harder to demonstrate with a simple system. Negative damping tends to make the motion increase, which means that energy is being added to the system from somewhere; therefore simple friction can never produce negative damping.

Nose wheel shimmy of an airplane is a good example of what happens if a system has a negative amount of damping. If the aircraft is moving along the ground at high speed, the nosewheel will eventually hit a pebble or something. The nosewheel is then no longer aligned with the direction of travel. By the usual "castering" principle, this causes a force that tends to return the wheel to its proper position (that is, the wheel exhibits positive stability). Unfortunately, many aircraft have too much stability, and too much inertia in the castering mechanism. The result is that the wheel tends to overshoot its equilibrium position and continue to the other side, going out of alignment in the opposite direction by an even greater amount. The result is an oscillation that quickly grows to large amplitude.

Note the relationship of stability and damping: when the wheel is being forced back toward alignment, the force is toward the equilibrium position (positive stability) but is in the same direction as the motion (negative damping).

To eliminate the shimmy problem, a hydraulic "shimmy damper" is installed on the nose wheel. [Figure 10.3](#) is cutaway drawing showing how a hydraulic damper works. It consists of an oil-filled cylinder, plus a pushrod attached to a disk inside the cylinder. When the pushrod moves from side to side, oil is forced to flow through the small holes in the disk. This creates a force proportional to the velocity of motion — i.e. damping.



[Figure 10.3](#): Hydraulic Damper

Sometimes the fluid leaks out of the damper, and even more commonly the linkages connecting the damper to the wheel become worn and loose. This makes the damper ineffective, whereupon you get a vivid demonstration of negative damping. A preflight check of the damper and linkages is easy and worthwhile.

Also... as discussed in [chapter 5](#), the airplane's rolling motion and pure vertical motion are normally very heavily damped, but this damping goes to zero and becomes negative at the stall.


10.4 Relationship of Stability and Damping

To reiterate: stability refers to a force that arises depending on the *position* of the system; damping refers to a force that arises depending on the *velocity*.

In old-fashioned terminology, what we call "stability" was sometimes called "static stability", and what we call "damping" was sometimes called "dynamic stability". What's worse, occasionally *both* terms were shortened to the single word, "stability", which was unnecessarily confusing.

Also, modern usage prefers "damping" not "dampening" — if you start talking about a "dampener" people will think you want to moisten the system.


Stability can be positive, zero, or negative; damping can also be positive, zero, or negative. A dynamical system can display any combination of these two properties — nine possibilities in all, as shown in [figure 10.4](#). In the top row, the bicycle wheel is dipped in molasses, which provides damping. In the middle row, there is no damping. In the bottom row, you can imagine there is some hypothetical "anti-molasses" that provides negative damping.



[Figure 10.4](#): Stability and Damping — Possible Combinations

10.5 Oleo-Pneumatic Struts

A great example of a device that provides a force that depends on position *and* a force that depends on velocity is the *oleo-pneumatic strut*, which is widely used on landing gear as a combination spring and shock absorber. It consists of a piston in a cylinder filled with both oil ("oleo") and air ("pneuma"), as shown in [figure 10.5](#). If the piston is moved up into the cylinder, the air at the top of the cylinder is compressed. (The hydraulic oil is essentially incompressible.) This "air spring" creates a force that depends on the position. As the piston moves, the oil in the hollow part of the piston is forced to flow through the holes in the disk, creating a force that depends on the speed of motion, using the same principle as the damper discussed previously.



[Figure 10.5](#): Oleo-Pneumatic Strut

It is important that the strut contain the right amount of air *and* the right amount of oil. Problems can arise more easily than you might think.

Suppose that over time, some of the oil leaks out of the strut on your airplane.¹ Your friend, Murgatroyd Fudpucker, borrows the plane and notices during preflight that one of the struts is low — that is, not enough of the piston is protruding from the cylinder. Murgatroyd gets out a bicycle pump and adds air to the strut. Everything now looks fine. During future preflight checks, a passive glance the strut will give you the impression that everything is OK.

Alas, things are not OK. The problem is that *oil* has been replaced with *air*. Since air is a thousand times more compressible than hydraulic oil, the amount of force it takes to make the strut ``bottom out" has been greatly reduced. If you or Murgatroyd makes even a slightly hard landing, the piston will smash against the end of the cylinder, metal to metal. This has roughly the same effect on the airframe as hitting it with a sledgehammer. Repairs could be very, very expensive.

Therefore, if there is any chance that the airplane has been mis-serviced since the last time you flew it, you should check not only the height of the struts, but also their springiness. To check a main-gear strut, lift up the wing a few inches and then let it drop. Similarly, to check the nose strut, lift up the nose (perhaps by pushing down on the tail) a little ways and then let it drop. If any strut compresses more than it should (e.g. if it comes anywhere close to bottoming out), do not fly the airplane until the strut has been properly serviced with air *and* oil.

There is a thin coating of oil on exposed part of the piston, which collects dust. When the piston is shoved into the cylinder, the O-ring will scrub the dirt down the piston and cause it to collect in a ring called the *scrub line*. Observing the scrub line can tell you how close the strut has come to bottoming out recently.

Please do not get the impression from the foregoing discussion that ``air is bad" and ``oil is good". I discovered an airplane recently where nose strut contained no air at all, but contained several inches too much oil instead. Once again, the preflight checklist calls for checking the height of the strut — which was completely normal. Fortunately, I noticed that the strut had no springiness whatsoever; trying to compress a solid column of hydraulic oil is like trying to compress cast iron.

To reiterate: you should make sure that the struts contain the right amount of air and the right amount of oil. Servicing a strut isn't very tricky; it just has to be done right.

10.6 Oscillations

Whenever a system has positive stability but not enough damping, you can expect to see oscillations.

10.6.1 Analysis of Dutch Roll

As remarked in [section 9.3](#), the airplane has only a small amount of stability around the roll axis. You may be wondering why designers don't fix this problem by increasing the slip-roll coupling. The answer is that they are worried about Dutch roll.

Dutch roll is a messy combination of rolling, slipping, and yawing. As we shall see, this combined motion is less damped than the pure rolling, slipping, or yawing motions would be.

A moderate amount of Dutch roll never killed anybody, but it does tend to provoke nausea, especially in passengers.

The Dutch-roll oscillations typically have such a short period (a couple of seconds) that it is a challenge for the pilot to overcome them by working the controls. A spiral dive, on the other hand, develops much more slowly. Therefore if it comes down to a compromise between roll-axis stability and Dutch-roll damping, designers generally increase the damping at the expense of the stability.

To understand where Dutch roll comes from, and how to fight it, gives us an opportunity to combine and apply most of the things we have learned about equilibrium, stability, and damping.

The rolling and yawing motions associated with Dutch roll are shown in [figure 10.6](#); we will discuss the slipping component in a moment.




Figure 10.6: Dutch Roll

The wingtip yaws forward, then rolls up, then yaws backward, then rolls downward, then repeats. The opposite wingtip does the same thing, 180 degrees out of phase. Imagine pedaling a bicycle backwards.

To analyze the damping of the Dutch roll system, we must remember that energy is force times distance; by the same token power (energy flow) is force times velocity. The component of the force in the direction of the velocity is the only thing that matters; the component in the perpendicular direction doesn't count.

We begin by using [figure 10.7](#) to analyze the forces that affect the rolling motion. The velocity and position of the wingtip is shown in red; net changes in the lift vector are shown in blue.

At point A in the figure, the wing is going upward. That means it has less angle of attack than normal (and in particular, less angle of attack than the opposite wingtip). The reduced lift corresponds to a net force opposite to the velocity, and therefore energy is being removed from the system. At point C, a similar analysis applies. The wingtip is descending, creating more angle of attack and more lift than normal. This corresponds to a net force which is once again opposite to the velocity, removing energy from the system. This is the same roll damping mechanism as discussed in [section 5.4](#).




Figure 10.7: Dutch Roll — Roll Forces

At point B, the wingtip has less velocity than normal, and less lift, while at point D the wingtip has more velocity and produces more lift. There is no effect on the damping, because the forces are perpendicular to the velocity.

We continue by using [figure 10.8](#) to analyze the forces that affect the yawing motion. At point B in the figure, the vertical fin/rudder is wagging to the right. This changes the rudder angle of attack, opposing the motion. This is the same yaw damping mechanism discussed in [section 8.3](#). Also, at this point, the port wingtip has less drag than the other, because it is moving backwards. Both of these effects take energy out of the system, providing damping. The same processes produce damping at point D also.




Figure 10.8: Dutch Roll — Yaw Forces

At point A, there is a little less induced drag on the port wingtip because it is flying at reduced angle of attack. This has no effect on the damping, because the force is perpendicular to the velocity.

Also at point A, there is a yawing force because the airplane's heading is not aligned with its direction of travel; the tail is too far to the left. This provides yaw-axis stability but does nothing for the yaw damping, because the force is perpendicular to the velocity.

The analysis of point C is analogous to point A.

If the yawing and rolling motions were the whole story, Dutch roll would be no problem. According to the analysis so far, there is lots of positive damping. The Dutch roll would quickly die out.

Unfortunately, nature is not so kind, as we discover when we take the sideways motion of the aircraft into account. Refer to [figure 10.9](#).




Figure 10.9: Dutch Roll — Slip Causes Problems

At point B in the figure, the left wingtip is at the highest point in the cycle. The airplane is banked to the right. The wings' lift vector is inclined to the right, so there is a rightward component of lift. In fact, during the whole half-cycle from point A to point C there is at least some rightward force. Since the airplane has lots of inertia and not much damping² with respect to pure sideways motion, the rightward velocity just increases and increases during the whole half-cycle. The maximum rightward velocity is achieved near point C.

During the next half-cycle (from C via D to A) the airplane is banked to the left. The leftward force reduces the previously-acquired rightward velocity to zero, and then builds up a leftward velocity. The sideways velocity is zero at point D, and the maximum leftward velocity is achieved near point A.

Note that like any other lightly-damped oscillator (such as a pendulum, for instance a playground swing set) the maximum rightward *force* occurs when the plane is at its maximum leftward *position*.

The final ingredient is the slip-roll coupling.³ A certain amount of slip-roll coupling is highly desirable because it is a necessary part of the process that produces roll-axis stability ([section 9.3](#)).

The bad news is that the slip-roll coupling contributes a negative amount of damping to the Dutch roll mode. The rightward velocity is maximum at point C, producing a leftward-rolling moment. The force is in the same direction as the roll velocity, so it adds energy to the Dutch roll.

Analogously, the leftward velocity is maximal at point A, producing a rightward-rolling moment. This, too, is in the same direction as the roll velocity, contributing negative damping.

So slip-roll coupling presents designers with a dilemma: it increases roll-axis stability, but decreases (Dutch) roll damping.

The simplest way a designer can resolve this dilemma is to notice that roll-axis stability depends on both slip-roll coupling and the long-tail slip effect. Therefore if you have a problem with Dutch roll, decrease the slip-roll coupling and increase the long-tail slip effect, for instance by making the tail boom longer and reducing the rudder area. As a rule of thumb, you can tell just by looking at a short-coupled airplane that it will have a problem with underdamped Dutch roll.

The other (all too common) design choice is to sacrifice stability. Most airplanes wind up with very, very little roll-axis stability. Consequently spiral dives are a constant threat.

Note: the Dutch roll motion has some similarities and some differences when compared to the uncoordinated wing rocking exercise discussed in [section 16.7](#). The latter is sometimes referred to as the Dutch roll exercise. The two patterns of motion share the concept of slipping to one side and then the other. The main difference is that the exercise calls for using the rudder to prevent any heading change, whereas true (natural) Dutch roll involves considerable yaw along with the banking and slipping.

[10.6.2](#) How to Fight Oscillations

Since this book is intended for pilots, not designers, we should discuss how the pilot should use the controls in order to oppose obnoxious oscillations.

First, bit of simple advice: in an airplane that is susceptible to Dutch roll, be extra careful to avoid uncoordinated usage of ailerons and rudder since that would unnecessarily put energy into the Dutch roll mode.

Once Dutch roll gets started (due to turbulence, or klutzy control-usage, or whatever), it may be hard to stop. In some airplanes you may be able to improve the situation as follows: If the rudder pedals are moving because of the sideways force that the Dutch roll puts on the rudder, then you should rest your feet firmly on the pedals to prevent them from moving. This will increase the stability and (more importantly) the *damping* of the yaw axis.

If that doesn't suffice, you can try to fight the oscillations by direct intervention. This requires some skill and lots of attention.

You should *not* think about correcting the *position* of the wing. If you deflect the ailerons to the right at point *D*, the wings will return to level (point *A*) sooner, but you will be applying a force in the same general direction as the velocity, increasing the velocity and the energy of the Dutch roll mode.

As we have seen, the airplane has plenty of stability and not enough damping, so what we need is a force that depends on the velocity, not the position. Therefore the ailerons need to be deflected to the left when the left wing has its maximum upward velocity, near point *A*. You should apply the deflection before point *A* and remove it after point *A*. Similarly, you should apply right aileron (smoothly) a little before point *C* and neutralize them (gradually) after point *C*.

A similar analysis applies to rudder usage. Don't try to correct the position. Instead, you need to apply right rudder at the point where the nose is swinging to the left with the maximum velocity (point *B*); by the same token you need to apply left rudder when the nose is swinging to the right with the maximum velocity (point *D*).

The same logic applies to phugoid oscillation ([section 6.1.12](#)), and to pilot-induced pitch oscillation associated with a botched landing.⁴ When the nose is high, you should not push on the yoke to correct the nose-position; you should anticipate that the position will very soon over-correct all by itself. If the nose is high and dropping (or about to drop), you need a judicious pull on the yoke to prevent the pitch attitude from overshooting.

You should tell yourself that you are fighting the velocity, not trying to correct the position. This is because you need to increase the damping, not the stability.

¹

On a retractable-gear airplane, you can lose all the oil, even the oil inside the hollow piston, more easily than on a fixed-gear airplane.

2

In a system with lots of damping and not much inertia, like a spoon in molasses, the velocity tends to be proportional to the applied force. In the other extreme (lots of inertia, little damping) we can apply Newton's second law without worrying about frictional forces — therefore the acceleration is proportional to the force and the velocity accumulates as long as the force is applied.

3

That is, a slip produces a rolling moment — by means of e.g. dihedral, sweepback, tall rudder, and/or shadow effects, as discussed in [section 9.2](#).

4

Actually it is much easier to fight pitch oscillations by direct intervention, since they happen more slowly.

11 Slips, Skids, and Snap Rolls

You should learn from the mistakes of others, because you'll never have enough time to make all those mistakes yourself.
— Ben Franklin

11.1 A Lesson on Snap Rolls

One fine spring day I was instructing a student who had about 5 hours experience. This was her first lesson in slow flight, but she was doing really well: she was maintaining the assigned altitude, the assigned heading, and the assigned airspeed (a couple of knots above the stalling speed). She was also doing a good job of keeping the inclinometer ball in the center, which required considerable pressure on the right rudder pedal because of the high power and low airspeed. I was really enjoying the flight, but suddenly I developed a feeling that there was something wrong. Gradually it dawned on me what the problem was. The problem was that the airplane was upside down.

Here's what had happened: her right foot had gotten tired, so she just removed it from the pedal — all at once. This produced a sudden yaw to the left. Naturally the left wing dropped, so she applied full right aileron. The nose was dropping, too, so she pulled back sharply on the yoke. The next thing anybody knew, we were upside down.

I took the controls and rolled the plane right-side-up. We lost about 500 feet of altitude during the maneuver. The student asked ``What was THAT?" and I said ``That was a pretty nice snap roll".

This is indeed the recipe for a snap roll: starting from a speed slightly above the stall, apply a sudden yaw with the rudder, apply opposite aileron, and pull back on the yoke. SNAP! — One wing stalls and the plane rolls over. In our case, we didn't roll exactly 180 degrees — ``only" about 135 degrees — but that's upside down enough for most people. It took a fraction of a second.

In due course the student completed her training and got her license. She's even still speaking to me. There are a number of points to be learned from this adventure:

- Although the airplane we were flying (a Cessna 152) has a reputation for being a docile aircraft, you should remember that even pussycats can bite! It is all too easy to picture the same thing happening just after takeoff, with insufficient altitude available for a recovery. This is the classic stall/spin accident that figures so prominently in the accident statistics. Don't get complacent — it could happen to *you*.
- This is why we practice slow flight and stalls a few thousand feet above the ground, and make sure there are no other aircraft nearby.
- This is why we make sure that fire extinguishers, chocks, tow bars, etc. are secure, not floating around in the back of the plane. I don't want to be picking them out of my ear during the snap-roll recovery.
- This is why we insist on maintaining coordinated flight (except for intentional slips, which are a special case). Keep the ball in its house. Don't apply right aileron without also applying right rudder. Don't apply left aileron without applying left rudder.

Let me reiterate: Piloting an airplane at low speeds requires using the rudder pedals. If you don't know how to do this correctly, you have no business trying to land, take off, or anything else.

11.2 Intentional Slips

A *slip* is performed by lowering one wing with the ailerons, and then applying opposite rudder. This is called ``top rudder" because you are pressing the rudder pedal on the same side as the raised wing.

Because the rudder is deflected, the air will be flowing somewhat *across* the fuselage. This creates much more drag than the usual streamlined flow *along* the fuselage. This is often useful for dissipating energy, e.g. for making the aircraft descend more rapidly on approach. This is always done at an airspeed well above the stall.

The crosswise flow not only creates drag (a rearward force) but also creates a sideways force that tends to change the direction of flight, as discussed in [section 8.10](#). A slip, therefore, can be viewed as a boat turn in

one direction (because of the crosswise flow), possibly combined with an ordinary turn in the other direction (because of the bank angle).

If you match the rudder deflection to the bank angle just right, no net turn results. This is called a *nonturning slip*.

Nonturning slips are used during crosswind landings, as discussed in [section 12.9](#). The upwind wing is lowered using the ailerons and the opposite rudder is used to prevent the aircraft from turning. The idea is that the bank determines the direction the airplane is *going*, while the rudder determines the direction the airplane is *pointing*. The idea is to make sure, despite the crosswind, that the direction of flight *and* the axis of the airplane are both aligned with the runway.

The definition of a slip ``to the left" versus a slip ``to the right" is a bit arbitrary and hard to remember. The following table may help. In a slip to the left:

``left" slip, matching statement


the airplane is moving toward a point that is somewhere to the left of the nose
the air is hitting the left side of the fuselage
the inclinometer ball is displaced to the left
if this is a nonturning slip, you are banked to the left

mismatching statement

the nose is pointing to the right of the direction of flight
you are applying right rudder
the slip string is displaced to the right
if you are not banked, you are making a boat turn to the right

Because of the potential for confusion, I try to avoid the term ``slip to the left" entirely. Instead, I might say ``let's make a boat turn to the right" or ``let's lower the left wing and perform a nonturning slip".

Some folks try to define the terms *side slip* and *forward slip*, but it is not worth making such a distinction. In [figure 11.1](#), depending on your intentions, you could be making a side slip to runway 11 or a forward slip to runway 9. But you could change your intentions at any point before the flare. During the approach, you are performing a nonturning slip toward the runway intersection, and that's all that need be said. Aerodynamically speaking, there is no meaningful difference between forward slip and side slip. Both are synonymous with nonturning slip.



[Figure 11.1](#): Side Slip Indistinguishable From Forward Slip

In some aircraft the airspeed indicator is grossly perturbed by a slip, as mentioned in [section 2.12.7](#).

In a typical Cessna 152/172/182, depending on the amount of slip, the airspeed can easily be off by 20%, which means the energy is off by 40%. This is enough to cause real trouble. In some less-common aircraft, you can send the airspeed needle below zero.

Suppose the static port is on the left side of the fuselage, and suppose you are in a slip to the left (the kind that requires pressing on the right rudder pedal). In this situation, the left (upwind) side of the fuselage is a high-pressure point. This high pressure cancels some of the dynamic pressure in the Pitot tube, so the airspeed indicator will lose airspeed.

Now suppose you are in a slip to the right. The static port is now on the downwind side. This will *not* be a low-pressure point. In fact, it could have almost as much high pressure as the other side, because of pressure recovery, as indicated in [figure 4.9](#). More likely, there will be relatively little pressure recovery, as illustrated [figure 4.10](#), and the static port will measure something rather close to the real static pressure.

Therefore: In a slip toward the side with the static port, expect the airspeed indicator to lose a lot of airspeed. In a slip toward the other side, expect a smaller loss.

Better yet, ignore the airspeed indicator and determine the angle of attack by looking out the window. Observe the pitch attitude relative to the direction of flight. Don't forget that because of drag caused by the slip, the new direction of flight will be angled more downward. Find a new landmark that remains a fixed

angle below the horizon.

11.3 Skids

The term *skid* denotes a particular type of slip that occurs when the airplane is in a bank and the uncoordinated airflow is coming from the side with the raised wing. Typically this happens because you have tried to speed up a turn using ``bottom rudder'', that is, pressing the rudder pedal on the same side as the lowered wing.

I will use the term *proper slip* to denote a slip that is not a skid.

If you have plenty of airspeed, the aerodynamics of a skid is the same as the aerodynamics of a proper slip. In both cases there is air flowing crosswise over the fuselage. However, you should form the habit of not skidding the airplane, for the following reason.

If the aircraft stalls, any slight crosswise flow will cause one wing to stall before the other. In particular, having the rudder deflected to the right means the aircraft will suddenly roll to the right. If the aircraft is in a 45 degree bank to the right and rolls another 45 degrees in the same direction (because you were applying right rudder pressure), it will reach the knife-edge attitude (wings vertical). If on the other hand you were holding top rudder (still holding right rudder but banking to the left this time), a sudden roll of 45 degrees would leave you with wings level (which is a big improvement over wings vertical).

If the wings are level, you can make a proper slip to the left or to the right; a skid is impossible by definition.

*** Bottom Rudder: Right vs. Wrong**

It is appallingly easy to set up a situation that leads to an unintentional skid. Suppose you are ready to make a left turn from base to final. You start the turn improperly, by applying a little left rudder. The crosswise airflow pattern acting on the dihedral of the wings will cause the airplane to bank to the left and make a relatively normal turn in the desired direction. You absent-mindedly maintain the left rudder pressure, so the bank continues to steepen. You decide to apply right aileron to prevent further steepening of the turn. That's all you need: you are in a skidding left turn, holding left rudder and right aileron, at low altitude. If you stall, you'll never be heard from again. Seriously, folks, this could happen to *you*.

Never apply more bottom rudder
than is needed to center the inclinometer ball.

There are only a few cases where bottom rudder is appropriate, for example:

1. If you are already turning to the left and use left aileron deflection to steepen the turn, you will need left rudder deflection in proportion to the aileron deflection, because of adverse yaw and yaw-axis inertia, as discussed in [section 8.8](#).
2. In a steady bank to the right during a low-airspeed, high-power climb, you may need some right rudder deflection to compensate for engine torque effects (mainly the effect of the helical propwash hitting the vertical fin and rudder, as discussed in [section 8.4](#)).
3. In a long-winged airplane in a steep turn at low airspeed, the long-tail slip effect will require you to hold bottom rudder to maintain coordination, as discussed in [section 8.9](#).
4. In a multi-engine airplane with one engine inoperative, you need to bank the airplane toward the side with the working engine, and deflect the rudder toward the same side, as discussed in [section 17.2.5](#).

Note that in all these cases you apply only enough bottom rudder to maintain coordinated flight. Do not skid!

11.4 Anticipate Correct Rudder Usage

As discussed in [chapter 8](#), there are four or five things that can cause the airplane to yaw. Your job is to use the rudders to eliminate the unwanted yaw, so that the airplane is always pointing the way it is going.

The objective is to *anticipate* how much rudder is required in various circumstances, so you aren't constantly correcting for errors.

The hardest thing to deal with is yaw-axis inertia. The rule is: rolling to the left requires left rudder; rolling to the right requires right rudder. The amount of rudder pressure should be proportional to the rate of roll. Adverse yaw complicates the situation, and requires rudder deflection whenever the roll rate does not match the aileron deflection. To a fair approximation the two effects can be covered by the rule: "rudder deflection proportional to aileron deflection".

Note that (unlike yaw-axis inertia) adverse yaw occurs even if you aren't turning. Suppose that a wind gust causes the left wing to drop. You immediately use right aileron to raise the wing. Right rudder is required. Don't get the idea that rudder is only required when you intend to turn.

Another tricky case arises when you make your first left turn after takeoff. You are holding a large amount of right rudder pressure because of the helical propwash, and you need to apply left aileron. Rather than using actual left rudder pressure, it probably suffices to use a reduction in right rudder pressure. This is harder to learn than it sounds. You may find it more convenient to *maintain* whatever right-rudder pressure is required to compensate for the helical propwash, and to make left turns by applying countervailing force on the left rudder pedal.

If you have a rudder trim control, by all means use it to compensate for the helical propwash effect.

11.5 Perceiving Slip, Perceiving Coordination

11.5.1 Looking Out the Side

To learn good coordination, first practice looking out the side. When you roll into a turn, you should see the wing go up or down like a flyswatter. If it slices down-and-forward, or up-and-backward, you are not using enough rudder.

You can control the airplane quite nicely while looking out the side. You can judge pitch attitude by the angle the wing chord makes with the lateral horizon. You can judge bank angle by the height of the wingtip above or below the horizon. And, as just mentioned, you can judge coordination by watching for forward or backward slicing motions when you roll into or out of a turn.

It is really important to be able to do this. Just for starters, there is no way you can do a decent job of scanning for traffic if you can't control the airplane precisely while looking out the side.

I even have my students do fancy things like stalls (and stall recoveries) while looking out the side.

Don't be a "gauge junkie" — the sort of pilot who can't even fly a rectangular traffic pattern except by reference to the directional gyro. When making a 90 degree turn, identify a landmark 90 degrees from your original heading and turn toward it. No gauges are required.

11.5.2 Looking Out the Front

The next step is to learn how to perceive correct coordination while looking out the front. This requires having a precise visual reference. There are several ways to arrange this.

Start by getting the airplane trimmed for straight and level flight at a reasonable airspeed, headed toward a definite point. In the figure, the plane is headed toward a point a couple of degrees to the right of the mountains.




Figure 11.2: Finger Used as Heads-Up Display

You can now use your finger as a reference, as shown in [figure 11.2](#). Rest your hand on the top of the instrument panel and align your finger with the straight-ahead point on the horizon.


Another option is to use a mark on the windshield, as shown by the red wedge in the figure. It really helps to

have a mark that falls very close to the line from your dominant eye to the aim point on the horizon, so if you can't find a scratch or a bug-corpse in just the right point, you should *make* a mark. You can use a grease pencil, a washable marking pen, a bit of tape, or whatever.

A single reference of this sort only works if your head is in the right position — wherever it was when you established the reference. Since you need to move around to look for traffic, be careful to move back into position before using the reference.

If you want to get really fancy, you can use *both* a finger and a mark on the windshield. That makes it easy to detect if your head is out of position. This also helps rule out the image from your non-dominant eye (although the easiest thing is to close that eye if it is confusing you).

[Figure 11.3](#) shows how the situation should look after rolling smoothly into a turn to the right with 30 degrees of bank.



[Figure 11.3](#): Heads-Up Display, Starting a Turn

You should use the visual reference as your primary indicator of pitch attitude and heading. Throughout the roll-in, turn, and roll-out, the rate of turn (i.e. the rate of heading change) should be proportional to the amount of bank. As the bank increases, the rate of turn should increase.

The rate of turn should be proportional to the amount of bank.

It is common mistake to think that the airplane should simply pivot around the roll axis and *then* start turning. If you look closely, you will see in [figure 11.3](#) that the sight line has already moved to the right a little. This represents the amount of turn that occurred during the roll-in. (If you roll in more slowly, this amount will increase.) Remember: The rate of turn should be proportional to the amount of bank. If the sight mark initially stands still (or backtracks!) and only later starts turning in the proper direction, it means you aren't applying enough rudder to compensate for yaw-axis inertia (and adverse yaw).

During the roll-out, the same rule still applies: the rate of turn should be proportional to the amount of bank. As the bank goes away, the rate of turn should go away. (Of course, the proportionality factor always depends on airspeed, but at each airspeed there is a definite proportionality between bank and rate of turn.) If you neglect to compensate for yaw-axis inertia, the nose will overshoot (yawing toward the continuation of the turn).

Summary: Don't let the nose backtrack on roll-in. Don't let the nose overshoot on roll-out. The rate of turn should be proportional to the amount of bank. The yaw-axis inertia and adverse yaw lead to the rule: rudder deflection should be proportional to aileron deflection.

You can see in [figure 11.3](#) why we went to the trouble of putting a mark on the windshield, rather than using, say a bolt on the cowl at the location marked by the cross in the figure (near the end of your thumb). Such an off-axis reference will not exhibit a rate of turn proportional to the amount of bank. As you can see, the problem is that the cross necessarily rotates a little ways to the outside of the properly-coordinated turn. If you tried to prevent this reference from swinging to the outside of the turn, you would be applying too much rudder during the roll-in. The amount of the error would depend on the angle between the bogus sight line and the actual roll axis — which depends on the shape of the airplane and the height of the pilot.

Once you have learned to make really good turns using the roll-axis sight mark, you should gradually learn to do without it. Make a point of imagining where the mark *would be* relative to other visual references such as the cowl, the window-frame, et cetera.

Later (after making a few hundred coordinated turns) you should be able to do it with your eyes closed, just by knowing how the controls should feel.

By the way: as you may have noticed, the sight line in [figure 11.3](#) is slightly above the horizon. This is because you need to pitch up a little bit to deal with the load factor in the turn.

Notes: (1) If you make a mark on the windshield, use a bright color, since black is too hard to distinguish from traffic. (2) The best time to make the mark is before takeoff. Taxi into position at the end of a long taxiway, and make a mark that lines up with the horizon at the far end. Even if the pre-flight mark is not perfect, it will facilitate making a better mark later.

[11.5.3](#) Using the Inclinometer Ball

The inclinometer ball will remain almost centered throughout the roll-in, turn, and roll-out if everything is done correctly. There are several reasons why you should not over-emphasize this instrument: (1) The response of the ball to coordination errors is sluggish and complex, so you have to be quite an expert to get useful feedback from it. In particular, I see lots of pilots who apply approximately the right amount of rudder, but apply it too late or too early. Diagnosing such errors using just the ball is nearly impossible; other references are more informative. (2) In general, anything that can be done by outside references should be done by outside references. (3) When rolling into or rolling out of a turn, there will be a force on the rudder which must be balanced by a horizontal component of lift (i.e. a slight bank) in order to maintain zero slip. See [section 17.2.4](#) for an explanation of why having the ball in the center is not exactly what you want (but nearly so).

The inclinometer ball definitely is helpful for providing information about a long-term slip — in particular, for telling you how much rudder trim to dial in during a high-power / low-speed climb. Especially in an unfamiliar airplane, it can be hard to tell whether one wing is down a little bit, without referring to the ball.

[11.5.4](#) Using the Seat of Your Pants

The phrase "flying by the seat of your pants" has become such a common cliché that people forget its real meaning: you can use the sense of touch in your rear end to determine whether or not your control usage is properly coordinated.¹

The idea of using your rear end as an inclinometer might sound trivial or obvious; after all, even non-pilots can notice immediately if they sit down on a park bench that is inclined. But the non-pilots are probably cheating, using their sense of sight (referring to the horizon) and their sense of which way is up (based on the acceleration-detecting organs in the inner ear). In the airplane, as you roll into a turn, the situation is much more challenging. First of all, you want the load vector (gravity plus centrifugal force) to be directed straight down into your seat (perpendicular to the wings, not to the horizon) — so visual reference to the horizon doesn't tell you what you need to know about inclination. Secondly, the organs of your inner ear are sensitive not only to the load vector but also to the rate of roll — so they don't tell you what you need to know, either.

This is a good illustration of why learning to fly an airplane is hard: the airplane is inclined (relative to the horizon) but it is not inclined (relative to the load vector). One sense (sight) conflicts with two others (inner ear and seat of pants).

Because the sense of sight is so dominant, the visual references discussed in [section 11.5.1](#) and [section 11.5.2](#) are the typically the easiest way to learn proper coordination. But you should also pay attention to what the seat of your pants is telling you. If you are being sloshed side to side as you roll into a turn, there is something wrong. It may help to close your eyes while the instructor makes a series of coordinated and uncoordinated turns.

While we are on the subject of the sense of touch: as you get experience with a particular airplane, you will learn how much force is required on the rudder to go with a certain amount of force on the ailerons (depending on airspeed, of course). Once you've got the feel of the controls, you should be able to make a decent turn without much thought or effort.

Flying by the seat of your pants may sound like a throwback to the days when airmail was carried in fabric-covered biplanes, but it is a useful technique even in modern instrument flying. Proper coordination is still important, and modern airplanes still suffer from yaw-axis inertia and adverse yaw, especially at approach speeds. As you maneuver to stay on the localizer, you don't want to be looking at the inclinometer ball — you've got too many other things competing for your visual attention.

[11.5.5](#) Intentional Slips

The previous sections concentrated on how to maintain coordinated flight. Sometimes, though, you want to perform a slip. You might want to get rid of some energy, or to align the airplane for a crosswind landing. The procedure is straightforward.

1. Make a note of the pitch attitude and direction of flight, since in some aircraft the airspeed indicator is perturbed by a slip, as discussed in [section 11.2](#). You will have to maintain angle of attack by looking at the angles themselves.
2. Make a note of which way the airplane is *going*. For a crosswind landing, maneuver so that the motion is aligned with the runway.
3. Make a note of which way the airplane is *pointing*. Call this heading *A*.
4. Using the rudder, yaw the nose to a new direction. Call this heading *B*. For a crosswind landing, choose this to be aligned with the runway.
5. Bank the airplane as required to keep it going the same direction as before.

The difference between heading *A* and heading *B* is the slip angle.

11.5.6 Slip Angle versus Bank Angle

Do not confuse slip angle with bank angle. In fact, they are perpendicular. That is, slip involves a rotation around the yaw axis, while bank involves a rotation around the roll axis, as defined in [section 19.6.1](#).

Although in a non-turning slip you can perhaps judge the amount of slip by the amount of bank, in general perceiving the bank angle is a rather poor substitute for perceiving the slip angle.

If you don't use the rudder, then

- A roll, i.e. a *change* in bank, can be a major cause of slip, because of adverse yaw and yaw-axis inertia, as discussed in [section 8.8](#).
- In contrast, a *steady* bank produces a relatively minor amount of slip, via the long-tail slip effect, as discussed in [section 8.9](#).

Using the rudder and ailerons, you can perform a wings-level boat turn, which involves a slip angle with zero bank. See [section 8.10](#).

In a twin with an engine out, you can have a turn with no bank and no slip, or a slip with no bank and no turn, or (preferably) a bank with no turn and no slip. See [section 17.2](#).

Causes of slip include:

- rudder deflection,
- aileron deflection and roll rate,
- asymmetric thrust, and
- tight turns (via the long-tail slip effect)

Causes of turn include:

- bank, and
- slip (via the boat-turn effect)

*** Coordination Exercises**

[Section 16.7](#) discusses some good coordination exercises.

11.6 Summary

- A proper slip results from applying more top rudder (or less bottom rudder) than required for coordinated flight.
- A skid results from applying more bottom rudder (or less top rudder) than required for coordinated flight.

A skid is more dangerous than a proper slip, because it is more likely to flip you upside down if anything goes wrong. Therefore, never apply excess bottom rudder (no exceptions). To say it another way, never try to speed up a turn with the rudder (no exceptions). Never try to roll out of a turn without applying coordinated rudder (no exceptions). Right aileron deflection requires right rudder deflection; left aileron deflection requires left rudder deflection.

1

In common usage, the phrase is a metaphor for any situation where the practitioner has such a good feel for the situation that quantitative information is superfluous — or where the practitioner is forced to rely on imprecise indications because quantitative information is unavailable.

12 Landing

Pilots spend a lot of time doing ``traffic pattern work" — a series of touch-and-goes. Non-pilots imagine this as being analogous to driving into a parallel-parking space, then immediately pulling out, driving around the block, and repeating the process — over and over again.

Landings involve procedures and perceptions that are just a little bit different from those involved in other phases of flight. A few of them are discussed in this chapter. (Special procedures for forced landings are discussed in [section 15.1](#).)

12.1 Planning the Approach

Use a checklist. I know a lot of pilots who fastidiously use a written checklist for preflight, but rely on memory for the approach and landing checklists. This is not good, but it is fairly easy to see how it happens: During preflight you are not strapped into your seat, and you are not busy flying the airplane. You can always take a minute to find the POH and read through it. In contrast, when you are setting up for a landing, the book is likely to be somewhere behind the back seat and you're likely to be too busy to go looking for it.

Therefore, here are some constructive suggestions. Pick the one(s) you like best:

- Make a *pocket checklist*. Print up the checklists you are going to need during flight and fold them in such a way that they fit in a pocket. During preflight, put the list in a convenient pocket.
- If you habitually use a lap desk, tape a copy of the in-flight checklists to the lap desk itself.
- If space can be found, glue a copy of the in-flight checklists to the instrument panel.

If you fly more than one airplane, make sure you have an appropriate pocket checklist or lap-desk checklist for each of them. As you progress in your pilot career, you will be flying progressively more complex aircraft, and if you persist in using the same old checklist you will get into trouble some day. Some aircraft have retractable landing gear; some don't. Some aircraft have cowl flaps; some don't. Some aircraft require using carburetor heat; some don't. Some aircraft require switching on the electric fuel pump; others forbid it.

Discipline yourself to pay attention to the checklist. Don't just keep it in your pocket as a good-luck charm.

The approach checklist should actually cover *three* things: approach, landing, and go-around. At the point where you decide to perform a go-around, you will be in no mood to go looking for a checklist.

By the same logic, by the time you are established on downwind in preparation for landing, it is not an appropriate time to be reading checklists. Therefore, the practical way to use the approach checklist is to review it *before* entering the traffic pattern. A few miles from the airport, read the checklist, think about it, and commit it to memory. Say it aloud several times if you like.¹ Short-term memory is considerably more reliable than long-term memory. Remember that the checklist is not a ``do-list"; you don't have to do each item at the moment you read it on the checklist.

* Other Planning Issues

In flight, you know you have to land sooner or later, but you should never allow yourself to get into a situation where you think you have to land on *this* runway *right now*. If you are approaching a soft, narrow, short runway with gusty crosswinds and the setting sun in your eyes, it might be a lot safer to land somewhere else. You might have to get a ride from the second airport back to the first, or you might just wait on the ground until conditions improve.


12.2 Judging Left or Right

Let's consider how things are supposed to look on final approach. One important ingredient is to be correctly lined up left/right. The task of getting lined up with a far-away object, without any intermediate guideposts, is unfamiliar to most people.

[Figure 12.1](#), [figure 12.2](#), and [figure 12.3](#) show how the runway looks if you are lined up too far to the left, perfectly on the runway centerline, or too far to the right (respectively).



[Figure 12.1](#): Lined Up Too Far Left



[Figure 12.2](#): Lined up On The Centerline



[Figure 12.3](#): Lined Up Too Far Right

The distinctions are easy enough to perceive, once you learn how. In all cases, one of the key ideas is to notice that point A lies directly above point B. That means you are lined up on the line from B to A. In particular, we see that in [figure 12.1](#) and [figure 12.3](#), you are exactly half a runway-width to one side. That is, you are lined up on one of the runway edge lines. If you continue with such an approach, you will mow down all the runway edge lights.

If while on final you perceive that you are lined up left or right of the extended centerline, you should *not* just fly directly toward the point of intended landing. Instead, you should fly over to the extended centerline *now* and then follow it to the runway. The objective is to be *traveling in the right direction* when you arrive at the runway.

As discussed in [section 12.6.2](#) and [section 12.11.1](#), you will not be able to see the runway centerline during critical parts of the flare, touchdown, and initial rollout. You need to maneuver to reference to the runway *edge*. You should start applying this skill on short final. If the runway is 40 feet wide, you should say to yourself ``I'm lined up 20 feet this side of the edge line.... I'm lined up 20 feet this side of the edge line...".

Land on the center line
by reference to the edge line.

Don't fixate on the centerline — it will disappear during the flare.

[12.3](#) Judging High or Low; Rule of Thumb

Even more important than having the left-or-right alignment is having the proper up-or-down alignment of the approach path. There are several ways to do this.


One of the worst ways is to use ``*local tricks*'', such as passing over the pond at 1500 MSL and then passing over the old red barn at 1000 MSL. Such an approach procedure doesn't work too well when you visit other airports.

The smart way to control the slope of the glide is to observe and control the slope angle directly. On an instrument approach, the electronic glideslope needle defines a 3 degree angle for you. At some airports there is a visual aid such as a VASI to define the angle for you. At most airports, though, provide no guidance at all, so you need to rely on your own perception of angles.

Most people are terrible at judging angles using the unaided eye. Therefore I recommend the following rule of thumb:

A thumb at arm's length subtends four degrees.


Specifically, the rule of thumb refers to the distance between the last joint and the end of the thumb, as shown in [figure 12.4](#).



[Figure 12.4](#): Rule of Thumb — 4 Degree Glideslope Angle

To use this rule, hold your thumb at arm's length, and arrange it so your sight line over the end of the thumb extends to the forward horizon, as shown in the figure. Then the sight line over the last joint of the thumb will be four degrees below the horizon. If this sight line extends to your chosen aim point, you know you are on a nice 4 degree glideslope.

In order to make clear the geometry of the situation, [figure 12.4](#) shows how your eye, your thumb, etc. will appear as viewed by your copilot. [Figure 12.5](#) shows how it looks from your own point of view.



[Figure 12.5](#): Rule of Thumb, Pilot's View

Note that for reasons discussed in [section 12.7.2](#), the aim point is generally not the runway threshold.


Your thumb may not be exactly the same size as mine, but if your thumb is smaller your arm is probably shorter and the angle is probably close to four degrees. In any case, you should learn what angle is subtended by your own thumb² — it comes in really handy.

Another application of this "rule of thumb" is to help perceive the destination of a power-off glide, as described in [section 15.1.4](#).

The next question is, how do you know you are actually *following* the 4 degree glideslope, as opposed to merely passing through it? Answer: as long as you remain on that glideslope, the aim point will *remain* four degrees below the horizon.

This is the correct strategy: throughout the final approach segment, your chosen aim point should remain below the horizon by the desired number of degrees.³ To say it the other way, if the angle between the horizon and your aim point is changing, then your *intended* destination is not your *actual* destination.

If the angle from the horizon to the aim point is increasing, you are going to land long; if the angle is decreasing, you are going to land short — unless you do something. The logic of this is shown in [figure 12.6](#).



[Figure 12.6](#): Landing Long or Short

The airplane in the figure is flying directly toward point X. It will overfly point A but land short of point Z. As the airplane moves from position 1 to position 2, the angle of A below the horizon increases to 90 degrees and beyond. The angle to point X remains constant, while point Z appears to move closer to the horizon.

If you are on final and perceive the aim point shrinking up toward the horizon, you probably need to add power. Conversely, if you see the angle growing (3 degrees... 3.5 degrees... 4 degrees...), you probably need to reduce power and/or increase drag.

Given that the angle shouldn't change, what sort of angle is suitable? Within the reasonable range (three to six degrees) it usually isn't critical which angle you choose. Here are the main considerations:

If you make a too-steep approach, it makes the flare maneuver more difficult and more critical. Also, some aircraft have so little drag (even in the landing configuration) that they have a hard time staying on a steep glideslope, unless they get help from a headwind.

Conversely, if you fly a too-shallow approach, you need to worry about running into obstructions. It also leaves you with fewer options in the event of an engine failure on final.

Generally, if the angle from the horizon to the aim point is less than three-quarters of a thumb (less than three degrees), you are flying a too-shallow approach. Conversely, if the angle is more than a thumb and a half (more than 6 degrees), you are flying an abnormally steep approach.

In all cases you should be extremely sensitive to *changes* in the angle, since that tells you whether you are going to land long or land short.

12.4 Judging Pitch Attitude and Angle of Attack


Now we come to the most critical task of all: you must control the angle of attack. This is important in all phases of flight, but especially so on final approach when you are intentionally rather low and slow.

12.4.1 Use Outside References and Trim

One way to maintain a definite angle of attack is to carefully perceive and control both the pitch angle and glideslope angle, as shown in [figure 12.7](#).

As discussed in [chapter 2](#), for any given flap setting the angle of attack depends on the difference between the pitch attitude and the direction of flight. Therefore if you maintain a definite value for those two angles, you are also maintaining a definite value for the angle of attack.

Trimming the airplane for the desired angle of attack and flying with a light touch on the controls is also exceedingly helpful in maintaining a definite value for angle of attack; see [section 12.12](#).



[Figure 12.7](#): Perception of Pitch and Glideslope Angles

To make sure the value in question is the *correct* value, you should look at the airspeed indicator every so often, but that should constitute only 10% of your looking. The other nine looks out of ten should be directed toward the outside, such as the angles in [figure 12.7](#).

Controlling angle of attack is even more important than controlling the left-or-right and up-or-down alignment of the flight path. If you show up at the runway slightly misaligned, or slightly long, it is usually not tragic and it is usually obvious how to solve the problem (perhaps by going around). On the other hand, if you lose control of the angle of attack, your flying career could end quite suddenly.

12.4.2 Observe and Control More Than One Thing


We see that there are three vertical angles of interest:

- the glideslope angle, i.e. the angle of the aim point below the horizon,
- the pitch attitude, i.e. the angle of the nose below the horizon, and
- the angle of attack, i.e. the angle of the aim point above the nose.

As discussed below, if you perceive and control any two of these angles, you automatically control the third.

Some pilots (especially students) try to oversimplify the situation by worrying about only *one* of the three angles. This leaves the other two angles completely uncontrolled. [Figure 12.8](#) shows three examples of what can happen if you control only one angle, namely the aim point relative to nose:

1. The lowest airplane has the aim point in the right place on the windscreen. However another angle, namely the glide slope angle, is wrong, so you hit the obstruction.
2. The middle airplane is just right. All three angles have their correct values. You got lucky.
3. The highest airplane once again starts out with the aim point just the right angle above the nose, but this does *not* mean that the angle of attack is correct, because the airplane is not actually moving toward the aim point. Another angle, namely the angle between the aim point and the horizon, is too big and (what's worse) it's changing. To keep the aim point the "right" angle above the nose, you foolishly keep pushing harder and harder on the yoke. At every point along this curved path you've got too much energy, but you don't know it because you are only watching one angle.



[Figure 12.8](#): Controlling Only One Angle

For a typical person in a typical airplane, on final approach you can easily see the aim point over the nose. If one day the nose of the airplane comes up and blocks your view of the aim point, you should notice immediately and be at least somewhat alarmed.

There are several possibilities. The most alarming ones are:

1. Possibly your pitch attitude is too high (meaning you might be about to stall).
2. Possibly you are not really moving toward your chosen aim point (meaning you are about to land long).
3. Possibly you have both problems (long and slow).

Less-disastrous possibilities include the following:

1. If you ever find yourself on approach with too much airspeed and too little altitude, it is OK to raise the nose and zoom back up to the correct glideslope. During this correction maneuver, the nose will (temporarily!) block your view of the aim point. Still, it remains a topic of concern: if the nose comes up like this, you should have a special reason, and it must be very temporary.
2. If you switch to an airplane with a longer, wider, and higher snout, it might block your view during a normal approach.
3. If you have a short torso, you might have trouble seeing the aim point even if your copilot can see it easily.⁴
4. If you use less than full flaps, it will make the problem worse.
5. An unusually large headwind will make the problem worse.

Note that the converse does not hold; maintaining a proper view of the aim point does not solve all the world's problems, as was illustrated by [figure 12.8](#). To control the airplane properly, you absolutely must perceive and control more than one angle.

In theory, you could concentrate on any two of these angles and let the third one take care of itself. On the other hand, it's not really any extra work to keep track of all three, and each one is interesting for its own special reason:

- I watch very closely the angle between the horizon and the aim point, because this one is related to energy. Since energy problems cannot be solved quickly, it really pays to notice small changes in this angle as early as possible.

- As my second angle, I watch the nose relative to the horizon. This is a strong habit. I *always* watch the nose relative to the horizon, whether I am climbing, descending, or flying level.
- The third angle (the position of the aim point on the windshield) is particularly interesting because it is related to angle of attack. It has the nice property of remaining more-or-less⁵ constant from one approach to the next, whereas the other two angles will change quite a bit depending on whether it is a steep approach or a shallow approach. Committing this angle to memory makes it possible to land without using the airspeed indicator.

Additional discussion of too-steep or too-shallow approaches can be found in [section 12.3](#).

12.4.3 Correct for Wind

There is one more ingredient in this recipe: the wind. As we shall see, in the presence of wind your direction of flight relative to the ground is not the same as your direction of flight through the air. You need to be able to perceive both.

Suppose you are on a nice 3 degree glideslope, doing 90 knots in no-wind conditions. Your direction of flight is 3 degrees below the horizon and the relative wind is therefore originating 3 degrees below the horizon. Now suppose a headwind of 30 knots springs up. You add power to remain on the 3 degree glideslope. Your flight path relative to the ground is still three degrees below the horizon, but the flight path through the windy air is only *two* degrees below the horizon.


[Figure 12.9](#) may clarify the situation. The approach commences from a point 1 mile from the runway and 300 feet up; this constitutes a 3 degree glideslope. In the absence of wind, the approach is flown as shown in [figure 12.10](#). You have 90 knots of true airspeed (90 K_{TAS}) and 450 fpm of descent rate. You will reach the runway in 40 seconds.

As shown in [figure 12.11](#), in the presence of wind you have only 60 knots of groundspeed — two thirds as much as in the no-wind case. In order to stay on the 3 degree glideslope, you must descend at two thirds of the rate. This is why you had to add power.


At the reduced groundspeed, it will take you an entire minute to reach the runway. At the end of that minute, the small hot-air balloon that is in the middle of the runway in [figure 12.12](#) will have been blown a half mile, and will meet you right at the runway threshold. Therefore your path through the air is not aimed toward the threshold, but is aimed toward the balloon. Your direction of flight through the air is only two degrees (not three degrees) below the horizon.

The relative wind is the reciprocal of the direction of flight through the air. The wing doesn't care about your groundspeed; it only cares about the angle of attack, which depends on the relative wind. To maintain the proper angle of attack the pitch attitude will be one degree higher than in the no-wind case.

Conclusion: First, you need to perceive your direction of flight relative to the ground, so you can be sure you will arrive at the aim point as intended. Second, you need to perceive your direction of flight through the air, so you can know what pitch attitude is required to give the desired angle of attack. If you are descending into a headwind, you will need less rate of descent; in any situation where you have less descent you will need less nose-down attitude.



[Figure 12.9](#): Three Degree Glide Slope



[Figure 12.10](#): Approach Without Wind (Ground View)




Figure 12.11: Approach With Wind (Ground View)




Figure 12.12: Approach With Wind; Path Through the Air

Note that the scheme of estimating the relative wind using the ratio of vertical speed to airspeed gives the correct answer even when nature's wind is blowing. As shown in [figure 12.13](#), you have a normal airspeed and a reduced VSI indication while plodding down the glideslope into the wind. See [section 2.11](#), including [figure 2.12](#), for some discussion of how this looks on the instruments.




Figure 12.13: Vertical, Horizontal Speeds Determine Angle (Side View)

12.5 Other Perceptions

Instruments: About one look out of ten, you should look at the airspeed indicator on final approach. The other nine looks out of ten, you should look outside, judging the angles as described above. During the flare, you should definitely be looking outside, not at the gauges. You want to land the airplane at a very high angle of attack. You will have to perceive the angle of attack using outside visual cues. During the flare, the airspeed indicator doesn't tell you anything about angle of attack (as discussed in [section 2.12](#)) or anything else you need to know. I once asked an experienced airline captain to tell me at what airspeed his airliner touched down. He said ``I don't know; I never looked. I always have more important things to look at''. That was a good pilot's honest answer.

Wind drift: On the base leg, you should make it a habit to check your wind drift. Normally you are being blown away from the airport, meaning that after you turn onto final you will have a headwind. If you are being blown toward the airport, watch out!

Groundspeed: It is not a very good idea to control your speed according to your impression of motion over the ground. What you care most about is angle of attack. The groundspeed perceptions will change according to

- the amount of headwind
- day versus night landing
- the model of airplane
- density altitude

Fortunately, angle perceptions remain more or less the same. Of course you should remain somewhat sensitive to groundspeed; if it seems unreasonably high you might have a tailwind. You should go around and take a good look at the windsock before trying again. Also remember that the vertical speed indications and power settings required on final approach tell you a lot about the headwind or tailwind.

12.6 Basic ``Normal'' Landing

Your Pilot's Operating Handbook should specify a ``normal'' landing procedure. It would probably be more accurate to rename it the ``basic'' landing procedure, for a simple reason: Many pilots are based at short,

unpaved, or crosswindy airports. For them, the basic procedure is definitely not their ``normal" procedure. The basic procedure should be thought of as the basis, the foundation on which other techniques are built.

In any case, here are the elements of the basic landing procedure: (1) the final approach, (2) the flare, and (3) the rollout.

12.6.1 Short Final

The main aspects of the final approach were discussed in previous sections.

12.6.2 Flare

The term *flare* refers to the part of the flight where you are raising the nose, from the nose-down attitude on final approach to the nose-high attitude at touchdown.

Throughout the flare process, raise the nose smoothly. It is a common mistake to raise the nose stepwise, that is, to raise the nose a little bit, see what happens, and then raise it a little bit more, and so forth. You should not ask yourself ``How much should I raise the nose?" It is much better to ask yourself ``At what rate should I be raising the nose?"

At each point in this process, you need to worry about three timescales: how long is it until ...

1. ... your flight path becomes horizontal
2. ... you reach the proper airspeed for touchdown
3. ... you reach ground level

Those are the three main dependent variables that are the result of the maneuver. Correspondingly, the three key independent variables that you use to control the maneuver are

1. ... the airspeed you have before starting the flare
2. ... the height at which you begin the flare
3. ... the rate at which you raise the nose

Typically you make the decisions in that order: First you pick an airspeed. That determines the height at which you must flare (the faster the speed the higher the flare). Then you adjust the rate accordingly.

In ideal conditions, you can schedule it so that all three things happen at the same time. For any given airspeed, if you start your flare at the right height and raise the nose at the right rate, you can arrange that by the time you reach ground level, you are just beginning to fly horizontally, and your attitude is just right for touchdown.

If the altitude, direction of flight, and attitude are just right, they imply that your angle of attack, airspeed, and energy are just right, too.

In less-than-ideal conditions, you should not attempt this ideal three-way timing. This is because in the real world you need to worry about wind gusts. You don't want a wind gust to come a long and rob you of your airspeed while you are still several feet above the ground, in the round part of your roundout.




Figure 12.14: Basic Landing Procedure

Therefore, in real-world conditions you should arrange that items (a) and (b) happen at the same time, and item (c) happens later, as shown in [figure 12.14](#). That is, the flare really has two parts:

- During the first part, called the *roundout*, you are making the transition from steadily descending flight to horizontal flight.
- During the second part, called *skimming*, you are moving along horizontally, a foot or less above the ground, waiting for the airspeed to bleed off so you can touch down at the proper airspeed.

Continue skimming, gradually raising the nose, until the angle of attack has increased to the point where you can land on the main wheels, with the nose wheel definitely in the air.⁶ To say it the other way, a flat, ``three-point" landing, with all three wheels making contact at the same time, is proof that your angle of attack is much too low and your airspeed is much too high.

If you find that the skimming phase lasts longer than necessary, then you started with too much airspeed and/or you began the flare too late. If you had too much airspeed on final, next time get rid of it earlier.

Every so often I get a student who thinks it is a good idea to wait until the last possible moment and then raise the nose all at once. I call this a ``square flare". Even though you can get away with this under some circumstances, it is a bad idea for the following reasons:

1. There is no margin for error. If you misjudge, and wait a little too long to perform the square flare, you will make an airplane-shaped hole in the runway.
2. It puts the instructor in an unpleasant situation. If you don't start the flare at the proper time, I can't just sit there, hoping you can see the altitude and rate of descent, and hoping you will do the square flare at the last moment. I have to take control of the airplane, which will hurt your feelings if you think it wasn't necessary.
3. The square flare doesn't work in all circumstances. Yes, you can get away with it in certain light trainers when your airspeed is much faster than your stalling speed, but in an airplane with a higher stall speed, the wing can't develop enough lift to force such a sudden change in direction of flight.
4. You can't reliably know how much to pull back. If you move to a different brand of plane, or if your plane is unusually lightly loaded, or if you fly the approach at an unusual airspeed, the square flare will go awry and you'll have no time to compensate.

There is no point in learning the square-flare technique (which will work in just a few airplanes, some of the time), when with the same amount of effort you can learn a technique that works in all sorts of airplanes, and gives a much greater margin of safety.

Remember, good pilots are judged on their smoothness, not their quickness.

In the proper touchdown attitude (in most airplanes), the nose will block your forward view. You will not be able to see the runway centerline. Therefore, during the latter part of the flare, during the touchdown, and during the initial parts of the rollout, you will have to guide the plane by reference to the runway *edge*. Otherwise, one of two things will happen: (1) If you manage to keep the centerline in view, you will touch down with much too low a pitch attitude and much too high a speed. (2) If you raise the nose anywhere near enough, you will lose sight of your reference and become an unguided missile.

If the stall warning horn comes on during the skimming phase, when you are flying horizontally a few inches above the runway, it is a good sign. You will be touching down shortly.

Conversely, if the stall warning horn comes on early in the roundout, when you are still several feet above the runway and descending, it is a bad sign. You should add power immediately. Adding power helps in two ways: (1) The power-on stalling speed is lower than the power-off stalling speed (because of the propwash over the wings). This might give you enough lift to arrest the descent. (2) The added power contributes to the energy budget, so you can rebuild your airspeed.

12.6.3 Timing the Flare

How do you recognize when it is time to begin the flare?

Let us begin by mentioning a few *unhelpful* answers to this question.

1. You could wait until you see the hair on the instructor's neck stand on end, then begin the flare. This is not good preparation for flying solo.
2. Many books suggest beginning the flare at about the height of a typical hangar. This doesn't work very well if you visit some place that has bigger hangars, smaller hangars, or no hangars at all. It also isn't very reliable at night.
3. Some people like to flare at about half the height of a typical tree. Alas, trees work even worse than hangars, for similar reasons.

4. You could wait until the width of the runway subtends a certain angle in your field of vision. This will get you into trouble if you visit some place with a wider or narrower runway.
5. You might think of using the perception of the ground rushing past, which does depend on height. Alas, this is hard to perceive, and is unacceptably sensitive to the amount of headwind.
6. You could try to use the depth perception that comes from having two eyes. However, human binocular stereopsis is absolutely useless at distances of 20 feet or greater. By the time this depth perception comes into play, it's too late. Wiley Post was blind in one eye, but that didn't prevent him from making good landings.

Here is something that actually helps: Use your sense of *timing*. At each moment on short final, ask yourself how much time t remains until you would, at the current rate, reach zero AGL. When this time t reaches the special value t_F (about two seconds), start your flare. (The exact value of t_F will depend on what sort of airplane you're flying, and other factors.)

Of course the actual flare will take longer than t_F — roughly twice as long. That's because t_F refers to what *would* happen if you forgot to flare. During the actual flare, your descent rate is reduced, so you take longer to descend.

This timing technique has some nice properties. It works on wide and narrow runways both. It works during daytime and nighttime both. It causes you to flare at a greater-than-usual height if you have a greater-than-usual vertical speed.

Now all you need is some way to perceive how much time t remains. You don't need to know the height in feet or the descent rate in feet per second; all you need is *some* quantity that perceptibly changes as you approach zero AGL. [Figure 12.15](#) shows one such quantity. The left side of the figure is what you should see when you are on final, at a definitely nonzero height. The letters $ABCD$ and $WXYZ$ represent landmarks along the side of the runway. In particular, for night landings you would use the runway lights as landmarks.



[Figure 12.15](#): Perceiving Zero Height

The important thing to notice is that the landmarks are not all colinear. In particular, BDZ is a triangle that covers nonzero area in your field of view.

Now, in contrast, imagine that you are on your hands and knees on the runway, so that your eye is just at the same height as the runway lights, about 12 inches AGL. Suppose that landmarks A and W are behind you, but you can still see the others. As shown in the right side of the figure, all the landmarks have become colinear. The erstwhile triangle BDZ has flattened out and now has zero area.

Of course you never actually fly with your eyes at zero AGL. Therefore you need to observe the *rate* at which triangle BDZ is gradually flattening out. By combining this rate perception with a sense of timing, you can decide when to begin the flare.

You can practice this perception indoors: Put a book on a table, then lower your head until the corners of the book-cover all line up.

12.6.4 Touchdown and Rollout

Don't land with the brakes applied. Of course your feet must be on the rudder pedals; just make sure you aren't accidentally depressing the brake pedals even a little bit. Wait until there is plenty of weight on the wheels (i.e., after the nosewheel is on the ground) before applying the brakes.

At touchdown and thereafter, the airplane should be sufficiently well centered that the centerline is between the main wheels. On a narrow runway you have no choice, but on a wide runway you should land on the centerline anyway. See how close you can come. Make it a matter of self-discipline and pride.

The touchdown should be gentle enough that the nosewheel *stays* in the air during touchdown and during the first 50 feet of the rollout. This is a good way of proving to yourself (and to all the kibitzers in the airport lounge) that you were in complete control of the landing. To say it the other way, if you hit with a lot of

vertical momentum, it will force the nosewheel down like a mouse trap. See also [section 12.11.7](#).

Stay in control during the rollout. Remember, the flight isn't over until the aircraft is tied down. The NTSB files are full of reports of pilots who made a decent touchdown and then (a quarter mile later) stopped paying attention and had an accident.

After you have taxied clear of the runway, perform the after-landing checklist. This will include items such as carburetor heat off, flaps retract, cowl flaps open, strobes off (for night taxiing, so you don't blind everybody), boost pumps off, et cetera.

12.7 High-Performance Landing

This section discusses the tradeoffs you must make when the field is short, obstructed, and/or plagued by gusty winds.

As always, plan ahead. In particular, you should do your short-field *takeoff* planning before landing at an unfamiliar short field, since in many airplanes it is quite possible to get into a field that you can't get out of.

The key elements of a high-performance landing are:

1. use the right configuration,
2. touch down at the right point,
3. touch down at an appropriately low airspeed, and
4. use the brakes effectively.

12.7.1 Use the Right Configuration

As discussed in [section 5.5](#), extending the flaps has five main effects:

1. Flaps decrease the stalling speed.
2. Flaps increase drag.
3. Flaps increase the incidence.
4. Flaps increase the washout.
5. Flaps perturb the trim speed.

These influence the landing in various ways:

- 1)
Having a low stalling speed is always good.
- 2a)
The typical short field is not just short, it's obstructed. Because of the obstructions, you want to make a relatively steep approach. Because of the steep approach, you might need the drag that comes with full flaps.
- 2b)
Your aim point will be not very far down the runway, so a steep approach helps keep you within power-off gliding range. If you lose power and need to glide a long ways, retract the flaps.
- 3)
Increased incidence means that (other things being equal) the pitch will be lower. (Remember: pitch + incidence = angle of attack + angle of climb.) Extending the flaps makes it easier to see over the nose but makes it harder to have the nosewheel in the air at touchdown.
- 4)
Increased washout increases roll damping so the airplane handles more nicely near the stall.
- 5)
In an ideal airplane, you would be able to make power changes and configuration changes without perturbing the trim speed. But in most airplanes, when the flaps are extended (and not otherwise), every power change affects the trim. One of my students pointed out to me that when I was flying with the flaps extended, every time I moved the throttle I simultaneously nudged the trim wheel with my thumb. I had been unaware that I was doing it, but it seems like a very sensible habit. You know

compensation is going to be needed, so why wait?

Also note that in many light aircraft, the last notch of flaps produces its full share of incidence and its full share of drag, but has only a small effect on the stalling speed. Therefore if you didn't need the last notch for energy management on final, you've got very little reason to extend the last notch at all, unless the field is very short and you need to get rid of every last knot of stalling speed.⁷

A gusty wind or a strong crosswind is a good reason using less than full flaps. Compared to full flaps, reduced flaps has the following consequences:

- For any given airspeed⁸ you will touch down at a higher pitch attitude. This means that if a gust during the "skimming" phase (after the roundout) causes you to touch down a little sooner than you intended, you will still touch down on the main wheels. This is good, because the main wheels can take a much bigger load than the nosewheel.
- By the same token: For any given pitch attitude you will touch down at a higher airspeed. In most respects, touching down at a higher airspeed is bad, but one might make the following argument: Since the sideways force of the crosswind on the fuselage is largely independent of your forward airspeed, and since your rudder authority etc. are proportional to airspeed squared, touching down at a higher airspeed gives you more authority to combat the crosswind. Therefore, if you are worried about running out of control authority, you might consider using less flaps, maybe even no flaps. The tradeoff is that even a modest increase in touchdown speed means you will use up significantly more runway. You must take this into account.
- As mentioned above, the stall speed increases. This is 100% bad. Even if you want a higher touchdown airspeed, you still would like it to be as far as possible above the stall. Remember that the effect of the flaps on the *incidence* (retract = nose-high = usually good) is different from the effect of the flaps on the *stall* (retract = bad).

Finally, while we are discussing configuration: the landing gear is an important part of the landing configuration. Please don't forget this. Double-check it on short final.

12.7.2 Touch Down at the Right Point

In the presence of obstructions, a relatively steep approach will allow you to use more of the runway. Consider for example a 50-foot tree quite close to the beginning of the runway pavement. If you use a six-degree approach slope, it will block you from using the first 500 feet of the runway. If you were to use a three-degree glide slope instead, it would block twice as much of the runway. You can get information about obstructions from the Airport/Facility Directory and other sources. Also, whenever a runway has a displaced threshold you should suspect it is displaced because of obstructions.

If your airplane requires a 1000-foot landing roll, and you are landing on a 2000-foot runway, you should arrange things so that you use the *middle* two quarters of the runway. That gives you a safety margin at each end. It doesn't make sense to put all your margin at one end or the other.

For an extreme short-field landing, your margins will be much smaller. In this case, your touchdown point will be beyond, but only very slightly beyond, the runway threshold. You must allow for the fact that your aim point will not be the same as your touchdown point, since the flare carries you forward several hundred feet beyond where the a straight-line extrapolation of your approach path would go. The correct procedure is to aim your approach path a corresponding distance *short* of the intended touchdown point. In extreme cases, the aim point may even be ahead of the runway threshold, as shown in [figure 12.16](#).




Figure 12.16: Extreme Short-Field Landing Procedure

On any runway, long or short, pick a definite spot and hit it as accurately as you can; don't just land "somewhere" down the runway. Then every landing can be practice for a short-field landing.

No matter what spot you pick, you want to land in the zone that starts at this spot and extends 100 feet or so beyond. If you have unfavorable conditions (such as gusts, wind shear, and/or an inexperienced pilot), the zone will be larger. Make sure the far end of this zone leaves enough room for the rollout, plus a safety

margin.

If the field is so horribly short that you need to choose an aim point that is near the threshold, or ahead of it, choose a glide slope that is steep enough that you can fly it without engine power. (Or, better yet, go find a more reasonable runway somewhere else. At any field where you can *depart* with reasonable safety margins, you should be able to land with considerable margin at each end. See [section 13.8.2](#).)

At any field that is not horribly short, you should make a power-on approach, and *not* aim for the very beginning of the runway. There are a couple of reasons for choosing a spot farther down the runway: (a) it gives you more obstacle clearance, and (b) if you should ever have engine trouble on final, you would have a much better chance of being able to make a power-off approach to the very beginning of the runway.

I often fly at a rather short, obstructed field: 1700 feet after the displaced threshold. We see people run off the end of the runway all too often. If you over-react to the possibility of such an over-run, you might be tempted to make an extreme short-field approach, so you would have the largest possible amount of runway ``left over''. With some skill and some headwind, you can land a Skyhawk using only the first half of the runway.

Much of flight safety depends on margins and on backup plans. At every phase of flight I ask myself how many things would have to go wrong at this point before I ran out of options.

So why put all the safety margin at the far end? What about the near end? A few years ago I decided that rather than using the first half of the runway, I would use the *middle* half of the runway.

Sure enough, a few months after making that decision, I was with a student who incapacitated the engine⁹ on 1/2 mile final. At that point we were close enough and high enough that I could glide toward the weeds as shown in [figure 12.16](#), flare, and land on the runway with several inches to spare.

12.7.3 Touch Down at a Low Speed

*** Choose the Right Runway**

The following statement may seem obvious, but it's worth emphasizing: Avoid landing with a tailwind! In an airplane that normally touches down at 50 knots, you will use up *more than twice* as much runway with a ten-knot tailwind than with a ten-knot headwind. Roughly speaking, the amount of runway consumed during rollout depends on the *square* of your groundspeed at touchdown.

Your Pilot's Operating Handbook should tell you how much runway will be consumed, as a function of headwind and other variables. Make sure your chosen runway is long enough. Include a safety margin, because the numbers in the book are based on flawless pilot technique, and you don't want to put yourself in a situation where perfection is required. Also, for reasons to be discussed below, even if you have a headwind, make sure you could safely land on the chosen runway *without* a headwind.

*** Compensate for Windshear and Gusts**

Proper management of your airspeed during a short-field approach is complicated and tricky. You have some difficult compromises to make. A low airspeed gives you the best short-field performance, but a higher airspeed gives you highly desirable protection against stalling if there is a gust or a windshear (or a lapse in pilot technique).

Your Pilot's Operating Handbook should specify the speed to use for short-field landing. This is the indicated airspeed you want to have when you begin your flare. In ideal conditions, you could trim for this speed early in the final approach leg, and maintain this speed all the way to the flare. In real-world conditions, however, the wind makes speed management much trickier.

Therefore, you need to include the following steps when planning your approach:

- Obtain a good estimate of the surface winds at the airport. This could come from the ATIS, AWOS, windsock, tower controller, other pilots, or whatever.
- Resolve the total wind into components, so you know what headwind and crosswind to expect during landing. You can use the methods of [section 14.2.2](#), but rotate your point of view so that you measure relative to runway heading, which usually differs from your current heading (since you usually plan the approach before turning onto final).

- Figure out what is the largest possible amount of airspeed that you could lose to a sudden gust or windshear on short final. (Gradual losses are no problem, and even sudden losses on long final are a relatively minor problem.) Call this amount the ``gust allowance". If it is larger than the headwind component, it means you are faced with the possibility that the headwind could shear to a tailwind, and therefore you should divert to a longer runway; you don't want to make a short-field landing with a tailwind.
- Your airspeed on short final should be equal to the short-field approach speed given in the POH *plus* the gust allowance.

On final (as always!) trim for the appropriate speed and fly with a light touch; this will greatly help you recognize when a windshear occurs, as discussed in [section 12.12](#).

If your approach speed includes a gust allowance and the expected gust does occur, then you are in good shape. Assuming you are at the right altitude and assuming you are not expecting any further windshear, you can just raise the nose and retrim. You are now flying at the handbook approach speed just as if there had been no gust and no gust allowance. The rest of the approach should be straightforward. (You typically need to make a slight power reduction, because in the absence of the headwind you will arrive at the runway sooner, so staying on the glideslope requires less power.)

On the other hand, if the gust does not occur, you will arrive at the runway with too much airspeed. Fortunately, though, if you have followed all the steps above, the gust allowance is less than the headwind component, so your *groundspeed* is less than the calm-wind short-field groundspeed, and you if you proceed to land your rollout shouldn't consume any more runway than it would in the calm-wind case.

The foregoing describes the correct procedure, in which you anticipated the windshear. Let's now consider various situations that could arise if you have forgotten to include a gust allowance in your approach speed.

1) Suppose you are flying at the handbook's short-field approach speed when a gust or windshear robs you of ten knots. If this happens on long final, several hundred feet above the ground, it is no big deal. You have lots of altitude and lots of time. You can regain your airspeed by diving about 60 feet, according to the law of the roller coaster ([section 1.2.1](#)). At this point you are on a new glide path which is 60 feet lower than the old one. This will take you to a point about 600 feet short of where the old one would have (assuming a 6 degree glide slope), but you can correct for this by increasing the power, re-intercepting the desired glide path, and then reducing the power.^{[10](#)}

2) Now suppose you suffer a similar unanticipated loss of airspeed when you are only 50 feet above the ground. In this case you have a definite problem. At this point you are on (or below) the desired glidepath and below the desired airspeed. You have a critical energy shortage. You have nothing to gain by pulling back on the yoke; if you try it you are likely to wind up as a statistic — one more ``unexplained" stall/spin accident. The proper way to deal with it is to apply full power, as discussed in [section 1.4](#). Simultaneously, dive to regain airspeed. Dive as much as you can without hitting anything, and then proceed with a go-around. Do not attempt to salvage this approach. Instead, go around and set up a proper approach, including an allowance for the windshear.

Beware decreasing headwind on final.

*** Compensate for Density and Weight**

Suppose you are flying at less than standard weight. For reasons discussed in [section 2.12.4](#), the angle of attack will be the same but the indicated airspeed will be less. The percentage change in speed should be half the percentage change in weight.

Now suppose you are landing at a high-altitude airport, where the air density is less. For reasons discussed in [section 2.12.3](#), the angle of attack will be the same and the indicated airspeed will be the same — but the true airspeed will be greater, the vertical speed will be greater, and the ground speed will be greater, by about 2% per thousand feet of density altitude. Because of the groundspeed, you will consume more runway, about 4% per thousand feet of density altitude. Your POH should contain a chart or table with more accurate information.

Note that in all cases, being able to accurately perceive the angles is a big help.

12.7.4 Use the Brakes

To stop in the shortest possible distance, the procedure is as follows:

- Touch down on the main wheels as always,
- lower the nose wheel fairly soon thereafter,
- retract the flaps,
- apply the brakes, firmly but not skidding, and then
- pull back on the yoke a little.

The reasons for these steps are as follows:

The amount of braking force that a tire can provide is directly proportional to how much weight is on the tire. As a consequence, you want to make sure there is as much weight as possible on the wheels before applying the brakes. If the nose is in the air, the wings are still supporting part of the weight of the airplane. Lowering the nose reduces the angle of attack. Retracting the flaps also reduces the angle of attack, since it reduces the angle of incidence.¹¹

A skidding tire provides much less braking force than a non-skidding tire. You never have anything to gain by allowing the tire to skid. Furthermore, skidding can very quickly lead to loss of directional control. If you think the tires might be skidding, release the brakes so they stop skidding, re-establish directional control, then reapply the brakes.

In addition the loss of braking effectiveness, skidding is very destructive to the tires — it quickly grinds away one part of the tire. The loss of rubber shortens the life of the tire, and the loss *all from one place* throws the tire out of balance. An out-of-balance tire tends to thrash itself off the pavement, reducing braking and steering effectiveness.

The idea of pulling back on the yoke during braking is simple: it increases the weight on the main wheels (which is where the brakes are). The main wheels are now supporting their normal share of the weight of the airplane, plus whatever down-force is being developed by the elevator, plus whatever share was previously being supported by the nosewheel. The idea is not to lift the nosewheel off the ground, just to bring its share of the weight *almost* to zero.

See [section 12.6.4](#) for additional discussion of the rollout, including the case of a not-very-short runway.

12.7.5 Summary: High-Performance Landing

For a short-field landing (compared to the basic landing described in the previous section) ...

- the aim point is short of the touchdown point.
- the approach is steeper.
- the airspeed is less (at corresponding points throughout the approach and roundout).
- the skimming phase is shorter or nonexistent.
- you lower the nose wheel sooner.
- you apply the brakes sooner and harder.

These points can be seen by comparing [figure 12.16](#) to [figure 12.14](#).

12.8 Soft-field Landing

If the field is soft or bumpy, it is important to touch down (1) as gently as possible, with the smallest possible vertical speed, and (2) with the lowest possible groundspeed. (In gusty-wind conditions, these two objectives are somewhat in conflict, and the first one should get priority. That is, it is better to touch down with a tiny bit of extra horizontal speed, rather than to risk ``dropping" the airplane into the mire with any appreciable vertical speed.)

The key element of soft-field technique is to use engine power during the flare and touchdown. This helps in two ways: first of all, the propwash over the wings lowers the stalling speed, meaning you can touch down at

a lower speed, and secondly, the power allows you to fly horizontally over the runway for an extended time, descending very slowly, gently ``feeling for the runway".

The approach to a soft-field is basically the same as a normal approach. The only differences are as follows:

On short final, after you are assured of reaching the field, you should extend full flaps. (They're optional on normal landings.)

Fairly late in the flare maneuver, you should add a little bit of power, just enough to maintain level flight, or a little bit less. The required amount of power is remarkably small. You are in ground effect, so there is very little induced drag, and you are moving slowly, so there is very little parasite drag. If you add too much power, the airplane will accelerate or climb, which is not what you want. You will be much too busy to look at the engine gauges during this maneuver, so use your ears: you can learn to recognize the right amount of power by its sound.

When the main wheels make contact with the ground, friction will cause the airplane to decelerate, possibly quite rapidly. This friction will also create a torque that tends to slam the nosewheel into the ground, so you generally have to pull back on the yoke to prevent this. Also, you can anticipate that the deceleration will drive your body forward (relative to the plane) at just the moment where you want to be pulling back, so tighten your shoulder harness and brace yourself.

As soon as possible after touchdown, reduce the power to idle.

As always, when taxiing on a soft surface, keep the airplane moving. If you stop, the airplane might sink in, and you will be unable to get it moving again.

During the rollout, and during taxiing on rough surfaces, it is usually a good idea to pull the yoke all the way back. The remaining airspeed and/or the propeller blast acting on the tail helps to reduce the weight on the nosewheel. This is important because (1) the nosewheel is usually more vulnerable to damage than the main wheels, and (2) more importantly, if the nosewheel drops too heavily into a pothole it could result in a prop strike.

Here's an advanced technique: if you are taxiing toward an abrupt bump, such as the edge of a piece of pavement, keep the yoke all the way back and apply a blast of power during the few feet leading up to the bump. If you do it right, in some aircraft the propwash hitting the tail will allow you to ``pop a wheelie", lifting the nosewheel almost (or perhaps entirely) off the ground. As soon as the nosewheel is over the bump, reduce the power back to idle.

If you are based at a paved airport, the ideal way to learn soft-field procedure is to fly somewhere that has a paved runway *and* an unpaved runway. Land on the paved runway, then practice soft-field taxiing and takeoffs before trying soft-field landings. This way your first experience with a soft bumpy runway comes at the lowest speeds rather than the highest speeds.

12.9 Crosswind Landing

Immediately before landing, the airplane is moving through the air, and is hardly affected by the ground. During the landing process and afterward, the airplane is moving along the ground — and is still affected by the air.

During the landing process and afterward, we want the airplane to be *moving* straight down the runway, and we also want the axis of the airplane to be *pointing* straight down the runway. These are two separate requirements; especially in the presence of a crosswind it is all too easy to have the airplane moving in one direction and pointing in another.

The usual way to meet all the requirements is to land in a slip. (An exception is discussed at the end of this section.)

Suppose for sake of discussion that the crosswind is coming from the right. Early on final approach you observe that in order to keep the airplane's *motion* aligned with the runway, the airplane's *heading* is pointed a few degrees to the right. This in normal, coordinated flight; the airplane's heading is aligned with the relative wind.

It is a very bad idea to touch down with the airplane's heading aligned with the relative wind as opposed to the runway. This will create a huge sideways force on the landing gear, and could knock the tires right off

their rims. If the tires survive, they will create a sudden large force in the direction you are pointing. This will cause the airplane to scoot off the upwind side of the runway.[12](#)

You need to change the direction you are pointing so that it is aligned with the runway, not the relative wind — and you need to do it *without* changing the direction you are going. Changing the heading is quite easily done using the rudder. You press the rudder pedal on the downwind side (the left pedal in this case). Rudder alone does not suffice, since as soon as the heading is changed the wind will strike the side of the fuselage and start blowing the airplane off course — an undesired boat turn (as discussed in [section 8.10](#)). Therefore you must lower the upwind wing; the lift vector is now tilted toward the upwind side, providing a force that counteracts the wind on the fuselage.

You are now ready to touch down. Land on the upwind wheel. Land on the upwind wheel![13](#) You should keep the ailerons and rudder deflected even after touchdown. Keep rolling along on one wheel for a while; as the airplane decelerates you will need to apply more and more aileron deflection in order to maintain the bank angle. Remember, you need that bank angle to provide the force that resists the wind.

Land on the upwind wheel.

Only after the upwind wheel has considerable weight on it should you allow the downwind wing to settle. At this point the aircraft is no longer banked. The friction of the wheels on the runway is the only force resisting the sideways force of the wind. The amount of sideways friction a tire can produce is proportional to the weight on it, which is why you must not level the wings until there is plenty of weight on the wheel(s).

At no point should you neutralize the ailerons. Keep the ailerons deflected so that the upwind wheel remains firmly planted on the runway. As airspeed decreases, you will need progressively more aileron deflection to create the required amount of force.

To reiterate, the overall sequence should be:

1. Lower the upwind wing and apply downwind rudder.
2. Land on the upwind wheel.
3. As the lift dies away, the weight of the airplane will force the other main wheel onto the ground.
4. Then you can let the nosewheel come down.

During this whole process you need to maintain pressure on the downwind rudder pedal, to counteract the weathervaning tendency ([section 8.11](#)). As soon as there is weight on the nosewheel, the nosewheel steering becomes effective, adding to whatever steering the aerodynamic forces on the rudder have been providing. Therefore at this point you can expect to suddenly need somewhat less pedal deflection.

Maintain appropriate aileron and rudder deflection during the rest of the rollout, and during taxiing as well. Remember, the flight isn't over until the airplane is tied down.

The question arises: at what point should you make the transition from coordinated flight (on final) to slipping flight (for touchdown)? Some pilots prefer to establish the slip on short final or even earlier; the idea is to have time to get the "feel" of the slip. My recommendation, though, is to begin the slip at the same time you are beginning the flare, not much earlier. The rationale is: (1) A strong crosswind is usually accompanied by a considerable headwind component, delaying your arrival at the runway, in which case an early slip is the last thing you need. It just creates drag which steals energy and aggravates the tendency to land short.[14](#) (2) The winds near the ground are never the same as the winds aloft, so any slip established on final will have to be changed during the flare anyway.

Be sure to correct for whatever crosswind is *actually* there at each point, not the crosswind you were expecting. Crosswinds are notoriously variable. As you descend and as you travel down the runway, you move in and out of the lee of trees and buildings.

If the crosswind is really strong and/or variable, you might consider using less than full flaps, as discussed in [section 12.7.1](#).

Note: There are some exceptional cases where landing in a slip is not practical. An example is the Boeing 737-400, which has huge engines mounted below the wing, and a relatively narrow wheelbase. You have to

land with the wings level; otherwise one engine would hit the ground. The trick is to maintain coordinated flight as long as possible. The direction of motion will be aligned with the runway, but the heading will not, until the very last moment. Then, use the rudder to align the heading with the runway. Deft aileron usage is needed to maintain wings level during the yaw maneuver, because of the unequal wingtip velocity. The remaining few seconds of flight will be an uncoordinated wings-level boat turn, but you hope not to turn very much. The idea is to be on the ground before the sideways force imparts any significant sideways velocity. This technique is not recommended for typical general-aviation aircraft. It's more work than necessary, and in a light aircraft the sideways velocity builds up too quickly.

12.10 Going Around

Before you begin the approach, at the time you review the landing checklist, be sure to review the go-around checklist.

If you're not prepared for the go-around,
you're not prepared for the approach.

There are many situations that call for a go-around. You should think about this in advance and establish guidelines for yourself so that you can begin a go-around *immediately* when the need arises.

If you need to go around, don't wait until the last moment. If you are rolling toward the end of the runway and are worried about running off the end into the trees, attempting a go-around will only make it worse. It is better to hit the trees when you are almost stopped than to hit the trees with almost enough energy for a go-around. An early go-around is good, but a late go-around is worse than nothing.

Here are some guidelines. You can imagine exceptions; for instance if you are flying a glider it is hard to perform a go-around. So you should come up with guidelines adapted to your situation. The point is that you should think about the go-around decision in advance. The accident records contain many examples of people who got into trouble because they spent too long deciding whether or not to go around.

- If you think you might be too high and/or too fast to land within the predetermined zone as mentioned [section 12.7.2](#), go around. Do not try to ``salvage" the approach by using extra runway.
- If your attempted landing results in a bounce, go around. Many accidents start with a harmless bounce; the salvage attempt results in running off the runway or making a disastrously hard landing.
- If there is a possibility of tailwind, make a low pass to check out the windsock, then return for landing.
- If there is a possibility of wildlife on the runway, make a low pass to scare them off, and another to make sure they are gone, then return for landing.
- If the crosswind is so strong and/or gusty that you have doubts about being able to keep the airplane on the runway, go around. Indeed, don't just go around, go away. Find another runway that is wider and/or more aligned with the wind.
- If another airplane pulls onto the runway when you are on short final, [15](#) break to one side and go around. Do not stay on the centerline; break to one side so you can keep an eye on the situation, since the other airplane remains a collision threat even while you are going around. Do not try to land farther down the runway (hoping that the other pilot will hold in position). Do not try to land short (hoping that the other pilot will take off and get out of your way).
- On an instrument approach, if you can't see the runway at certain predetermined points, a go-around is mandatory.
- If you find yourself on final with the landing gear not down and locked, or otherwise find yourself not in the correct landing configuration, go around. This is usually a sign that you didn't follow the landing checklist, and you need to take the time to re-run the *whole* checklist, to see what other things you may have missed.
- Do not wait until you are sure you need a go-around; that's the wrong question. The right question is whether you are sure things are OK for landing. If you are unsure, go around.

- If ATC says go around, go around. If your copilot or instructor says go around, go around. If you disagree, go around and argue about it later. (There are exceptions; don't let yourself get talked into attempting a dangerously late go-around.)

If ATC clears you to land, that does not prohibit you from going around. For instance, if your gear is not down, ATC would prefer to see you go around rather than land gear-up. Similarly, if ATC clears you to "land and hold short" of a runway intersection, they would prefer see you go around early rather than skid through the intersection at the last moment.

Energy mismanagement is the most-common reason for go-arounds. This is a good reason for evaluating your energy situation early and often. Ask yourself: are we high and fast, or low and slow? Fixing an energy problem is easy if you start early, but it is hard or impossible if you start late. Also remember:

An early go-around is good, but
a late go-around is worse than nothing.

When you begin the go-around, do it right. Don't add "some" power; add full takeoff power.

In a Cessna 152, 172, or 182 with flaps extended, an increase in engine power will magically re-trim the airplane for a lower airspeed, as mentioned in [section 2.3](#). This is annoying when you make small power adjustments on final approach, and downright dangerous when you apply full power for a go-around. Your first defense (which works in all airplanes) is to watch the pitch attitude; if the nose starts to pitch up, push it down with the yoke. This is sometimes quite a hefty push. (Practice simulated go-arounds at a safe altitude every so often, so you won't be surprised.)

Take a look at the airspeed indicator. Raise or lower the nose as necessary to establish the proper airspeed for the go-around.

After you have done the right thing with the power and the angle of attack, start working on the configuration. If you are carrying full flaps, remember that the last notch contributes a lot of drag but doesn't contribute much to the stalling speed, so you want to retract that notch fairly early in the process. Also, retracting the flaps part way will help with the trim problems. Don't retract the rest of the flaps until you have a reasonable airspeed margin above the stall. To the extent possible, use the trim wheel to take the pressure off the yoke. (A yoke-mounted electric trim switch comes in very handy for this.)

Make sure you have established a positive rate of climb before retracting the gear. This rule arises because in some situations you may need to perform a "bounce and go" — that is, to touch down on the runway briefly before going around. It is much nicer to bounce on the wheels.

12.11 Learning to Land the Airplane

12.11.1 Maneuver by Reference to the Edge

As mentioned in [section 12.6.2](#) and [section 12.2](#), in most airplanes, the pilot cannot see the runway centerline when the airplane is in the proper attitude for touchdown. This comes as a shock to many student pilots.

Therefore, we want to *land on the center line by reference to the edge line*.

There are several good ways to learn to do this. Repeated out-of-control attempts to land the airplane are not the recommended way.

A trick that works beautifully in typical light Cessnas (150/152/172/182)¹⁶ is the following: taxi down to the end of a disused runway (e.g. the crosswind runway) or a long taxiway that resembles a runway. Taxi into "takeoff position" and shut down the engine. You remain in the left seat, while your instructor sits on the tail, raising the nose to touchdown attitude. You should sit there for several minutes contemplating the perceptions. Compare level attitude with touchdown attitude. You will note that in touchdown attitude, you will not be able to see the centerline or the right-hand edge of the runway, but you will be able to see the left-hand edge. Especially if you move your head a little toward your side of the airplane, you should be able to see the whole sideline — from the point abeam your position all the way to the far end.

You can study these perceptions during taxi. Fortunately, all landings are preceded by takeoffs. Especially in an unfamiliar airplane, you should consciously use the pre-takeoff taxi to practice taxiing on the centerline

without looking at the centerline. That has a certain Zen ring to it, doesn't it? The trick is to taxi by reference to the taxiway edge line on your side. If the taxiway is 40 feet wide, you should concentrate on taxiing 20 feet in from the left edge. The instructor may help by holding a chart in front of your nose, forcing you to control the airplane by reference to the sideline.¹⁷ Every ten seconds or so the chart will be moved aside so you can recalibrate your perceptions.

During taxi, you should also practice perceiving height. Ask yourself, ``how far below me are the wheels?" You will need to know that when it comes time for landing.

12.11.2 Hesitation Takeoff

Make sure you have an instructor with you, especially the first time you try this. At an airport with a nice long runway, taxi into position for take-off. Pull the yoke all the way back, as you would for a soft-field takeoff. Using full power temporarily, accelerate until the nose comes up to the attitude that corresponds to stalling angle of attack or slightly less. Then retard the throttle almost to idle so that your airspeed does not increase any more. Do not let the pitch attitude or the airspeed get so high that you actually become airborne. Do not raise the nose so much that the tail hits the runway. Then just taxi down the runway in this configuration. Make a careful note of the perceptions of height, pitch, heading, and left/right position relative to the edge line.

Make sure you don't run out of runway. One option is to close the throttle, stop, and taxi back. Another option is to add power and fly away. Be careful to maintain constant pitch attitude as you increase the power. This will require releasing some of the back pressure on the yoke. This is because the increased propwash increases the effectiveness of the tail.

12.11.3 Practice Maneuvering at Altitude

The traditional (but not the best) way to learn about landing the airplane is try it again and again until it comes out right.

Landing practice has its place, of course — but it is not the only thing, or the first thing, you should do. Especially if you are learning landings for the first time, or are learning to fly a new type of airplane, there is no point in practicing defective landings over and over. That just reinforces bad habits. Also, as Langewiesche ([reference 1](#)) pointed out, landings happen so quickly that there is very little time to learn anything.

Therefore, you should leave the traffic pattern. Go somewhere where you have more altitude and fewer other aircraft. Perform the familiarization exercises as described in [section 16.8](#).

You want to spend a fair amount of time practicing slow flight. This is the sort of thing you really want to learn in the practice area, not during an attempted landing. Landing involves flying very slowly, right next to the ground. You've got no business trying to fly slowly at three feet above ground level (AGL) if you don't know how to do so at three *thousand* feet AGL.

In slow flight, in the landing configuration, make a note of the angle of attack. This is the angle of attack you want to have when you touch down on the runway. Remember the pitch attitude that goes with this angle of attack. Observe the angle the cowling makes relative to the forward horizon, and observe the angle the wingtip makes relative to the lateral horizon. Since at touchdown you will be (I hope) flying purely horizontally (i.e. negligible vertical velocity), the pitch attitude tells you everything you need to know about the angle of attack (at any given flap setting).

You will probably discover that the angle of attack you want to have on final approach is halfway between the cruise angle of attack and the stalling angle of attack. This rule of thumb is related to the more widely known rule of thumb that approach speed should be about 1.3 times the stalling speed.¹⁸

This little fact (approach angle of attack is halfway between cruise angle of attack and stalling angle of attack) is more useful than it might seem. It means you can land the airplane — and I mean an on-the-numbers, short-field landing if necessary — even if your airspeed indicator has failed (or you just can't see it because your lights have failed at night). You should not consider yourself properly ``checked out" in an airplane until you know how to do this.

[Table 12.1](#) shows some airspeeds and angles for a typical general-aviation aircraft.¹⁹

	Airspeed (KCAS)	Pitch Attitude	Incidence	Angle of Climb	Angle of Attack
cruise (clean)	115	0.0°	4.5°	0.0°	4.5°
level V_Y (clean)	76	4.0°	4.5°	0.0°	8.5°
level (flaps)	76	0.0°	8.5°	0.0°	8.5°
slower (flaps)	70	2.0°	8.5°	0.0°	10.5°
descent (flaps)	70	-2.0°	8.5°	-4.0°	10.5°
flare (flaps)	decr.	incr.	8.5°	incr.	incr.
stall (flaps)	53	12.0°	8.5°	0.0°	20.5°

[Table 12.1](#): Landing — Airspeeds and Angles

On approach, the angle of attack is distinctly not the same as the pitch attitude. Don't be fooled; bear in mind that you probably have ten or a hundred times more experience in level flight than you do in descending flight. You're not flying toward the horizon any more; you're flying toward a point several degrees below the horizon. As you transition from level flight to a four degree descent, you need to lower the nose by several degrees in order to maintain the same angle of attack.

[12.11.4](#) Practice Flaring and Stalling at Altitude

The following is a great way to learn some of the skills that you need for landing the airplane.

Choose a safe altitude (3000 feet AGL or thereabouts) and designate it as the altitude of a ``virtual runway''. Starting at an altitude 500 feet or more above the virtual runway, set up a power-off glide in the landing configuration (gear and flaps extended) at the normal approach speed. Then, about 10 feet above the virtual runway, begin a flare, so that you wind up flying level, power off, at the virtual runway altitude. As the airplane decelerates, keep pulling back, cashing in airspeed to pay for drag, maintaining altitude. Continue pulling back until the airplane stalls. Then make a normal stall recovery.

The point of this maneuver is to learn at what rate you need to raise the nose during the flare to maintain level flight.

As a variation of the above procedure, you can practice ``soft field'' landings on the virtual runway. After you have flown horizontally at the virtual runway altitude for a second or two with zero power, add enough power to sustain steady level flight. See also next section.

Practice recovering from evil zooms ([section 12.11.8](#)) and other types of defective flare ([section 12.11.9](#)).

[12.11.5](#) Practice Flying in the Runway Environment

As mentioned above, the landing flare lasts only a few seconds, and if you do a hundred landings you still have only a few minutes of experience handling a flaring airplane. Practicing slow flight at altitude is a tremendous help. Practice this. However, don't expect it to do the whole job, because (a) the airplane handles slightly differently in ground effect, and (b) you need to learn to perceive alignment with the runway, altitude, descent rate, etc. very precisely, based on visual cues in the runway environment.

Before actually trying to land the airplane, go to an airport with a nice long runway and make a few low passes at a safe airspeed.

- 1) Make the first pass about 10 feet above the runway at approach speed.
- 2) Then try it about 5 feet above the runway, at approach speed.
- 3) As you gain skill and confidence, try it about 2 feet above the runway, at approach speed.
- 4) Then try it about 1 foot above the runway, at approach speed.

During these maneuvers, you will learn to judge your height above the runway, learn to maneuver the plane so that it is centered on the runway, and learn to use the rudder (and *opposite* aileron) to get the fuselage aligned with the direction of motion even in the presence of a crosswind.

Finally, after you know how to perceive and control what is happening in the runway environment:

5)

Fly down the runway one foot or less above the surface, at a low airspeed. This will be discussed at length in the next section.

Note that it is a *very, very bad idea* to fly 10 feet or even 5 feet above the runway at a low airspeed. It is OK to stall the airplane at 3000 feet AGL, and it is OK to stall it at 0.5 feet AGL, but it is definitely not OK to drop it in from 10 feet AGL.

12.11.6 Learn Soft-Field Procedure First

After you are comfortable with high-speed flight in the runway environment, and with flaring the airplane at altitude, and handling it in the touchdown attitude, it is time for the most important exercise.

Fly the approach to a nice long runway. As you flare, advance the throttle a tiny amount. The idea is to generate enough power to allow you to fly down the runway in ground effect, a small distance above the ground. This is the soft-field landing procedure, but it works just fine on paved runways, too.²⁰ Strive to maintain one foot of altitude. You should be able to hold this altitude within a few inches. As you become more proficient, try maintaining ever-lower altitudes with ever-finer precision.

The amount of power required is very small, perhaps only 100 RPM above idle. Because the airplane is in ground effect, induced drag is greatly reduced. Because the airplane is moving so slowly, parasite drag is very small.

Gradually raise the nose to the proper touchdown attitude, and keep flying down the runway at ``zero point five AGL''. If a gust comes along and drops you the last six inches, it will be a perfect landing.

Remember to keep a careful watch on the runway edge; in the proper touchdown attitude you won't be able to see the centerline and if you persist in trying to look out the front you will wander off to one side and mow down the runway lights.

Also, keep your wits about you — don't fly the whole length of the runway and run into the trees at the end. Make a timely decision to add power and go around, or chop the power and land.

Take the time to look down at the runway, to double check your perception of height. Look at the lateral wingtip against the horizon. Get rid of the notion that the landing is something that happens at a point in time. Landing is a process that lasts a goodly amount of time.

12.11.7 Nose-High Rollout

After landing, the nosewheel is supposed to stay in the air for a while. For practice, you can make it stay in the air quite a bit longer by adding a tiny amount of power. That creates a situation analogous to the hesitation takeoff described in [section 12.11.2](#).

Even if you don't add power, try to keep the nosewheel off the ground for as long as you can (provided you've got enough runway). This has two advantages.

1. When the nose is up in the air, the airplane produces relatively high drag. This called *aerodynamic braking*. It allows you to slow down without wearing out the brakes. On the other hand, aerodynamic braking is not as effective as real brakes, so if you are approaching the end of the runway, lower the nose and retract the flaps. This puts more weight on the wheels, and therefore allows you to apply the brakes more heavily without skidding.
2. This provides additional practice handling the plane in the proper touchdown attitude. You should try to learn from every landing.

Another suggestion: You will sometimes (alas) touch down with a too-low nose attitude, so that the nosewheel hits almost immediately. If this happens, gently raise the nose to the proper attitude. Again, the purpose of this is twofold: aerodynamic braking plus a reminder of what proper touchdown attitude looks like. If this causes you to become airborne again, it means that your touchdown speed was much too high, which is a valuable lesson. Just stop raising the nose, wait half a second, and the airplane will re-land.

12.11.8 Recovering from an Evil Zoom

Consider the situation where you flare too much, too late. That is, you fly down quite near the ground and,

while your airspeed is still several knots above the stall, you pull back on the stick quite a lot. The pitch attitude will become very much nose-up. If you allow this pitch attitude to persist, the airplane will zoom up a few feet and then stall. At this point, there is no way to prevent a crash. The usual stall-recovery procedure (diving to regain airspeed) will not work. You won't be able to dive enough, because the ground gets in the way.

This is a common and very serious mistake. It is a particularly evil type of zoom. (Some other books call it ``ballooning" but that seems like an insult to all the beautiful hot-air balloons and helium balloons in the world.)

Obviously you want to stay out of situations from which no recovery is possible. The solution in this case is simple: you absolutely must observe the pitch attitude. If you see a large nose-up pitch attitude, begin a recovery immediately. Do not wait to hear the stall warning horn. Do not wait to feel aerodynamic indications of a stall. Push the nose back down to the attitude that corresponds to slow flight (roughly 15 degrees nose up in typical airplanes) and apply full power immediately. You know that the airplane can fly level at full power in this attitude, so if achieve that attitude and that power setting soon enough (before you have lost too much airspeed) you will be fine. It is important to practice this procedure, as discussed in [section 16.17.6](#).

Do not try to salvage the landing. Go around!

You cannot recover from an evil zoom simply by reversing the process that got you into trouble. During the upward zoom and the downward ``reverse zoom", the airplane loses so much energy due to drag that you will not be able to arrest the descent in time for touchdown. To say it again: If you see a bad nose-up situation and try to recover just by pushing the nose way down, the airplane will dive right into the runway nose-first. This is an example of a pilot-induced oscillation, as discussed in [section 16.3](#).

You can reduce your chance of falling prey to an evil zoom by thinking about the pitch attitude at all times. You need to control attitude in the short term, as a means of controlling altitude in the long term.

[12.11.9](#) Salvaging an Imperfect Flare

Nothing is perfect. Sometimes the flare is noticeably imperfect, yet not so bad that a go-around is required. The number of possible imperfections is enormous, but it is worth discussing how to handle the most-common cases.

Remember that for any given airspeed on short final, there will be exactly one ideal altitude at which to begin the roundout, and one ideal rate at which to raise the nose.

Scenario #1: Suppose you begin the roundout a little bit too late, and/or raise the nose too slowly during the initial moments of the roundout. You can detect this by noticing that the ground is rushing up toward you and will reach you too soon.

Solution: Raise the nose at a slightly higher rate than usual, fast enough to arrest the descent in the available time. This results in an almost-nice roundout, just a little bit squared-off. At that point you have the right altitude (a few inches) and the right vertical speed (zero), but your airspeed will be too high. If there is enough runway available, just skim along until the airspeed bleeds off, then land. On a short runway, go around — the sooner the better.

Scenario #2: Suppose you begin the flare at about the right time, but you raise the nose at too great a rate during the first part of the roundout. (This can be considered a very mild version of the evil zoom discussed in the previous section.)

Solution: If you notice this early enough, you can salvage the situation. You should temporarily stop raising the nose. Hold a constant pitch attitude for a few moments. This constant pitch attitude will *not* correspond to a constant airspeed, nor a constant angle of attack, nor a constant vertical speed. The airplane will lose energy, lose airspeed, and develop an ever-increasing rate of descent. You may think that lowering the nose is the ``obvious" way to undo the error, but you should resist the temptation; by the time you have managed to lower it, you will be at too low an altitude with too great a descent rate. Therefore, just hold a constant pitch attitude. Adding a smidge of power (a) will keep things from happening too fast, and (b) means you will have more energy at the end of the roundout. (If you add *too* much power, then at the end of the roundout you'll have more energy than you need, causing prolonged skimming as discussed in the previous scenario.) As you fly along at constant pitch attitude, at some point you will see a combination of airspeed and descent rate that you recognize from your previous normal landings. At that point, resume raising the nose at an appropriate rate.

Scenario #3: Suppose you begin the flare too early. Your first indication that something is wrong might be the following: You are flying a nice circular looping path that will be tangent to the ground; that is, you will reach zero altitude at just the time you reach zero vertical speed. However, alas, you notice that in order to do that, you are raising the nose at a rate that will lead to a stall before the roundout is completed.

Solution: Add a little bit of power. During the rest of the maneuver, raise the nose at a reduced rate. (Once again, if you add *too* much power, it could eat up a lot of runway.)

12.12 Fly with a Light Touch

As discussed in [section 2.7](#), it is vitally important to fly with a light touch. This is good practice in all regimes of flight, but it is particularly important on approach. In particular, imagine you are conducting a short-field approach, which means you've got no excess airspeed. Suppose on long final everything is just right: the right direction of flight, the right pitch attitude, the right angle of attack, the right airspeed — and in particular, the right trim.

You can — and should — confirm that you've got the right trim by letting go of the yoke.[21](#)

Now suppose that on half mile final the airplane spontaneously pitches down.

The airplane is trying to tell you something! It is trying to tell you that it lost some airspeed — presumably because of a windshear. This is a very, very common thing to happen on final. You are presumably landing into the wind, and the headwind is almost certainly stronger at pattern altitude than it is on the ground. Therefore you are virtually guaranteed to encounter a decreasing headwind during the final descent. This will rob you of some airspeed. If you are lucky, it will happen so gradually that nobody notices. If you are not lucky, it will happen suddenly. A few knots will suddenly disappear from the airspeed indicator (which you may not notice) and the airplane will want to pitch down (which it is your duty to notice).

The all-too-common temptation is to pull back on the yoke, trying to maintain pitch attitude and (vainly) hoping to maintain constant angle of descent. This is not smart.

Remember: the airplane is trimmed for a definite angle of attack. If you pull back on the yoke, you are forcing the airplane to a higher angle of attack (and a lower airspeed). Since you were already trimmed for short-field approach speed, this is definitely not a good idea.

To reiterate: the yoke is not just a control carrying commands from you to the airplane — it is also a valuable sensor carrying information from the airplane to you.

You should make sure the airplane is at all times trimmed for the right airspeed (or, rather, angle of attack). You should be aware of (and wary of) any force you apply to the yoke, forcing the airplane off its trim speed.

Additional discussion of airspeed management, including compensation for windshear, can be found in [section 12.7](#).

12.13 Critique Your Own Landings

Some of my students learn faster than others. The ones that learn the fastest are the ones who have internalized a set of high standards (and even higher goals) and who have learned to critique their own performance. These folks give me a good feeling. I know that they will continue to get better even when I'm not in the plane — a pleasant contrast to those who get gradually worse when left to themselves, and depend on the instructor to get them back in shape.

The standards for a good approach and landing are reasonably easy to remember:

- Review the checklist before the approach. Make sure you are prepared for the approach, the landing, *and the go-around*.
- Use the right configuration (flaps, landing gear, carburetor heat, cowl flaps, et cetera).
- On final, track the extended centerline. (That is, don't just fly toward the threshold at some cockeyed angle.)
- On final, maintain the correct angle of attack. Flying with a light touch will help you with this.

- On final, maintain a constant angle of descent. That is, keep the aim point a constant angle below the horizon.
- Pick a definite spot along the runway. Pretend it is the beginning of a very short runway, and try to land beyond, but less than 100 feet beyond, that point.[22](#)
- At touchdown, the nosewheel should definitely be in the air.
- The touchdown should be gentle enough that the nosewheel *stays* in the air during touchdown and during the first 50 feet of the rollout.
- At touchdown and thereafter, the axis of the airplane *and* its direction of motion should be aligned with the runway. In a crosswind, this means landing on the upwind wheel.
- At touchdown and thereafter, the airplane should be sufficiently well centered that the centerline is between the main wheels.
- Don't (accidentally or otherwise) apply the brakes until there is plenty of weight on the wheels.
- Stay in control during the rollout.

If you can do all those things, you don't need an instructor to tell you it was a good landing.

[1](#)

If you are staying in the traffic pattern, doing touch-and-goes, you have to brief the landing checklist before takeoff.

[2](#)

To calibrate your thumb, you can use the following rather specialized facts: The standard face-plate on a household light switch is 4.5 inches tall. At a distance of 5 feet, 4 inches, it subtends four degrees. Therefore, measure the switch (just to make sure) and then stand with your eye 5'4" away and see how your thumb compares.

[3](#)

By the way, you should make sure your line of sight to the threshold is unobstructed. At night, if you can't see the green threshold lights clearly, you should suspect there is an obstruction between them and the airplane. Add power!

[4](#)

Extending the flaps will help. A booster-cushion on the seat may help. A modest slip might help you see past the side of the nose. A landmark *abeam* the aim point often comes in handy. If the problem is severe, you might want to choose a different model of airplane.

[5](#)

Of course this angle is not exactly constant; it depends on flap setting, it depends on whether your seat is adjusted extra-high or extra-low, and it depends on the amount of headwind.

[6](#)

Obviously this discussion does not apply to tail-wheel type airplanes.

[7](#)

A student pilot might be tempted to always extend full flaps so that all landings would have the same incidence and would therefore look the same, but a sophisticated pilot should be able to deal with the difference between full-flap and partial-flap touchdown attitudes.

[8](#)

This applies to airspeeds that you can actually achieve, without stalling, at the reduced flap setting.

[9](#)

In a Cessna, if you pull the throttle all the way out and then *bend* it down, it's stuck at idle and you can't fix it without tools.

[10](#)

If your previous aim point was halfway down a long runway, you could just choose a new aim point 600 feet ahead of the old one, but you should feel guilty about doing so. If you ever have to choose a new aim point, you should take it as a warning of poor pilot technique.

[11](#)

On some airplanes, the flap handle is distressingly close to the landing gear handle. Make sure you grab the right one. You don't want to retract the landing gear by mistake.

[12](#)

You might think the wind would always blow airplanes off the downwind side of the runway, but more often than not they end up on the upwind side.

[13](#)

It is a common mistake among beginners to roll the wings level just before touchdown (even though they had been maintaining the correct slip up to that point) — perhaps in the effort to make it ``look like" a normal no-crosswind touchdown.

[14](#)

Also: Heaven help you if try to ``stretch the glide" by pulling back on the yoke. If you stall out of a slip you will enter a spin, and there will not be enough altitude for a recovery.

[15](#)

If that happens on long final, don't over-react.

[16](#)

It doesn't work as easily on Pipers, because they require much more force on the tail to raise the nose.

[17](#)

While you are looking out the side, don't forget to make a note of how the wingtip looks against the lateral horizon. That provides very useful pitch and bank information.

[18](#)

As discussed in [section 2.12.6](#), you must multiply the calibrated (not indicated) stalling speed by 1.3, and then convert the product to an indicated airspeed.

[19](#)

For other, similar aircraft, the numbers will be similar. In a radically different type of airplane (e.g. a jet interceptor with short, highly-swept wings) the numbers will be radically different.

[20](#)

As discussed in the previous section, do not attempt this maneuver until you are proficient at judging altitude and maneuvering in the runway environment. Do not do anything that puts you at risk of a low-altitude stall until you are within a foot or so of the runway.

[21](#)

Don't take your hand away and start scratching your ankle; just open your grip to the point where you are not quite touching the yoke.

[22](#)

On a long runway, on a day with gusty crosswinds, this is the least important of the criteria. I'm willing to compromise a little on spot-landing performance if necessary to get a soft, slow, well-aligned touchdown.

13 Takeoff

Takeoff is optional.


Landing (sooner or later) is mandatory.

The most important part of taking off is making the decision to do so. Discussion of decisionmaking ([section 13.8](#)) will be postponed until after we have discussed normal takeoffs — not because it gets lower priority, but just because it's hard to appreciate an abnormal situation unless you understand the normal situation.

13.1 Simplest Takeoff

This section presents a "case study" of a takeoff in which the pilot has to do remarkably little work. (In subsequent sections we will describe ways in which you can get better results by doing a little more work.)

This procedure applies when you have a well-paved runway with plenty of length and no obstructions to worry about. As shown in [figure 13.1](#) and [table 13.1](#), part way down the runway you rotate so that the pitch attitude is about 7.5 degrees. You then just hold that pitch attitude. Period.



	Angle of Attack	Angle of Climb	Pitch Attitude	Incidence	Airspeed
Initial roll	4.5°	0°	0.0°	4.5°	small, incr.
After rotation	12.0°	0°	7.5°	4.5°	increasing
At liftoff	12.0°	0°	7.5°	4.5°	6% below V_Y
Initial climb	decr.	incr.	7.5°	4.5°	increasing
Steady climb	7.0°	5°	7.5°	4.5°	10% above V_Y

[Table 13.1](#): Simplified Takeoff

Remarkably, at the moment of liftoff, the pilot doesn't have to do anything. The plane lifts off when it is ready, that is, when it has enough airspeed to support its weight at a 12 degree angle of attack. This will occur a few knots below V_Y , assuming V_Y corresponds to a 8.5 degree angle of attack (which is pretty typical; see also [section 2.4](#)). To construct the last phase of the scenario (asymptotic climb), I made some additional assumptions, namely that your engine is just powerful enough to provide a climb gradient of 5° at a speed 10% above V_Y . In particular, I imagine climbing out with airspeed = 83 knots and vertical speed = 735 feet per minute, in an airplane where V_Y is 75 knots. These are certainly believable numbers.

Note that before liftoff, most of the engine power is going into increasing your kinetic energy; a little is needed to overcome drag, and none is going into potential energy. Then, in the initial climb, we have a funny situation where we are climbing and accelerating at the same time. Finally, in the asymptotic climb phase, most of the power is going into potential energy; some is needed to overcome drag and none is going to increase airspeed.

The technique just described is smooth, simple, and elegant, but it has drawbacks. It does not give optimal climb performance (see [section 13.3](#)), it can cause problems if there is a gusty wind ([section 13.2](#)) or a crosswind ([section 13.5](#)), and it can cause problems if climb performance is impaired for any reason ([section 13.8.1](#) and [section 2.9](#)).


13.2 Normal Takeoff

Imagine that you are using the simplified technique of the previous section, that is, rotating early and letting the airplane ``fly itself off'' whenever it is ready. Then imagine that just after liftoff, a gust of wind comes along and robs you of a few knots of airspeed. This will cause the airplane to settle back onto the runway. This is not elegant. To get around this, use a different procedure: do not rotate until the airplane has a few knots more than the liftoff airspeed. This means that liftoff will occur right then, while you are rotating. It also means that by the time you are airborne, you can stay airborne even if you lose a few knots.

Here is a second refinement: Most runways are not perfectly smooth. If the nosewheel hits a bump at 50 knots, it is likely to knock the nose of the airplane into the air, which has several disadvantages: (1) It will cause your passengers to be bounced around more than is necessary. (2) It could cause a premature liftoff. (3) It causes unnecessary wear and tear (and possibly outright damage) to the airframe. Therefore, you can *semi-rotate*. That is, fairly early in the takeoff roll you can rotate to a pitch attitude of 3 degrees or so. This is enough to get the nosewheel slightly off the ground, but not so much that the airplane will lift off (at any reasonable speed), and not so much that the nose will obstruct your vision (in most airplanes). This semi-rotation involves a pretty tiny pitch attitude compared to, say, proper landing attitude. When the airspeed reaches V_X or thereabouts, you should lift off, by rotating another few degrees.

Finally, here is a third refinement: You know that the airplane will climb more steeply at V_Y than it will at a higher airspeed. Therefore, during the initial climb, you should watch for the point where the airplane reaches V_Y . At that point, you should make a pitch adjustment: increase the pitch attitude a small amount (another 2.5 degrees, according to the numbers in our scenario) and trim to maintain V_Y .

The whole scenario is depicted in [figure 13.2](#) and [table 13.2](#)



[Figure 13.2](#): Normal Takeoff

	Angle of Attack	Angle of Climb	Pitch Attitude	Incidence	Airspeed
Initial roll	4.5°	0°	0.0°	4.5°	small, incr.
After semi-rotation	7.5°	0°	3.0°	4.5°	increasing
Just after rotation & liftoff	12°	0°	7.5°	4.5°	just above V_X
Initial climb	decr.	incr.	7.5°	4.5°	increasing
Steady climb	8.5°	6°	10.0°	4.5°	V_Y

[Table 13.2](#): Normal Takeoff

The last phase of this scenario assumes your engine can sustain a 6 degree climb gradient at V_Y . In particular, I imagine 800 feet per minute at 75 knots.

In the figure, the dotted-line flight path and the uncolored airplane show the results you would have obtained using the simplified procedure described in the previous section. Remember that by climbing out at V_Y you gain more altitude (per unit time) than you would at any other airspeed.

* Flaps for Normal Takeoff

Extending the flaps for takeoff will improve your ability to see over the nose. This is because it increases the incidence; therefore the airplane will fly at a lower pitch attitude (for any given angle of attack). If the Pilot's Operating Handbook recommends flaps for a short-field or soft-field takeoff, there's no law against using them even when the field is long and smooth.

* Perceiving the Airspeed

Choosing an attitude and letting the airplane ``fly itself off'' as described in the previous section has the

advantage that you don't need to look at the airspeed indicator, meaning you can devote all your attention to outside references. However, this can get you into trouble if you choose the wrong attitude (see [section 2.9](#)). Airspeed, not attitude, is your best information about angle of attack ([section 2.12](#)).

At the opposite extreme, certainly it is not a good idea to devote *all* of your attention to the airspeed indicator. Fortunately, you can use your eyes (to perceive your speed relative to ground references), your ears (to perceive the sound of the engine and the sound of the wind on the airframe), and your fingertips (to perceive the forces on the yoke). This means you can get qualitative information about airspeed while keeping most of your attention focused outside. Every so often, though, you should glance at the airspeed indicator to supplement the qualitative information with quantitative information.

13.3 Obstructed-Field Takeoff

This section describes the procedure to use when you have a well-paved runway with an obstruction relatively nearby in the departure area.¹


Plan the takeoff carefully. Take into account density altitude, runway slope, headwind or lack thereof, et cetera. Make sure you know the value of V_X under these conditions, and choose a suitable rotation speed V_R as discussed below.

Use the proper flap settings, as specified in the Pilot's Operating Handbook. Here's a useful cross-check: on most light aircraft, when you extend the flaps for an obstructed-field takeoff, you will observe that the angle of the flap matches the angle of a fully-deflected aileron.

Start at the beginning of the runway. If the taxiway leads you onto the runway some distance from the beginning, you will have to back-taxi on the runway, back to the very beginning.

Open the throttle smoothly, but not so slowly that you use up significant amounts of runway before the engine reaches full power. Some people advocate using the brakes to hold the aircraft stationary until the engine comes up to full power, but this is rarely necessary; if you open the throttle properly the airplane will move only a few feet while you're doing so.²

As shown in [figure 13.3](#) and [table 13.3](#), you should choose a rotation speed V_R at or near V_X — that is, quite a bit higher than what you would use for a soft-field takeoff ([section 13.4](#)) or even a normal takeoff. The idea is to use the wheels to support the weight of the airplane until you have built up a lot of energy. It's OK to semi-rotate a little bit, to take some load off the nosewheel, but you don't want the wings to be producing significant lift until you're ready to climb away. Then rotate smoothly to the "climb-out" pitch attitude, whereupon the airplane will lift off immediately. Climb away at V_X . Trim for V_X . After you have cleared the obstruction, you can accelerate to V_Y . Finally, after you have reached a comfortable altitude, you can accelerate to "cruise climb" speed and trim again.



[Figure 13.3](#): Obstructed-Field Takeoff

	Angle of Attack	Angle of Climb	Pitch Attitude	Incidence	Airspeed
Initial roll	4.5°	0°	0.0°	4.5°	small, incr.
climb	13.0°	7°	15.5°	4.5°	V_X

[Table 13.3](#): Obstructed Field Takeoff

In the last phase of the example scenario, I imagine a climb rate of 780 fpm at 63 knots, which gives a climb gradient of 7 degrees.

In the figure, the dotted-line flight path and the uncolored airplane show the results you would have obtained

following the normal-takeoff procedure, that is, accelerating while climbing and then climbing at V_Y . Note that using by using obstructed-field procedure, you have not climbed as high, but you have better obstacle clearance because you have not flown nearly so far horizontally.

It may seem paradoxical that you get better obstacle clearance by staying on the runway *longer*, but it's true. The rationale is as follows: You want to pass over the obstruction at a reasonable altitude with a reasonable airspeed. This requires a certain amount of energy. To maximize energy you want to minimize drag throughout the maneuver. Keeping the airplane on the runway until reaching a high speed is rough on the airplane, but supporting its weight with the wheels usually involves less drag than supporting its weight with the wings. To say it another way: rolling resistance is less than induced drag, unless the field is quite soft or bumpy.

Once airborne, you want to climb at V_X for reasons discussed in [section 7.5.3](#).

The idea of choosing V_R to be equal to V_X is only an approximation. There are exceptions:

- For example, if you are facing a 20-foot-high billboard that is the only obstruction in the area, it is theoretically logical to zoom over at a speed several knots below V_X , then dive back down on the other side.³ Short-term altitude gain (as given by the law of the roller coaster) is more important than long-term rate of climb (as given by the power curve).
- On the other side of the coin, if the elevation of your departure airport is near the absolute ceiling of your airplane (so that you will have very little rate of climb once airborne) and if the runway is long and well-paved but obstructed, then it makes sense to say on the runway (or at least in ground effect) until the speed is well above V_X .

Still, for typical circumstances, choosing V_R at or near V_X is a reasonable guideline.

*** Skimming versus Wheelbarrowing or Flap-Popping**

The procedure outlined above (staying on the runway at high speed, with the flaps extended) may not be possible in your airplane. Depending on the incidence of the wings, the airplane may fly itself off well before you reach the desired rotation speed.

Usually the best way to deal with this situation is to let the airplane come off the ground, and then fly along in ground effect, rather like a soft-field takeoff.

Another possible procedure (which is usually *not* recommended) is to keep the flaps retracted until you are ready to leave the runway. Less flaps means less incidence. A big disadvantage is that ``popping" the flaps like this increases your workload at a time when there are lots of other things you should be attending to. Another disadvantage is that you run the risk of extending the flaps past the takeoff position to the landing position, creating lots of drag, which is really not what you want in this situation. If your POH calls for this procedure, go ahead, but be careful. Make sure you have some sort of detent to block inadvertent over-extension.

An ever worse situation arises if you try to keep the plane on the ground by pushing forward on the yoke. This is called *wheelbarrowing*. What happens is that while you are holding the nose wheel down, the main wheels come off the ground. You are counteracting the incidence with a negative pitch attitude. The steering becomes dangerously unstable. There is also a risk of the propeller striking the ground.

13.4 Soft-Field Takeoff

Sometimes you want to get the airplane airborne at the lowest possible airspeed, using the shortest possible takeoff roll. For example, gooey mud on the runway will cause tremendous amounts of friction on the wheels. The sooner you become airborne, the sooner you are free of that friction and the better you will be able to accelerate. Additional reasons for using soft-field procedure will be given below.

The procedure is as follows: Extend the flaps as recommended by the manufacturer; in the absence of a specific recommendation, extend the flaps so that they just match a fully down-deflected aileron. The idea is to get the most coefficient of lift without undue drag.

At the beginning of the takeoff roll, pull the yoke *fully* backward. Early in the takeoff roll, the nose will rise, as indicated in [figure 13.4](#). Allow it to rise to the pitch attitude that corresponds to the stalling angle of attack, or

slightly less. This is typically about 15 degrees nose up.




Figure 13.4: Soft-Field Takeoff

To maintain this pitch attitude as the aircraft accelerates, you will have to gradually let the yoke move forward. You will become airborne at a very low airspeed — roughly the stalling speed.⁴ If you were to maintain the liftoff attitude, a typical airplane will accelerate poorly while climbing poorly, but that's not what we want. (A lower-powered airplane might get into a situation where it can neither accelerate nor climb.) Instead, gradually lower the nose, so that you fly parallel to the ground, remaining one foot above the ground. As the aircraft accelerates in ground effect, the required angle of attack will decrease, so you will see the pitch attitude get lower and lower.

There are two ways of completing the maneuver.

- If the field is unobstructed, remain in ground effect until the pitch attitude (and angle of attack) have decreased to their normal takeoff values, as discussed in [section 13.2](#). Then climb while accelerating to V_Y just as in the normal takeoff.
- If, however, there are obstructions, it is better to remain in ground effect until the speed approaches V_X , then raise the nose and climb out while maintaining V_X as in the obstructed-field takeoff ([section 13.3](#)).

You may be surprised at how well soft-field procedure works. Just after liftoff, the airspeed is extremely low. In ordinary conditions of flight, your airplane might well have a negative rate of climb at that airspeed — yet in this case it not only maintains altitude, but accelerates. The special ingredient in this case is ground effect: a wing produces very little induced drag while it is in ground effect (that is, roughly, within one wingspan or less of the ground) for reasons discussed in [section 3.12.4](#).

Just after liftoff using this procedure,

1. there is no rolling friction because the wheels are not touching the ground;
2. there is very little induced drag because you are in ground effect; and
3. there is very little parasite drag because you are moving slowly; and
4. no power is being used for climb because you are moving horizontally.

The engine is producing full power, so if none of it goes into drag and none of it goes into climb, the airplane will accelerate like crazy.

There are many situations where this procedure is useful.

- If the runway is covered with mud, tall grass, sand, or snow, there can be troublesome amounts of friction against the wheels. Soft-field procedure allows you to transfer the airplane's weight from the wheels to the wings as early as possible, decreasing friction and improving acceleration.
- If the runway is rough and bumpy, the problem is not so much friction, but rather damage from hitting a bump at high speed. The sooner you lift off, the less harm to the airplane. Remember, the force involved in hitting a bump goes like the square of the groundspeed.
- Suppose the runway is perfectly smooth and firm, but very short — and suppose it is surrounded by open fields with lots of bumps but no serious obstructions. You can become airborne over the runway, and then accelerate in ground effect over the fields.
- Suppose you are attempting an ordinary takeoff from an ordinary field, but due to a gust (or perhaps even a lapse in pilot technique) you become airborne at a too-low airspeed. The best strategy is to accelerate in ground effect; you don't want to re-contact the runway (especially if there is a crosswind) and you don't want to try climbing at the too-low airspeed.

In all cases you must be careful to remain in ground effect until you have accelerated to a proper climb speed. If you try to climb at the liftoff speed you will have a big problem: in many cases, you will be unable to climb

out of ground effect. That is, as soon as you climb to a height where ground effect is no longer significant, the induced drag will become so large that you will be unable to climb *or* accelerate.

* **Brief the Passengers**

If you have passengers aboard who haven't seen a soft-field takeoff before, give them the courtesy of an explanation. Otherwise, they may find the procedure extremely disturbing.⁵ Just tell them you will lift off at a low airspeed and they fly horizontally for a few moments while you accelerate to the optimal climb speed. Tell them that (a) this is standard procedure for getting best performance, and (b) it minimizes jolts to the passengers.

* **Maneuver by Reference to the Edge Line**

Whereas in a normal takeoff you can guide the airplane by looking out the front, in a soft-field takeoff the nose will block your view during most of the maneuver. Therefore you must use the *edge* of the runway as your reference. Practice this skill during taxi. You will need this skill for landings and for soft-field takeoffs, but those aren't the best times to be learning it.

13.5 Crosswind Technique

There is not a "crosswind procedure" that you would use *instead* of normal procedure, soft-field procedure, or obstructed-field procedure. Rather, you use crosswind technique *in conjunction with* such procedures.

A crosswind takeoff is not as tricky as a crosswind landing, but it does call for some special care. Consider the following scenario: You are trying to take off in gusty conditions using the (over)simplified techniques of [section 13.1](#). You've already rotated, and are accelerating toward liftoff speed with the wings level. As the speed increases, the wings produce more and more lift, lightening the load on the main wheels. The wind is still blowing against the side of the fuselage as strongly as ever. The ability of the wheels to provide a sideways force to resist the wind is proportional to the downward load on the wheels.⁶ If you keep the wings level, there will necessarily come a point — prior to liftoff — where the wind overpowers the wheels and blows the airplane to the side, scraping the tires across the runway.

So, here are the correct techniques for handling a crosswind takeoff.

Regarding rudder usage: To counteract the airplane's weathervaning tendency ([section 8.11](#)), you must press on the downwind pedal to keep the plane going straight. Before rotation, both the rudder and the nosewheel contribute useful steering. After rotation, weathervaning continues, but the rudder has to do 100% of the steering. Therefore you can plan on applying a little additional pedal deflection following rotation.

Regarding aileron usage, there are two options:

1. The less-common method is more or less the reverse of a crosswind landing. During the takeoff roll, deflect the ailerons into the wind, to place more weight on the upwind wheel. The ailerons create force in proportion to airspeed squared, so at the beginning of the takeoff roll you will need *full* aileron deflection. As the airspeed builds up, gradually reduce the deflection. Rotate normally, maintaining appropriate aileron deflection, so that the downwind wing comes up while the upwind wing remains down. Keep the upwind wheel firmly planted, so that it can provide friction to resist the wind. Now the airplane is in a bank, trundling down the runway on one wheel; the sideways lift of the wings serves to counteract the force of the wind on the fuselage. As the load on the remaining wheel decreases to zero, the airplane will lift straight up.

Since the ailerons are deflected one way and the rudder another, you are commanding a slip. Indeed, the moment before liftoff you are (as desired) in a nonturning slip. The moment after liftoff you are still in this slip, but it is no longer what you need. You should level the wings and swing the nose to windward (to get the fuselage aligned with the airflow).

2. The more-common method is as follows: You deflect the ailerons into the wind, but not as much as in the previous method. The idea is not to transfer all the weight to the upwind wheel, but merely to equalize the weight, counteracting the wind's tendency to flip the airplane over onto the downwind side. To keep the wind from pushing you sideways, you keep weight on *both* wheels, delaying rotation until you have almost 100% of flying speed (rather like the obstructed-field takeoff procedure, [section 13.3](#)). You then rotate and fly away. This method is not optimal for soft or bumpy

runways (because it involves driving along the runway at high speed).

Again, immediately after liftoff you must make a heading change to establish a crosswind correction angle, so that the fuselage is aligned with the airflow.

Note that in both cases, the heading change that occurs right after liftoff is *not* a normal, coordinated turn. The motion of the center of mass is already aligned with the runway, so you do not want to change the direction of motion, just the heading. Use the rudder, not the ailerons.

After you have lifted off, you must take care not to settle back onto the runway. Since the airplane's heading is no longer aligned with the runway, re-landing would cause a severe sideways force on the landing gear.

As you climb out, you should expect that the crosswind will be stronger at altitude than it was near the ground. To compensate, make the appropriate heading changes.

13.6 Multi-Engine Takeoff

In a multi-engine airplane, an engine failure shortly after takeoff is a very critical situation. It places considerable demands on the pilot. Make sure you know what to do; brief yourself in detail before the takeoff. Engine failures and related procedures are discussed in [section 17.2](#).

Early in the takeoff roll, verify that both engines are developing the same amount of power. If the aircraft is trying to pull to one side, you've got a problem. Also, check the engine gauges to make sure (a) you've got the normal RPM on both engines, (b) you've got the normal manifold pressure on both engines, and (c) you've got the normal fuel flow on both engines. The instruments that measure these three quantities are usually a single gauge with two needles, so if you notice that the needles are *split* you've got a problem.

If anything funny happens while there is runway remaining ahead of you, close both throttles immediately and stop straight ahead. Even if you are airborne, close the throttles and re-land if there is sufficient runway length. Indeed, even if the remaining runway is not quite enough, you might want to land on it: Suppose that because of density altitude or whatever, your aircraft has poor single-engine climb performance. You will sustain vastly less damage if you land and run off the end of the runway at low speed, rather than making an unsuccessful attempt to climb out on one engine.

You really don't want to be airborne at a speed below V_{MC} , i.e. at a speed where you can't maintain directional control on one engine. In many aircraft, you should aim for a lift-off speed of V_{MC} plus 5 knots. To make sure you do not lift off too soon, you can delay rotation until reaching V_{MC} . You can semi-rotate earlier if you want; just make sure you don't rotate to a pitch attitude that will cause liftoff below the desired airspeed. After liftoff, climb while accelerating to V_Y (which ought to be greater than or equal to V_{YSE}).

In many twins, V_{MC} is essentially equal to the stalling speed. In others, however, it is considerably higher, which makes soft-field takeoffs problematic. You don't want to lift off at "the lowest possible airspeed" (like you would in a single) since if you lost an engine at that speed you'd have a big problem: uncontrollable yaw. It would be a lot safer to lift off at V_{MC} or higher, even if this means staying away from soft, bumpy fields.

13.7 Other Elements of the Takeoff

At a tower airport, you will need to get approvals and clearances before taxiing or taking off.

During the takeoff roll and climb-out, you will need to apply right rudder to compensate for the helical propwash, as discussed in [section 8.4](#).

In an aircraft with retractable landing gear, you have to decide when to retract them. It is *not* a good procedure to retract them the instant you become airborne. The reason is that sometimes things go wrong in the first seconds after liftoff, and you don't want to foreclose the option of re-landing on the remaining runway. Therefore the usual procedure is to retract the gear when it is no longer possible to re-land on the departure runway. You should say aloud the checklist item: "No more useful runway; gear coming up".

On a really, really long runway, it's OK to reduce drag by getting the gear up somewhat before you've flown all the way down the runway. However: (1) it's usually not worth the trouble, and (2) make sure that you're high enough that, in the event you *do* want to land immediately, you have time to re-extend the gear.

Very early in the climb, pick a landmark somewhere a few miles along your intended flight path, so you can

maintain direction of flight primarily by outside references. The upwind leg of the traffic pattern is supposed to be an extension of the runway centerline. Similarly, note the pitch angle relative the horizon, so you can maintain the proper angle of attack and detect any windshear. You can cross-check direction, pitch angle, and angle of attack using the directional gyro, horizon gyro, and airspeed indicator, but you don't want to spend more than a tenth of your time looking at gauges. You need to be looking outside to check for traffic.

When ATC gives you a takeoff clearance, it means you are supposed to have the runway to yourself. This applies to the runway itself, not to the airspace, so as soon as you are airborne, you are 100% responsible for seeing and avoiding other traffic. Even on the runway, it pays to keep your eyes open; there's always a chance that ATC has made a mistake, and an even bigger chance that some other pilot has pulled onto the runway without a clearance.

Upon reaching a comfortable altitude, say 500 feet AGL, there are a number of things that might need doing: If your aircraft has cowl flaps, check them. On a normal takeoff they will already be open, but on a go-around you will have to open them. This is also be a good time to throttle back to normal climb power, which is less than takeoff power on most aircraft with controllable-pitch propellers. This is also a good time to retract any remaining flaps. Finally, this might be a good time to accelerate from V_y to a nice cruise-climb speed.

You should not mess with the cowl flaps or other items until you are several hundred feet up. Turbulence might cause a pitch or bank excursion while your attention is distracted, or you might bump the yoke. At low altitude, basic aircraft control should get your undivided attention.

In some aircraft, the fuel-boost pumps should be turned off at 1000 AGL; in other aircraft they stay on throughout the initial climb. Other aircraft don't use boost pumps at all.

13.8 Decisionmaking

The most important thing a pilot can do to promote aviation safety is to know when to leave the airplane tied down. Don't pressure yourself — or let others pressure you — into making a questionable flight.

I advise all my passengers explicitly:

A flight can be delayed or diverted for many reasons, including weather, mechanical trouble, pilot fatigue, et cetera. If you feel they have to go or return at a particular time, you should make alternate arrangements. Different takeoff situations call for different takeoff techniques. You have to ask yourself:

- Should I be making this flight at all?
- Is the runway long and unobstructed, or is it short and obstructed?
- Is the runway firm and smooth, or is it soft and bumpy?
- Is there a significant crosswind?

Use a takeoff checklist that is appropriate to the particular aircraft you are flying (not a generic substitute). Some airplanes require the fuel boost pump on for takeoff, while others require it off. A C-152 requires 10 degrees of flaps for short-field takeoff, while a C-172 requires zero.

13.8.1 Monitoring Takeoff Performance (wrong)

I sometimes hear statements such as these:

- Statement #1 (wrong): ``On any runway, if you have attained 70% of your takeoff speed before you have used up 50% of the runway, then you will have 100% of your takeoff speed by the end of the runway."

People even claim to ``prove" statement #1, using physics plus a number of hare-brained assumptions, including:

1. Assuming friction is negligible. In fact, friction is much more important in the second half of takeoff roll.[Z](#)
2. Assuming the engine puts out constant thrust. Although constant thrust might be a fair approximation for jets or rockets, for piston engines (especially ones with constant-speed props) constant *power* is a better approximation. Therefore we expect considerably less thrust in the second half of the takeoff roll.
3. Assuming zero wind. This might be true sometimes, but it's certainly not safe to assume this in

general. With a strong enough headwind, you can attain 70% of flying speed with no engine power at all.

The following modified version is also wrong, and even more dangerous:

- Statement #2 (wrong): "On any runway, if you have attained 70% of your takeoff speed before you have used up 50% of the runway, then the takeoff will be successful".

A little thought shows this cannot possibly be correct in general. It cannot even be repaired by changing the percentages. As shown in [figure 13.5](#), consider a very, very long runway and a density altitude slightly above the airplane's absolute ceiling. You will be able to reach 100% of flying speed before you have used up even 10% of the runway. You will be able to take off and climb a few feet, but you will never be able to climb out of ground effect, no matter how long the runway. Therefore: forget any X% — Y% rule you may have heard.



[Figure 13.5](#): Takeoff Failure despite plenty of airspeed and runway

[13.8.2](#) Monitoring Takeoff Performance (right)

Suppose that you are on your takeoff roll, and several subtle things have gone wrong: (a) you have underestimated the density altitude; (b) for various reasons (see below) the engine is only producing 80% as much power as it should, even at this altitude; (c) the parking brake is partially stuck so the brakes are dragging; (d) you didn't notice a shift in the wind, so you now have a few knots of tailwind; (e) you didn't notice that the runway has a slight up-slope; and (f) your mother-in-law has stowed away in the back seat, so the airplane is 15% heavier than you planned for. You may not be able to complete the takeoff safely. The question is, can you somehow notice the performance deficit in time to abort the takeoff?

If you are familiar with the airplane, you should know how the engine is supposed to sound; if it sounds rough, have it checked. Similarly, you may know what engine RPM to expect early in the takeoff roll; if you get less, abort the takeoff and investigate.

Unfortunately, if you are not intimately familiar with the airplane, it can be very difficult to notice a performance deficit until it is too late. Careful planning and checking is required, as we shall see.

Using the Pilot's Operating Handbook (POH), calculate the takeoff ground roll distance that is expected for your takeoff conditions. Also calculate the landing ground roll distance for the same conditions. Choose a runway that is at least as long as the two distances *combined*, plus a suitable margin for error. Observe and note well what part of the runway should be consumed by the takeoff roll.⁸ Then commence your takeoff. If you are not airborne by the predicted point, close the throttle and apply the brakes immediately. Taxi back to the hangar and figure out what's wrong.

Do not attempt to use "extra" runway length to salvage the takeoff if there is a significant performance deficit. If you've got a deficit, you should figure out why, and the takeoff roll is no place to be doing complicated figuring.

Now let's consider the annoying situation where the available runway is just a little shorter than the aforementioned "takeoff plus landing" ground roll distance. The POH tells you that a takeoff should be possible, if everything goes right, but it does not tell you how to make a timely determination that you've got a problem. In such a situation, there are three possibilities. One is to change the situation; that is, you can offload some fuel, toss out some payload, wait for cooler air, and wait for more headwind — so that you can attempt a takeoff using the procedure described in the previous paragraph. The second possibility is to figure out how much runway your airplane should consume reaching various speeds *less than* flying speed, so that you can have earlier opportunities to abort the takeoff. This is a job for a test pilot; the typical POH does not provide such information, and takeoff performance is notoriously hard to predict accurately. Please do not try this; playing "amateur test pilot" is like playing Russian roulette. The third possibility, if you have any remaining doubts about your airplane's performance, is to stay home.

[13.8.3](#) Causes of Diminished Power

There are dozens of things that could go wrong with an aircraft engine.

- One of the exhaust valves could be burned or stuck, so it won't fully close.

- One of lobes on the camshaft could be worn, so a valve won't fully open.
- The magneto timing could be not quite right.
- There could be a bird's nest in the air intake.
- et cetera.

Such problems are not particularly rare; I have personally experienced the first four items in this list.

If some such thing goes wrong, the engine will usually *not* stop cold. It will continue to run, producing a fairly large percentage of its normal power. In flight, this resilience is clearly an advantage.

During takeoff, this resilience is a two-edged sword. Because the engine continues to develop lots of power, you might not notice the degradation. You might be tempted to take off with such an engine. This could lead to big trouble, especially on an obstructed-field takeoff.

13.8.4 Plan & Practice Rejected Takeoffs

There are many types of problems that you may not notice until you have begun your takeoff roll. Early in the takeoff roll, scan the airspeed, engine RPM, manifold pressure, and fuel flow to make sure you're getting reasonable readings.⁹

You should *always* plan your takeoff. This includes planning for a rejected takeoff, for reasons discussed in [section 13.8.2](#).

Be sure you practice this. The first few times the rejected-takeoff situation arises, your expectation of a normal takeoff will be so strong that it is difficult to accept the situation and close the throttle. A rejected takeoff is psychologically at least as difficult as a go-around. Actually, most single-engine pilots find it *more* difficult than a go-around, if only because it isn't given as much emphasis during training.

Instructors: here's an instructional technique: During preflight, brief the student on the procedures for rejected takeoff. Choose a runway that is plenty long. During the takeoff roll, wait until the airspeed is about half of the liftoff speed. Then slap a suction cup on the airspeed indicator and say, ``simulated airspeed indicator failure''.

If something seemingly minor happens early in the takeoff roll, reject the takeoff. The rationale is that during the takeoff roll you don't have time to make an intelligent decision about what is minor and what is not, so (assuming there is plenty of runway remaining) the safest thing is to stop now and think later. See also [section 15.3](#).

13.9 Summary

Four of the most-common takeoff procedures are related in a fairly logical way, as summarized in [table 13.4](#).

	Unobstructed	Obstructed
Well-paved	Semi-rotate early. Fully rotate at V_R . Climb while accelerating to V_Y .	Rotate at V_X . Climb at constant airspeed: V_X .
Soft	Hop into ground effect just above V_S . Accelerate horizontally (1 foot AGL) to V_R . Climb while accelerating to V_Y .	Hop into ground effect just above V_S . Accelerate horizontally (1 foot AGL) to V_X . Climb at constant airspeed: V_X .

[Table 13.4](#): Basic Takeoff Procedures

Additionally, in each of the four cases, you must take into account the crosswind if any.

Proper planning is important. A wise ``no-go" decision could save you a lot of trouble. Make sure you know the proper procedures, including the critical airspeeds. Make sure you know how much runway you will need. If, during the takeoff roll, it looks like you are getting less performance than you should, stop and figure out what's wrong. Practice rejected takeoffs.

Make sure you know what angle of climb you should expect. You need this to check obstacle clearance. This

also affects your choice of initial pitch attitude.

When choosing an initial pitch attitude, remember that pitch attitude is not the same as angle of attack. See [section 2.9](#) for information on the right (and wrong) ways to handle cases where the correct pitch attitude differs from what you expected.

Keep the aircraft properly trimmed and fly with a light touch. Don't forget the after-takeoff checklist.

L

[1](#)

In your Pilot's Operating Handbook, this is probably called ``short-field takeoff'. However, as we shall see, this is definitely not the right procedure for a short unobstructed field — it actually uses more runway than a normal takeoff. If you have a really short but unobstructed field, consider using soft-field procedure ([section 13.4](#)).

[2](#)

I like to avoid running the engine at high power when the airplane is not moving at all, since this tends to suck up rocks, damaging the propeller. If you are moving, by the time the rock gets off the ground you will be somewhere else, possibly escaping damage.

[3](#)

I wouldn't do this except in an emergency, because it would imply operating without adequate safety margins.

[4](#)

... but you shouldn't be looking at the airspeed indicator. It doesn't provide any useful information at these speeds.

[5](#)

Imagine how it looks: The airplane is airborne but not climbing, and you are flying directly toward the bases of the trees at high speed. Just when they're convinced they're about to die, you pop the nose up and climb out.

[6](#)

Most types of friction behave this way.

[7](#)

Indeed, if friction were negligible, airplanes would fly much faster and would use much less fuel.

[8](#)

There are several good ways to do this. (a) Some runways have standard markings every 500 feet. (b) Sometimes the required takeoff distance is a half or a third (or some other convenient fraction) of the total runway length. (c) Sometimes you can pace off the distance between runway lights, and then count lights. (d) Sometimes you just have to pace off the whole distance

[9](#)

Other problems you might notice during takeoff include: A door that is not properly latched may pop open as the airspeed builds up. A seatbelt hanging out can cause a very loud, aperiodic banging noise. Neither of these is aerodynamically serious, so don't over-react.

14 Cross-Country Flying

I'm not lost, I'm just uncertain of my position.

The term *cross-country flying* refers to essentially all flying that takes you beyond the immediate vicinity of the airport.

In cross-country flying, a number of basic skills assume added importance. For example,

- When you stay near your home airport, you can land and refuel whenever you want, but during a cross-country flight you need to plan ahead.
- When you stay near your home airport, you can land immediately if threatening weather moves in, but during a cross-country flight you need to do a lot more planning and a lot more en-route double-checking.
- When you stay near your home airport, you presumably know the length of all the runways and the layout of the traffic pattern, but it can be highly embarrassing to show up at another airport and turn left base when everybody else is using a right-hand pattern. It is also embarrassing to land a little long and a little fast and then discover that the runway is very short.
- And last but not least, you need good navigation. Navigation involves keeping track of where you are and finding your way to the destination. The three primary methods of navigation are pilotage ([section 14.1](#)), dead reckoning ([section 14.2](#)), and navigation by instruments ([section 14.3](#)).

14.1 Pilotage

The term *pilotage* refers to finding your way by reference to landmarks. This is a basic yet important pilot skill.

From the air, things look different than they do from the ground. It will take you a while to learn aeronautical pilotage skills. The rest of this section covers miscellaneous small hints.

14.1.1 Airports Make Good Landmarks

When you are planning a cross-country trip, it is advantageous to plan a route that passes over airports along the way. They make great checkpoints.

Airports make good landmarks.

If you fly over an airport, it is hard to mistake it for something else. Indeed, many airports have their name printed on one of the taxiways in twenty-foot-high letters, which pretty much eliminates all doubt as to where you are.

Even if you are not using the airports as navigational references, it is a great exercise to practice spotting all the little airports along the route. This is not easy; it is an acquired skill. Airports with grass runways can be particularly challenging, since it is hard to distinguish them from the surrounding fields. Hint: Look for the airplanes. If you see lots of airplanes parked on the grass, there's probably a runway nearby.

If you stumble across an airport that doesn't correspond with where you think you are on the chart, it probably means you are off course, but not necessarily. That's because some private strips and military fields are intentionally omitted from the charts.¹

Spotting airports at night is sometimes a challenge. Non-pilots often have the impression that airports ought to be brightly lit, but in fact they are not. An airport in the middle of a town will be about the darkest thing in town.

Major airports have fairly bright runway edge lights, but the lights are highly directional, so unless you are near the final approach course you may be unable to see them. Also note that the tower has control of the runway lights, and may well turn off all the lights on whatever runways are not being used at the moment.

Most airports have rotating beacons that flash white and green, alternately. However, it is surprising how

many airports have no beacons, inoperative beacons, or beacons that are so dim as to be useless.

Airport-spotting skill might come in very handy if you ever need to make a landing on short notice.

14.1.2 One-Dimensional Landmarks

Some landmarks (like airports, small towns, small lakes, etc.) are essentially point-like (zero-dimensional). Other landmarks (highways, railroads, coastlines) are essentially one-dimensional. In the latter case, you can readily see that you are *somewhere* along the landmark, but you will need additional information to know *where* along the landmark you are. Suggestion: the intersection of two one-dimensional landmarks makes a fine waypoint.

In forested areas, highways and railroads have the additional problem that you may not be able to see them unless you are nearly overhead.

14.1.3 Choose Distinctive Landmarks

In parts of the world where there are relatively few lakes and rivers, they make good landmarks. In other parts of the world, there are so many lakes and rivers that it is distressingly easy to misidentify them.

Similar words apply to highways: if there are a lot of them, their usefulness as landmarks is impaired.

14.1.4 Doglegs

When planning your first few cross-country trips, rather than planning to make a beeline from departure to final destination, plan a dogleg course that passes directly over a goodly number of airports and other landmarks along the way.

In general, if there is a long stretch without a 100% obvious landmark, plan a dogleg so that there is. Especially on hazy days, this simplifies life.

Even a rather crooked dogleg (say, 20 degrees off the beeline heading) adds only a few percent to the length of the trip.

14.1.5 Reality-Based Navigation

When you are at home, *planning* a flight, it makes sense to look at the chart and try to pick out a set of convenient, conspicuously-charted objects. This is called map-based navigation: you go from the map to the reality.

On the other hand, when you are in the plane, it makes a lot of sense to reverse the process: Look out the window and find some conspicuous object, and then see if you can find it on the map! This is called reality-based navigation: you go from the reality to the map.

14.2 Dead Reckoning

The term *dead reckoning* refers to navigating by keeping track of time, rate of travel, and direction of travel. To do a good job of dead reckoning, you need three instruments:


- watch or clock,
- airspeed indicator, and
- compass.

In addition, you will need an estimate of

- wind speed, and
- wind direction.

Before discussing the theory of this, let's do an example. Let's suppose you are airborne at 5000 feet, cruising at 110 knots (indicated airspeed) on a heading of 090 degrees. At 32 minutes after the hour, you arrive over Hackettstown, New Jersey, and your next checkpoint is Sussex, New Jersey. The ``winds aloft" forecast called for winds of 335 degrees at 25 knots. You need to know what heading to fly and how long it will take to reach the next checkpoint. The calculation that follows is a rough estimate that you can do in the cockpit. (Later on we'll see how to do more exact calculations in the peace and quiet of the flight-planning room.)

First of all, note the time. Write it on the chart near Hackettstown, as exemplified by the red ``:32" marked on the chart in [figure 14.1.2](#) While you are there, draw a line from there to the next waypoint (Sussex). This line, too, is shown in red in [figure 14.1](#). Look outside, checking for traffic.



[Figure 14.1](#): Course Line from Hackettstown to Sussex

14.2.1 Course

Next, you should estimate the *course* from your present position to the next waypoint. To do this in the cockpit, use your hand as follows: put your thumb on your present position (Hackettstown) and your long finger on the next waypoint (Sussex). Now move your hand (without rotating it)³ until your thumb is at the center of the compass rose. In this case, the Broadway VOR⁴ is convenient, since it is only a few miles from Hackettstown. Now look along the line from your thumb to finger, and see where it crosses the compass rose. In this case we find that it crosses at the tickmark that corresponds to 040 degrees, which we take as our approximate magnetic course.

In the absence of other information, this approximate course is your best estimate of the proper heading. This may not be exactly your optimal heading, but it is a reasonable approximation, certainly better than maintaining your previous heading. Turn promptly to your best-estimate heading and maintain it while carrying out the next steps of the calculation. If and when you have information about crosswinds ([section 14.2.3](#)) and VOR twist ([section 14.3.6](#)) you can refine this estimate. Check for traffic again.

14.2.2 Distance, Time, and Airspeed

When looking for a waypoint, such as your destination airport, it doesn't do you much good to be on course if you have already inadvertently passed the waypoint. Therefore, it is vital to know how far you have progressed along the course. This is just as important as staying on course, and perhaps not as easy. Consider the contrast:

- It is rather easy to notice that you are off-course by half a mile when passing a waypoint.

- It is more difficult to notice that you are a minute early or late when passing a waypoint.

Note that the distance error involved in the second case is many times larger than in the first case.


To say it another way, it is easier to notice an unforecast crosswind that is blowing you left or right of course than it is to notice an unforecast headwind or tailwind that is messing with your progress along the course.

To keep track of distance along the route using pilotage, you need an estimate of your groundspeed. Then, given speed and distance, you can figure out how much time it will take to get to the next waypoint.

The first step is convert indicated airspeed to true airspeed. In this case, 110 K_{IAS} is about 120 K_{TAS}.⁵

The next step is to account for the wind. We need to resolve the total wind into a headwind component and a crosswind component. We will use the face of the directional gyro as an analog computer to help solve trigonometry problems.

Recall that the wind was out of 335 degrees at 25 knots. Since these forecasts always use *true* azimuth, you need to convert 335 true to 347 magnetic. (Notice how the compass roses on the chart are rotated relative to true north if there is any doubt as to the sign and magnitude of the correction.) Now find 347 degrees on your directional gyro. It will be at about your 10:00 position, as shown in [figure 14.2](#). Now we are going to use the circular face of the DG as a map. We choose the scale factor such that the radius of the DG represents the magnitude of the total wind, 25 knots in this case. Imagine a vector from the "347 degrees" point on the DG to the center. This represents the total wind, as shown in red in [figure 14.2](#).



[Figure 14.2](#): Wind Calculation using Directional Gyro

The headwind component is represented by the projection of the wind vector onto a line that runs vertically across the face of the instrument (from your 12:00 position to your 6:00 position), as shown in yellow in [figure 14.2](#). In this case its length is about 3/5ths of a radius, which represents about 15 knots. Therefore your groundspeed is about 105 knots (true airspeed minus headwind component).

Now, we need to estimate the distance of this leg of the flight. There are two ways to do this.

Method one is literally the rule of thumb. The length of my thumb (from the last joint to the end of the nail) corresponds to ten nautical miles on sectional charts, almost exactly. You can calibrate your own thumb. In this case, the required distance is about two and a half thumbs, or about 25 nm.

Method two is sometimes more accurate. Again put your thumb and finger on Hackettstown and Sussex, respectively. Now move and rotate your hand (without changing the distance between thumb and finger) so that you can use the tick marks on one of the north-south grid lines of the chart as a reference. One minute of latitude is one nautical mile.⁶ Again the answer is about 25 nm.

It is easy to remember that a groundspeed of 120 knots corresponds to two miles per minute. At that speed, you would be there in 12.5 minutes. However, in this example your groundspeed is about 10% slower than that, so it will take about 10% longer, about 14 minutes. You therefore expect to pass over Sussex at 46 minutes past the hour.

[14.2.3](#) Crosswind Correction

Now we are going to calculate the crosswind component. Again we will use the face of the DG to help solve the trigonometry problem.⁷

Recall that the wind was out of 335 degrees at 25 knots, represented by the red vector in [figure 14.2](#). The

projection of this vector onto line that runs horizontally across the instrument (from your 9:00 position to your 3:00 position) represents the crosswind component, as shown in blue in the figure. The length of this component in this case is about 4/5ths of a radius, which represents about 20 knots. That is, we have a crosswind component of about 20 knots, from the left.

To convert the crosswind velocity component to a crosswind correction angle, you can use the information in [table 14.1.8](#) In the present example, you should turn the airplane 10 degrees to the left of course,⁹ that is, to a heading of 030 degrees (heading = course + wind correction).

Groundspeed (knots, real)	Groundspeed (knots, pi=3)	Crosswind Correction (knots per degree)
57	60	1.0
85	90	1.5
115	120	2.0
145	150	2.5
170	180	3.0

[Table 14.1](#): Crosswind Correction Angle

If you are off course, apply an intercept angle (heading = course + wind correction + intercept) as discussed in [section 14.3.4](#). The heading problem is now solved.¹⁰ Check for traffic again.

14.2.4 The Wind Triangle


Concept #1: According to Galileo's principle of relativity, you cannot measure a velocity by itself; you can only measure the velocity of one thing *relative* to another.

Concept #2: Velocity is a *vector*; that is, it has a *magnitude* and a *direction*. (In contrast, something that has only a magnitude, without direction, is called a *scalar*.)

There are three velocities involved in dead reckoning, as illustrated in [figure 14.3](#), and as shown in [table 14.2](#).

Vector	=	Magnitude & Direction
airplane velocity relative to the air	=	airspeed & heading
airplane velocity relative to the ground	=	groundspeed & track direction
air velocity relative to the ground	=	wind speed & wind direction


[Table 14.2](#): Relative Velocities



[Figure 14.3](#): Wind Triangle

Note that the word *velocity* always refers to a vector, and the word *speed* always refers to the corresponding scalar magnitude.

If you want to draw an accurate wind triangle, you must be careful to draw the headwind and crosswind components as projections along and across your *course* as is shown in [figure 14.4](#). (It is a common mistake to draw them along and across your airspeed vector instead.) With the help of such a drawing you can understand why a direct crosswind (that is, a wind directly perpendicular to your course) will slow you down a little bit: even if there is no headwind component, your groundspeed (the base of a right triangle) will be shorter than your airspeed (the hypotenuse).



[Figure 14.4](#): Definition of Headwind and Crosswind

[14.2.5](#) Discussion

Note that in the scenario presented above ([section 14.2](#)) there was basically no alternative to using the quick, approximate dead reckoning techniques ([section 14.2.2](#) and [section 14.2.3](#)) for choosing the heading. Consider the possible alternatives:

- Pilotage was great for determining your position over Hackettstown, but it didn't tell you the outbound heading.
- Fancy instruments don't always help. The Broadway VOR could tell you that you are northeast of Broadway, but, again, you already knew that. Even if the VOR were on the field at Hackettstown it wouldn't have told you what outbound heading to use. A VOR on the field at Sussex would have simplified the job of finding the *course*, but you still would have needed to apply a wind correction to find the right *heading*.

Here is a rundown of various flight-planning methods:

1. The dead reckoning methods outlined here — even though they involve various approximations — are good enough for most purposes.
2. One very good method is DUAT (Direct User Access Terminal). Its original purpose was to allow pilots to get weather briefing information on-line, but the contractor that provides the service has augmented it with a free flight planner. It will find a route for you automatically, then compute the courses, distances, headings, and times between waypoints. It gets the forecast winds aloft directly from the FAA weather computers. An important advantage of this approach is that you can, with very little effort, try several different routes and several different altitudes. Sometimes it's worth going a little out of your way to pick up a tailwind. Information about DUAT in general and the flight planning features in particular can be found at the DUAT provider's web site.
3. On those rare occasions when I can't connect to DUAT, I have a spreadsheet in my laptop that will calculate headings, times, fuel consumption, et cetera, given courses, distances, winds aloft, et cetera. You can download a copy of the spreadsheet (and instructions) from my web site.[11](#)

By the way, if you are curious about the details of how to compute wind triangles to high accuracy, you can reverse-engineer the formulas in the spreadsheet.

4. There exists a clever mechanical device called an E6-B which works like a slide rule for adding and subtracting vectors. I haven't used mine in years. When I'm in the airplane, I use the rules of thumb discussed above to estimate headwind components and crosswind components. When I'm not in the airplane, I use a computer.

It is important to be able to solve navigation problems while you are flying the airplane. When you are in the cockpit improvising a flight plan, an approximate solution *right now* is vastly preferable to an exact solution that would require many minutes of careful calculation. There are many reasons why you might want to (or need to) improvise a deviation from your preconceived flight plan. These include:

- The winds-aloft forecasts are never precise enough to allow super-accurate dead reckoning.

- You might want to deviate around some isolated but threatening cloud build-up.
- On an IFR flight, it is quite likely that your clearance will not be identical to what you filed, and it is quite likely that whatever clearance you get will be changed enroute. Oftentimes the controller is trying to do you a favor by offering to shorten your route, and you can save lots of time and fuel by accepting. Other times the controller is trying to do somebody else a favor by changing your route, and you will cause lots of trouble if you can't do a good job of flying the amended clearance.
- One of the advantages of a pilot license is that it gives you more freedom to go where you want, when you want. Being able to improvise a flight plan makes flying more fun.

Almost the only time you need really accurate dead reckoning is when you are taking FAA written tests. You will be allowed to use an E6-B or an "approved" electronic equivalent; a laptop will not usually be allowed. The tests sometimes contain questions where the right answer differs from the wrong answer by a tiny amount. In such cases, you must use the FAA-approved approximations¹² and no others.

On a real trip (as opposed to a written test), if you are planning to rely on highly accurate dead reckoning, such as flying to the Azores without a GPS, then you should probably have your head examined.

14.3 Navigating by Instruments

14.3.1 Don't Be a Gauge Junkie

Navigating by instruments does not relieve you of your responsibility to see and avoid other aircraft.

This is especially important in the vicinity of VORs. It is altogether too common for pilots approaching a VOR to have their heads "down and locked" — paying vastly too much attention to the Course Deviation Indicator (CDI) needle and not enough attention to other traffic. The more accurately you fly over the VOR, the more likely you are to hit somebody else who is trying to do the same thing.

By the time you are within a couple of miles of the VOR, you should know how much wind correction is needed. (The wind doesn't change at the station!) You should know the course to the station, so you should be able to get there (plus or minus a tenth of a mile) by dead reckoning. Therefore take up the correct heading and just hold it. Don't chase the needle. Look outside.

Keep track of your position on the chart. Draw a line on the chart representing your intended course, and make sure it doesn't come too close to any obstructions or special-use airspace.

14.3.2 Navigation Systems

Navigation systems in common use include:

- LORAN = Long Range Navigation. It uses the frequency band from 90 to 110 kilohertz.
- NDB = Non-Directional Beacon. It uses the frequency band from 300 to 1600 kilohertz (which includes standard AM radio). The aircraft instrument that receives and interprets the NDB signal is called an Automatic Direction Finder (ADF).
- VOR = Very-high-frequency Omni Range. It uses the frequency band from 108 to 118 megahertz.
- DME = Distance-Measuring Equipment. It uses the frequency band from 962 to 1213 megahertz.
- GPS = Global Positioning System. It uses a system of satellites transmitting on approximately 1.5 gigahertz.

The principles of operation of these systems will not be discussed in this book.

14.3.3 Off-Course Distance

The *Course Deviation Indicator* on a VOR receiver indicates off-course *angle* (two degrees per dot). In contrast, on a GPS the CDI reads directly in miles (or fractions thereof), which is much more convenient. If you know how far you are from the VOR, you have to do a little work to figure out the off-course distance.

The simplest way is to look at the chart. If you are ten miles from the VOR, the distance between tick marks on the compass rose tells you immediately what distance corresponds to a five-degree off-course angle. If you are nearer or farther from the VOR, the off-course distance (for any given angle) is proportionately smaller or

larger.

The other way is to use arithmetic. Suppose you are 57 miles from the VOR. Since there are 57 degrees in a radian, at this point each degree of off-course angle corresponds to one mile of off-course distance. At this point (57 miles from the VOR), a three dot deflection corresponds to being six miles off course, which is embarrassingly poor navigation. In contrast, suppose you are only a couple of miles from the VOR. Then the same CDI deflection (three dots, which is six degrees) corresponds to being off course by less than a quarter mile, which is perfectly fine navigation.¹³

Bottom line: when you are close to the VOR, do not overreact to small CDI deflections. Conversely, when you are far from the VOR, you must notice and react to rather small CDI deflections.

14.3.4 Intended Heading

Here is an example which is typical of conversations I often have with my students during cross-country flights:

- Question: ``What is your intended heading, and why?"
- Answer: ``Heading equals course plus wind correction plus intercept. The course is 040, there's about 10 degrees of crosswind from the left, and I need about 5 degrees of intercept since the course is off to my right. Therefore intended heading is 035 degrees (040 - 10 + 5)."

Note #1: This is a question about *intended* heading. It has almost nothing to do with the present *actual* heading. You should be able to answer this question without looking at the DG.

Note #2: The formula for the desired heading is always:

$$\text{Heading} = \text{Course} + \text{Wind Correction} + \text{Intercept Angle}$$

A ten-degree intercept angle is usually plenty. If you are a mile off course, a ten degree intercept angle will get you back on course in less than 6 miles, which should be just fine for typical enroute navigation. As you get better at navigation, you will be able to detect smaller off-course distances (say, half a mile), in which case a smaller intercept angle (5 degrees) will be appropriate. Small corrections are the mark of a pro.

If you are farther off course, say 2 or 3 miles, you can still use a ten degree intercept angle, which will get you back on course in 12 or 18 miles. If for some reason you need to be back on course sooner than that, you can use a larger intercept angle.

Usually, the reason you are off course is because you didn't do a very fastidious job of maintaining the correct heading over the last few miles. In such a case, the solution is straightforward: choose the right heading (course + wind correction + new intercept angle) and maintain it.

In other cases, you might have been blown off course by an unexpected wind. In such a case, you might want to revise your estimate of the crosswind correction angle. Therefore the intended heading will be course + *revised* wind correction + intercept.

14.3.5 Cross Radials

It is worth repeating what was said in [section 14.2.2](#): When looking for a waypoint, such as your destination airport, it doesn't do you much good to be on course if you have already inadvertently passed the waypoint. Therefore, it is vital to know how far you have progressed along the course. This is just as important as staying on course.

GPS, LORAN, or DME make it easy to keep track of your progress along the course. VOR and NDB stations, provided they are not directly ahead or behind you, can also provide progress information. Radials from such a off-course station will cross your course, and successive radials correspond to successive points along the course.

The question arises as to *how far* off your course the off-course station should be. If the station is too far away, you may have trouble receiving the signal. Also, the farther the station is away, the less precise will be the information you get from it, just because the same number of degrees will correspond to a longer distance. On the other hand, you don't want the station to be too close to the course. This is because you

want the cross radials to cross the course at a reasonably large angle (preferably 45 degrees or more); otherwise accuracy is impaired. Therefore, if you chose a station that is farther from the course its usefulness will extend over a longer portion of the flight.

14.3.6 Twisted VORs

In the region where I fly, every VOR is misaligned by several degrees. If you want to fly along an airway that is defined by, say, the 224 radial of a certain VOR, you need to hold a 227 heading (in no-wind conditions).

Here's why: In general, the radio-beams that a VOR radiates are not necessarily aligned with the actual magnetic directions. Presumably the transmitters were properly aligned when they were first installed, but some of them have not been re-aligned in over 35 years.

That's significant, because the earth's magnetic field changes over time. A VOR that was aligned with the magnetic directions several decades ago may disagree with the current magnetic directions by quite a bit. The FAA is ``supposed" to re-align them, but they've fallen rather far behind. Re-alignment is a lot of work: not only do you need a really big wrench to rotate the transmitter, but you need to revise all the navigation charts.

Your GPS will agree with your compass and disagree with the VOR. That's because GPS receivers have a database that can be updated with the latest map of magnetic variation.

Here are some of the implications:

- If you are using dead reckoning, when you select a course based on the charted airways or compass roses, add the VOR twist to the charted radial before using it.
- If you are flying by reference to VORs, add the VOR twist to whatever your VOR receiver is saying before using it to select a heading. That is, the nice rule (heading = course + wind correction + intercept) given in [section 14.3.4](#) must be replaced by an uglier rule: heading = radial + twist + wind correction + intercept.[14](#)
- If you are *not* flying by reference to VORs, but rather using GPS to fly from waypoint to waypoint along an airway, do not be surprised if your GPS indicates a bearing that differs from the charted VOR radial for that airway.
- If you are using a GPS or a landmark to check the accuracy of your VOR receiver at some arbitrary off-airway location, add the VOR twist to whatever your VOR receiver is saying, then compare that to the actual magnetic bearing. Usually, though, it is better to use the GPS to identify a named intersection on a VOR airway or on a VOR approach, and check the VOR against the published VOR radial (not GPS radial) to that point.[15](#)
- If you simply want to use a VOR to determine whether you are following an airway, or an instrument approach procedure, you do *not* need to worry about VOR twist. The charts conveniently tell you what VOR radial goes from point A to point B.

Here is how you can ascertain the VOR twist: The first step is to obtain the actual magnetic variation in your area. Often the easiest way is to look through the Airport/Facility Directory[16](#) and find a nearby airport that has been recently surveyed. That will give you the local variation as of a specified date. (Sectional Aeronautical Charts are another source of information about magnetic variation, but you never know how up-to-date that information is.) You can also get your GPS to tell you the local variation. The second step is to look up the VOR in the A/FD, and see what variation the VOR is aligned to. Subtract the VOR alignment number from the actual variation.

14.4 Combined Techniques

You should not hesitate to combine pilotage, dead reckoning, and navigation by instruments. For instance, you could use a VOR signal to stay on course left/right, and use time and groundspeed to measure your progress along the course. Conversely, you could use dead reckoning to stay on course and use a cross-radial to measure your progress along the course.

The combination of dead reckoning with pilotage is quite powerful. Dead reckoning helps you find your landmarks. Pilotage allows you to establish your position with certainty, so that small dead reckoning errors (which are inevitable) do not accumulate.

14.5 Staying Un-Lost

Here are some suggestions to help you keep from getting lost:

1. Keep track of your current position. I know this involves a certain amount of work, but remember: Staying un-lost is easier than getting un-lost. Consider the analogy: don't wait until you have cavities and then start brushing your teeth. By the same token, don't wait until you are lost to start keeping track of your position.
2. Write on the charts. When you pass a checkpoint, write the time on the chart, so later you will know how long it has been since you were there.
3. Choose unambiguous checkpoints. If you think you are heading for Jonesville, make sure you are not heading for Smithville instead. (A water tower with the word ``Smithville" in twenty-foot-high letters should make you suspicious.) If your checkpoint is a round lake, make sure there is not another round lake a few miles away.
4. Keep your DG aligned with the magnetic compass. Sometimes a DG will behave nicely for hours or even years. Then, just when you have become complacent, it will start precessing like crazy. Also note that radical maneuvering (steep turns, stalls, takeoffs and landings) can cause an otherwise well-behaved DG to lose a few degrees.
5. Identify the Morse code on each navaid that you use. I've seen lots of students tune up the wrong navaid, or an inoperative navaid, and then blissfully follow the needle. Don't just check that there is ``some" Morse code, make sure it is the *right* Morse code.
6. Make the VOR status flag part of your scan. If you lose the signal from the VOR (perhaps because you flew out of range, or perhaps because of a loose connection in your receiver) the CDI needle will settle in the center, making it look like you are doing an excellent navigation job no matter how far off course you are. The easiest defense is to notice that the ``To / OFF / From" flag is showing ``OFF".
7. Make a habit of getting enroute radar advisories (also known as flight following) from ATC. On my very first cross-country trip after getting my private pilot certificate, I was blissfully following road ``A" when I thought I was following road ``B". Unfortunately, road ``A" was leading right into the middle of a restricted area that was being used for air-to-air missile testing. Fortunately, I was getting advisories. A few miles before the restricted area ATC called me and suggested I make an immediate 90 degree turn. At that point I didn't even realize I was lost, so it took me a while to understand what ATC was saying. Eventually I took the hint and everything worked out OK.

14.6 Getting Un-Lost

This section gives some suggestions for what to do when you are uncertain of your position.

14.6.1 Basics

First of all, don't panic. Being slightly lost is not, by itself, a big-time emergency. However:

- Being lost and low on fuel is a big problem. Therefore try to get un-lost reasonably promptly, while you still have plenty of fuel.
- Being lost at night in mountainous country is a big problem.

14.6.2 When in Doubt, Climb

Low altitude causes lots of problems, including:

- You might run into obstructions.
- Your ability to see distant landmarks is limited.
- Your ability to receive VOR signals is limited.
- Your ability to use your communications radio is limited.

There are of course exceptions: For instance, you don't want to climb into a cloud layer unless you have current instrument-flying skills and a clearance. Similarly, you don't want to climb into restricted airspace without permission. Still, given the choice between running into a mountain and violating restricted airspace, the latter is preferable.

14.6.3 GPS or LORAN

Most GPS and LORAN receivers have a really nice feature: By pushing a button or two, you can display the name of the nearest airport(s), along with the bearing and distance from your present position to there.

These instruments will also, of course, tell you your latitude and longitude, but usually this is less convenient than the ``nearest airport" feature.

14.6.4 VOR Cross Radials or VOR/DME

If you know even approximately where you are (within a few dozen miles), pick a VOR in the area and tune it up. Draw a line on the chart, along the radial that the VOR is telling you. Then pick a second VOR and draw another line. The point where the lines intersect is your position. If the lines cross at a shallow angle, the precision of the fix will be poor, so when picking the second VOR try to pick one that is well off the line given by the first VOR.

Of course, if you have VOR and DME, the job is even easier.

14.6.5 Ask ATC

Here is an exchange I heard on the radio, back when I was a student pilot:

Voice 1: PSA 1705, cleared for visual approach.

Voice 2: Approach, PSA 1705 is unfamiliar with the area, requesting vectors to final.

Voice 1: Roger, PSA 1705, fly heading 350, vectors to final.

I smiled when I heard that. I figured if airline captains could ask for vectors, so could private pilots, and even students.

It's true: you don't need to declare an emergency. You don't need to admit that you're lost. You don't even need to *be* lost. You can just be slightly unfamiliar with the area.

ATC has radars. They can find you real fast, and give you a vector toward wherever you want to go. Even without radar, some flight service stations can find you by doing ``direction finding" on your VHF radio transmissions (although this system is slowly dying of neglect).

Don't worry about getting ``blamed" for being lost. ATC would much rather have a lost pilot who is talking to them than a lost pilot who isn't talking. You should be embarrassed enough to be motivated to do better navigation next time, but not so embarrassed that you hesitate to ask for help this time.

To contact ATC, if there is any doubt¹⁷ about what frequency to use, call up on 121.5 MHz. Practically every ATC facility can receive and transmit on that frequency. Yes, it is the ``emergency" frequency, but it is not so special that you should be the least bit hesitant about using it.

¹ ... presumably as a way of discouraging uninvited guests.

² Use a pencil, so that you can erase and re-use the chart for your next flight.

³ The idea is that the direction from your thumb to finger is the same before and after the move. You may want to practice this skill. Here's a good exercise: find a nice straight Victor airway on the chart. Put your thumb and forefinger on it, twenty or thirty miles apart. Then move your hand to a compass rose that is not on the airway, and read off the heading. Finally, look back at the airway and see what the ``official" heading of the airway is. With a few minutes practice, you should be able to get within a couple of degrees.

⁴ See [section 14.3.2](#) for a list of navigation systems and their acronyms.

⁵ In many airplanes, there is a device like a circular slide rule built into the airspeed indicator that allows you to calculate the true airspeed, given the indicated airspeed, altitude, and temperature. Otherwise, you could perhaps use your E-6B. Typically though, it suffices to use the simple rule: two

percent per thousand feet.

[6](#)

That's how nautical miles were originally defined, and that's why aviators use them. A nautical mile is about 1.15 statute miles, or about 1.85 kilometers. A *knot* is defined to be a nautical mile per hour.

[7](#)

To practice these techniques on the ground, you can use the compass roses on the chart, in the same way that you would use the DG if you were in the airplane. That is, draw a line from the edge of the compass rose to the center, representing the total wind, and then resolve it into components parallel and perpendicular to your course line.

[8](#)

You don't need to memorize the table. You just need to remember that there are about 57 degrees in a radian; then you can figure out the rest at a moment's notice. The figuring is even easier if you approximate p by 3.0 (i.e. 60 degrees per radian). You can speed things up by remembering the conversion factor (degrees per knot of crosswind) that applies at the typical groundspeed of your airplane. Also, you can usually simplify the calculation by comparing the crosswind component to your airspeed (as opposed to groundspeed), unless the headwind or tailwind component is really large.

[9](#)

To calculate a more-accurate course, see [section 14.3.6](#).

[10](#)

If you want to get technical about it, you should recalculate the crosswind component after correcting for the crosswind, but you will find that this is almost never a significant correction.

[11](#)

See [section 20.4](#) for more information about this.

[12](#)

In particular, the test-makers apparently believe $p = 3$ (i.e. 60 degrees per radian). I remember one question where you would be marked wrong if you used $p = 3.14$ (i.e. 57 degrees per radian).

[13](#)

... enroute navigation, that is. In contrast, if you are performing an instrument approach, you might want to do better, but that is a different subject.

[14](#)

If you forget to account for twist, you can drive yourself crazy trying to figure out the winds aloft. Suppose you are following an airway that is twisted by +3 degrees, flying at 120 knots in no-wind conditions. Then your instruments will seemingly indicate 6 knots of ``wind" from the east when you are northbound on the airway, and seemingly indicate 6 knots of ``wind" from the west when you are southbound on the same airway!

[15](#)

This is the most accurate check you can do, but does not quite meet the requirements for the mandatory 30-day receiver check — unless you check two VOR receivers at the same time and sign it off as a cross-check.

[16](#)

The A/FD is published as a small book every 56 days. The same information is also available on the web, which is often more practical. For example, search for ``JFK VOR variation".

[17](#)

At the other extreme: if you are already talking to ATC on a given frequency, you should undoubtedly use that frequency.

15 Emergency Procedures

Q: ``What should I do if the door comes open in flight?"

A: ``Fly the airplane".

15.1 Engine Out Procedures

This section discusses what you should do if your engine quits while you are airborne.¹ This mainly applies to single-engine airplanes; additional procedures for multi-engine airplanes are covered in [chapter 17](#).

15.1.1 Emergency Checklist

It is important to have an emergency checklist. You should commit it to memory, and review it right before each flight. Do not wait until you are confronted with a ``deafening silence" to figure out what is on the emergency checklist, and why.

If your aircraft manuals do not provide a suitable emergency checklist, you might consider adopting something along the following lines:

Emergency Checklist

- **Aviate** — best-glide K_{IAS} and trim; configure.
- **Navigate** — pick a field; turn toward it.
- **Investigate** — carb. heat, boost, tanks, primer, mags, mixture
- **Communicate** — 7700, 121.5 or current.
- **Secure** — gear up for short, soft or water;
throttle, mixture, mags, master, tanks — off;
belts — snug.

This list has been designed to make it easy to memorize. You should make every effort to commit it to memory, so that if somebody wakes you up in the middle of the night and asks you ``what is the emergency checklist" you should be able to shout, instantly, ``AVIATE, NAVIGATE, INVESTIGATE, COMMUNICATE, SECURE!"

The first item, **Aviate**, is clearly the first priority. No matter what happens next, you want to be in control of the aircraft when it happens. There are lots of scenarios where an engine failure results in a critically low airspeed (especially if somebody is dumb enough to try to maintain the pre-failure pitch attitude, or (worse) the pre-failure altitude while deciding whether or not there has been a failure). If the airspeed is low you *must* re-establish the proper glide speed² immediately, even if it means cashing in some precious altitude.

The opposite extreme is possible, too; namely it is quite possible that at the time the emergency begins, the aircraft is going much faster than the best-glide speed. This is not so immediately dangerous, but the longer you take to establish the proper glide speed the more energy will be thrown away in the form of unnecessary parasite drag. In this case, gently zoom upward, converting airspeed to altitude. Retrim.

In addition to trimming for the correct airspeed, you should configure everything else appropriately, as discussed in [section 15.1.2](#).

The second item, **Navigate**, is clearly next in importance. In [section 15.1.4](#) there is a discussion of clever techniques for judging which fields are within gliding range — but you should not pick a field at the limits of this range if there is anything suitable that is close. In particular, start by looking down at a 45 degree angle, or even straight down. If it is right below you, it is probably within gliding range!

The next item is **Investigate**. Sometimes when the engine quits, you know immediately what the problem is. Ninety percent of the time, the problem is fuel-related, so you should reflexively switch tanks and turn on the boost pump as appropriate. Then turn on the carburetor heat, because it is only effective while the engine is still warm. Then go left-to-right across the panel, checking everything in turn. Make sure the primer is in and locked. See if the engine runs better on the left magneto, right magneto, or both. See if it is happier with a leaner or richer mixture. In most cases the propeller will keep turning, just due to the action of the relative

wind, even in the total absence of engine power -- but if it stops, use the starter to get it going again. Give everything a once-over before spending too much time on one particular item, unless you are pretty sure you know what the problem is. And above all, don't forget to fly the airplane.

The next item is **Communicate**. If you are already in contact with a controller, it is almost certainly a good idea to stay on that frequency. If, on the other hand, you have any doubt about what frequency to use, go immediately to the international distress frequency, 121.5 MHz. That's what it's there for.

Similarly, if you have been assigned your own squawk code, don't bother to change it unless ATC asks you to. On the other hand if you are presently on the all-purpose code 1200, do not hesitate to switch to the emergency code, 7700. That rings alarm bells (literally) at ATC and highlights you on the controller's radar screen.

Some people argue you should Communicate even before you Investigate. Certainly if you are in instrument meteorological conditions you should tell the ATC of your predicament even before you Navigate, (1) so they can vector you to a landing field and (2) so they can clear out the airspace below you.

The fifth item on the list is **Secure**. It is amazing how easy it is to forget this item. Wouldn't you hate to make a beautiful power-off approach to an ideal field — and then forget to extend the landing gear?

At 100 feet AGL, make sure you pull the throttle and mixture to idle cut-off. The main reason is that you don't want the engine to roar back to life just after touchdown. This could easily happen if (for example) there had been a fuel shortage, and the flare freed up some fuel from a corner of the tank. The reason for doing it at 100 feet AGL is to give the engine a chance to cool down, reducing the risk of a post-crash fire. Closing the fuel-tank shutoff valve helps reduce the risk of fire — but in most planes it is *not* a sufficiently quick way of stopping the engine so be sure to pull the throttle and mixture also.

Shutting off the engine will be difficult; it will require overcoming a huge psychological barrier. After all, you've spent the last several minutes trying to restart the engine, and now you are supposed to shut it off. Make sure you have made this decision in advance: promise you will shut the engine off at 100 AGL.

Switching off the master also reduces (somewhat) the risk of fire, but in an aircraft with electric flaps and/or landing gear, you might want to save the master switch for last.

- Aviate.
- Navigate.
- Investigate.
- Communicate.
- Secure.

15.1.2 Configuring for Glide

In the ``clean" configuration, the airplane will be able to glide much farther, perhaps twice as far as in the ``dirty" configuration. If you start out at low altitude, twice nothing is nothing, so it may not be worth messing with the configuration; just configure for landing and be done with it.

On the other hand, if you start out at a reasonable altitude and are trying to glide a long ways, then you want flaps retracted, landing gear retracted, and propeller in the coarse pitch (low RPM) position.

Some other books say that once the flaps are down, you should leave them down; they point out that at a low airspeed (below the bottom of the green arc) retracting the flaps will cause an immediate stall.

I look at it somewhat differently. This situation actually arose on my private pilot checkride. I was at 1000 AGL, with two notches of flaps extended for landing. Then Tower asked me to extend my downwind. By the time I was able to turn final, I was nearly two miles from the airport. At this point the examiner caused a simulated engine failure.

I went through the following thought process:

- It's a long way to the airport. If I don't do everything right, we won't make it.

- It is important to glide at the right airspeed. I know what the best glide speed is in the clean configuration, but I have no idea what the best glide speed is in the current configuration.
- I'll bet there is no such thing as best glide in this configuration. I've got to get the flaps up.
- I'm really slow, near the bottom of the green arc. If I just retract the flaps, I might stall.
- The airplane stalls at a definite angle of attack. No airplane ever stalled at zero angle of attack. It's mathematically impossible.
- I've got full control over angle of attack. Watch this!

At that point I shoved forward on the yoke. Zero angle of attack. Zero *G*. The examiner started gently floating out of his seat. I retracted the flaps all at once. I continued the zero-*G* pushover until we approached the canonical best-glide airspeed. Then I raised the nose, trimmed for best glide, and quickly ran the rest of the emergency checklist. I even rolled in some left rudder trim.

The glide took us to a place in the weeds about 100 yards short of the runway. I flew right down into ground effect and then flared. While skimming in ground effect I extended the flaps. When we reached the runway the stall warning was already on. I plopped onto the runway. We were stopped before reaching the big painted number.

The point of this story is this: If you need to glide a long ways, retract the flaps. Just do it in such a way that you don't stall.³

You can, of course, glide with flaps and/or gear extended if you want to make a steep approach to a nearby field.

Also, when you are through gliding (i.e. when you are ready to flare), make sure the flaps are extended, so you can touch down at the lowest possible speed.

For landing on water, in most airplanes you want the gear up. For landing on most other surfaces, you want the gear down. Don't wait until the last moment to put them down; with the engine off it might take longer than usual to get them down. Make sure you know how to use the manual gear extension system on your airplane.

15.1.3 Return to Airport?

We now focus on the special case of engine failure shortly after departure, since that is a relatively common and very critical case. Many people are tempted to turn back to the airport, but this is not usually the right answer.

The right answer depends on many factors, including:

- the wind
- the length of the runway
- the capabilities of the airplane
- whether or not partial power is still available
- the capabilities of the pilot

Every situation is different, so the following analysis can't possibly fit them all exactly. On the other hand, it is worth your while to plan in advance. Know what your options are. For each phase of flight, make sure you have a backup plan ("Plan B") appropriate to the situation. Be ready to carry out Plan B at a moment's notice.

Here is a piece of simple but important advice: if you can land straight ahead, do so. As an extreme example, consider this: a small plane departing from runway 31L at JFK (length: 14,600 feet) could climb to 500 feet, lose the engine, and still land straight ahead on the same runway with plenty of room left over.

Here's another piece of simple advice: don't turn back unless you are sure you can make it — and there are lots of situations where you *can't* make it.

For example, consider a fully loaded Cessna 152. It has a power-off glide ratio of ten to one. Unfortunately, in no-wind conditions the climb gradient is *less* than ten to one. Therefore, even if the airplane could turn on a dime, at every point on the return trip the airplane would be below where it had been on the outbound trip. Then when you take into consideration the altitude lost while getting the airplane turned around, it is easy to see why the airplane cannot possibly return to the point where it left the ground.

Under such conditions, the farther you have flown on the departure leg, the more options you have for an off-airport landing, and the more impossible the turnback becomes.

An important factor to take into account is that a simple 180 degree turn does not suffice to return you to the departure runway. The airplane will travel an appreciable distance *sideways* during the turn. You won't need to do a full-blown procedure turn, but you will need to do some additional maneuvering that makes an already-bad situation worse.

Given a sufficiently long runway, the airplane may be able to return to a point on the runway closer to the departure end — which is just fine. Again, imagine departing runway 31L at JFK, and climbing straight ahead. Suppose the engine quits at a point 1/2 mile beyond the departure end of the runway. At that point you should have more than a thousand feet of altitude. You should be able to reverse course and make a downwind landing near the beginning of runway 13R⁴ even though you could not glide back to the point where you lifted off.

A modest headwind on departure will help keep the airplane near the airport during the outbound leg, and will help hurry it back to the airport during the return trip. This sounds wonderful, because it increases the possibility that you can glide back to the runway. The trouble is that (whether or not you make it back to the airport) you are faced with a downwind landing. Even a modest amount of tailwind (say 15 knots) can have a tremendous effect. Suppose your airplane is capable of touching down at 55 knots. If you land into the wind, you have a groundspeed of 40 knots, but if you land downwind you have a groundspeed of 70 knots. Runway usage depends on the *square* of the groundspeed, so the downwind landing will use *three times* as much runway: $(70/40)^2 = 3.06$. Also, in a collision, the amount of damage and injury is typically proportional to the square of the groundspeed — so if you turn downwind and *don't* manage to land on the runway you are in very big trouble indeed.

Here's another option for you to consider: suppose that your airport has a second runway running crosswise to the active runway. If your engine fails somewhere over the cross runway, you might be able to turn 90 degrees and land on it.

Even in less ideal cases, it is quite likely that a crosswind landing on a different runway (or even a taxiway) is easier and safer than a downwind landing on the departure runway.

If you are really concerned about engine failure during the departure climb, and the airport is the only safe landing zone for miles around, you should begin a gentle turn almost immediately after liftoff. Then if the engine quits, you're closer to the airport and you've got a more convenient heading. I don't recommend this in general, because engine failure is not the only consideration. For starters, we need to worry about causing a mid-air collision in the pattern. A turning departure climb toward the traffic-pattern side of the runway would cause you to enter the downwind leg from below at just about the point where inbound traffic is entering from the 45 degree leg.

In many cases, engine trouble results in partial rather than total power loss. You have to decide whether continued operation of the damaged engine is safe, but if so, it gives you some more options. Even if the aircraft is not capable of climbing on the remaining power, the rate of descent may be dramatically reduced. Think of the aircraft as a noisy glider with a good glide ratio. It may be capable of "gliding" to places that a totally unpowered aircraft could not.

I reiterate: don't try to turn back to the airport unless you are sure you can make it, and in most typical cases you can't. You should find a nice road or field and put it down under control. It helps a lot if you have practiced forced landings, so you know the power-off performance of your airplane and how to land from non-ideal approaches.

Another very serious consideration is this: reversing course smoothly wastes valuable time and energy, whereas reversing course quickly requires radical maneuvering. Nobody wants to recommend that pilots perform radical maneuvering at low airspeeds close to the ground in an unplanned situation. It might help you return to the runway, but there is a much greater chance that it will provoke a stall/spin accident. An

off-airport landing is not usually fatal, whereas a stall/spin accident usually is.

Another reason for not attempting to turn back is that most people are so surprised by an engine failure that it takes them a few seconds to regain their wits. During this interval, precious time, energy, and distance have been wasted, so even though it might originally have been possible to turn back, it no longer is.

Indeed, (unless you are very well trained) your first reaction will be completely wrong — not just late, but dead wrong. When the engine quits, the airplane's nose will tend to drop. The flight path has changed from, say, a 10-to-1 climb to a 10-to-1 descent, so to maintain constant angle of attack the nose *must* drop a huge amount — 12 degrees. You may think ``Gee, I don't want the nose to drop" and may be tempted to pull back on the yoke. This is a sure way to kill yourself. Please, do not think of the yoke as the up/down control. When the engine quits, the airplane is going to descend. The only question is whether you will spin in, or glide to a controlled touchdown. Remember, you can survive an off-airport landing if you touch down under control.

The obvious reason why you want to maintain control is that the rate of descent you get in a normal glide is much less than what you get if you stall and let it ``drop in" — not to mention what you get in a spin.

A less obvious but still very important consideration is this: If the airplane is not under proper control, it is likely that one wingtip will hit before anything else. This will cause the aircraft to cartwheel, causing tremendously more damage and injuries than if you had landed under control and just skidded to a stop.

It is important to have a plan. At the airport(s) you use regularly, scout out the territory near the departure paths, and formulate a plan for where you will land if the engine quits. (Further discussion of the power-off glide appears in the next section.)


In any case, you need a plan for what to do with the controls. Your first priority is to maintain a proper angle of attack. Do not attempt to hold the nose up; let the nose go down (or push briefly to help it go down). Fine tune the pitch attitude and trim to maintain the best-glide airspeed.

Land into the wind if possible.

15.1.4 Power-Off Glide Perception and Planning

In a forced-landing situation, your glide path will be rather steep. The lift-to-drag ratio of typical Skyhawk or Cherokee (in best-glide configuration) is about 10-to-1, which corresponds to an angle of six degrees. This is perhaps twice as steep as a typical power-on approach. It is even somewhat steeper than the typical ``power off" approach, since that normally really means ``engine idle" and an engine at idle produces noticeably more power (and less drag) than an engine that is really off.

We can use the rule of thumb: a thumb at arm's length subtends four degrees.⁵ As shown in [figure 15.1](#), you can just barely glide to any point that is a *thumb and a half* below the horizon (assuming no-wind conditions). If your engine quits, you should run your thumb around the horizon to help you select possible landing sites.



[Figure 15.1](#): What Landing Sites Are Reachable?

In the presence of wind, the circle of possible landing sites will be shifted downwind from the circle described in the previous paragraph. Suppose you are gliding at 60 knots (airspeed) into a 30-knot headwind. Your groundspeed has been cut in half, so your glide will be twice as steep as in the no-wind case. Your destination must now be at least three thumbs below the horizon.

It is well worth knowing how far you can glide. Suppose you are roughly one mile up, so that you can glide ten miles. There is more area in the ring between seven and ten miles away than there is in the entire disk between zero and seven. (On the other hand, if find a close field, go for it. Don't glide a long way just because you can.)

Now let's assume you have picked a field and are gliding toward it. Now your ability to perceive angles really

pays off. Whereas in a normal approach you use engine power to maintain a pre-determined glide angle and destination, in a maximum-distance glide the glide angle is fixed and you want to perceive what the destination is.

Here is the key idea: there will be *some* landmark that remains *some* fixed angle below the horizon, and that is the point toward which you are gliding. Pick a point. If its angle down from the horizon is decreasing, you will land short of that point. If its angle is increasing, you will overfly that point — unless you do something.

There are two critical reasons why you always need to explicitly identify the point to which you are gliding. For one, you need this information for angle of attack control. Remember, the angle of attack is the angle between the wing and the direction of flight. If you can't perceive the direction of flight you won't be able to perceive the angle of attack (see [chapter 2](#)). Secondly, you need to know whether your present glidepath will cause you to land long or land short; the earlier you can perceive this the sooner you can make any necessary correction.

A small correction early is better than a big correction late.


You should not pick a field that is right at the limit of the airplane's gliding ability, because you'll wind up short if anything goes wrong. Pick a field that is substantially closer than the limit, since you can always lose altitude by circling, adjusting the length of the base leg, extending flaps, slipping, et cetera.

Strategic turns are appropriate early in the game; flaps and slips are more appropriate on short final.

S-turns on final are almost never the best way to eat up unwanted altitude. Small-angle turns have almost no effect, and large turns take too long to perform, and take your chosen field out of sight temporarily. Furthermore, after two turns (one to the left and one to the right) you will be back on your original heading, but offset laterally; you need to make two more turns to get back on course. If you have time and altitude to do all that, there are better things to be doing.

If you are on long final and can (using flaps and/or slips) keep the field a constant angle below the horizon, you are all set. Glide straight on in.

If you are on long final with excess energy, or if you are approaching the field from a substantial angle relative to the intended direction of landing, do not aim directly for the field. Aim for the so-called *base key point* ([figure 15.2](#)), i.e. the point where the base leg begins.



[Figure 15.2](#): Forced Landing from Base Key Point

From the base key point, you have a lot of options. If you arrive with the ideal amount of energy, you can fly a nice base leg and then turn final. If you arrive with slightly more or less energy than that, you can angle the base leg away from or toward the field. Fly along the base leg until the desired destination is the appropriate angle below the horizon, then turn final.⁶

Another point that is made in the figure is that you can use the *width* of the landing field to your advantage. You may be unaccustomed to this, since at an airport the tradition is to land always as close as possible to the runway centerline.

Also, do not plan your final approach to take you to the threshold of your chosen field. No matter how long or short the field, aim for a point one third of the way along the field. Remember there are lots of things that could steal energy from your glide, and you *really* don't want to land short.

It is better to hit the trees at the far end at 20 knots than to hit the trees at the near end at 60 knots.

The energy, the expected damage, and the expected injuries all depend on the *square* of the airspeed, so the 20-knot collision involves *nine* times less energy than the 60-knot collision.

The main reason why you might wind up landing short (despite a well-planned glide) is the infamous decreasing headwind on final, as mentioned in [section 12.12](#). If you are gliding at 70 knots, and you lose just ten knots due to a windshear, you will have to descend 60 feet to regain your proper airspeed.⁷ That 60 vertical feet (at a ten-to-one glide ratio) corresponds to 600 feet of horizontal travel. If you are unprepared for it, finding yourself 10 knots slow, 60 feet low, and/or 600 feet short during a forced landing is no fun.

[15.2](#) Preventing Emergencies

[15.2.1](#) Safety Margins

You need to have *margins* of safety and *layers* of protection. You don't want to be in a situation where any one failure causes harm. You always want to have not just plan A, but also plan B, plan C, plan D, et cetera.

For example:

- The left magneto is a backup for the right magneto, and vice versa.
- You have multiple sources of navigation information (pilotage, dead reckoning, VOR, GPS, ATC radar) which can be cross-checked against each other.
- During preflight, you use a dipstick to cross-check the fuel gauge. Enroute, you use the expected fuel-burn rate and a clock to cross-check the fuel gauge.

You should pay attention to anything that peels away one or more layers of safety. Keep track of how many layers remain.

If one magneto fails, park the airplane until it is fixed!

If the fuel gauges cannot be trusted, park the airplane until it is fixed! Do not rely on clock-and-dipstick alone. By the same token, do not rely on the gauges alone.

[15.2.2](#) Fuel Management

Far and away the most common reason for losing engine power is *fuel mismanagement*. This includes running out of fuel as well as contamination of the fuel. The good news is that such problems are relatively easy to prevent.

It pays to be careful. I am pretty methodical about checking the fuel sumps. I used to check for little droplets of water at the bottom of the sampler. Once, after completing the check, I was about to dump out the sampler and begin the flight, but I decided to take a closer look. Then I noticed tiny drops floating at the *top* of the sampler. It turns out the entire tube was full of water, with just a tiny bit of fuel mixed in. I went back to the plane and got another tube of water, and another, and another. It turns out that the fuel vendor had just switched to a new pump/tank system, and had sold me more than a quart of water (at AVGAS prices!) along with the fuel.

So, here are some suggestions:

- Check the fuel sumps if the aircraft has been sitting overnight or longer. Humid air in the tanks can condense at night. The problem is worse if the tanks are less than 100% full, because that leaves more room for air. Because of the daily temperature changes, new air gets into the tank each day. The condensed water, hiding underneath the fuel, does not re-evaporate — it just accumulates day after day.
- Check the sumps if the aircraft has been refueled since the last flight.
- If you do detect water after refueling, notify the fuel vendor immediately, so that one of your less-meticulous fellow pilots doesn't get harmed. Then, wait a few minutes and check your tanks again. A fair amount of water can be suspended in the fuel in the form of tiny droplets that take a while to settle out.

- Check the fuel sumps if the airplane has been sitting in the rain even for a short time. The filler caps have been known to let water leak in.
- Check the color of the fuel. Different octane grades are color-coded differently. The color is rather pale, so it may help to look lengthwise down a long column of fuel.
- Check the odor of the fuel. If it smells like jet fuel, watch out. There have been many cases where an airplane that runs on AVGAS has been mis-fueled with Jet-A. The engine may run on the mixture for a while, but it will be rapidly destroyed because of detonation in the cylinders.
- If you suspect here is a mixture of Jet-A along with AVGAS in your tank, here's how you can check: Put a drop or two of the suspect fuel on a piece of paper. For comparison, put a similar amount of known pure AVGAS on the paper nearby. The AVGAS should evaporate rather quickly. The Jet-A, if any, will remain behind, leaving a translucent spot on the paper.
- There are several ways to detect a sample that is 100% water. For one, the water will not have the right color, since the fuel color-code dyes are insoluble in water. Secondly, water has a noticeably different density and viscosity — it just doesn't "slosh" the same way. It also doesn't evaporate at the same rate. Last but not least, you can *add* a drop of water to your fuel sampler and make sure it goes to the bottom.
- Before each flight, peer into the tank to make absolutely sure the fuel *quantity* is OK. Cross-check what you see against the gauges. One fellow I know bought gas and got ready to take off, but noticed that the tanks were nearly empty. The service crew had refueled the wrong plane. Another fellow filled the tanks on Sunday and went to fly again on Friday. The tanks were nearly empty because of a leak.
- Don't switch tanks just before takeoff. On typical airplanes the engine can run for *two or three minutes* just using the fuel stored in the carburetor and engine sumps. That's just long enough to get you into big trouble if you use tank "A" for taxi and runup, and then switch to tank "B" for takeoff. What if tank "B" is contaminated? What if it is empty? What if there is a blockage in the lines? What if you accidentally select "Off" instead of "B" during the switch? Et cetera, et cetera.... If there is a problem with tank "B", you'd like to find out about it before starting your takeoff roll. If you absolutely must switch to tank "B" for takeoff, do a duplicate runup on that tank, and wait long enough to consume the fuel in the lines and sumps and prove that you are actually getting fuel from tank "B".
- By the same token, it isn't smart to switch to a new tank on final approach. Plan ahead; do your tank-switching at an altitude and at a location where if something bad happens you'll have a chance to do something about it.

15.3 Dealing with Emergencies

As suggested in the epigraph to this chapter, the first step in dealing with any in-flight emergency is to *fly the airplane*.

This sounds simple and obvious, but there have been far too many cases where the aircraft stalled or flew into the terrain because the pilots were too busy fussing with something that should have been only a minor distraction.

For instance, what should you do if the door comes open during takeoff? Answer: fly the airplane. No general aviation airplane I know of will crash because the door is ajar. The door will be held open about an inch or so. There will be enough suction to make it rather hard to close that last inch. There will be a bit of noise and a bit of a draft, but perhaps less than you might imagine. You might not even notice at first. In any case, the safest thing to do is to ignore all the details and return for landing. When you are safely stopped you can fiddle with the door as much as you like.

As another example, suppose you are just about to turn onto final approach when you notice that only two of the three landing gear are indicating "down and locked". What should you do? Go around! Do not try to debug the landing gear on final. For that matter, do not try to debug *anything* in the traffic pattern — it is too close to the ground and (usually) too congested. Get out of there. If there is a control tower, don't forget to tell ATC what's happening: "Tower, Five Seven Tango has some uncertainty with the landing gear.⁸ We'd like to leave the pattern while we investigate". Then climb to a reasonable altitude, away from the airport, and take your time fixing the problem.

1

Engine trouble (or other trouble) during the takeoff roll is discussed in [section 13.8](#), especially [section 13.8.4](#). Being lost, which might or might not be an emergency, is discussed in [section 14.6](#).

2

As part of the pre-takeoff briefing, remind yourself of the best-glide airspeed for the airplane you are flying. It varies a lot from plane to plane. Find this speed on the airspeed indicator and point to it. (Don't just remember the numerical value, because in an actual emergency you will probably be so excited that you lose your ability to remember numbers. Your ability to remember geometrical relationships will be less impaired.)

3

In retrospect, I wish I hadn't pushed quite so hard. A half-G pushover would have been more than sufficient, and would have kept the examiner from floating out of his seat. But basically I had the right idea.

4

i.e. the reciprocal direction on the same piece of pavement.

5

See [section 12.3](#) for more discussion of how to measure angles.

6

Note this method gives you control of your destination using at most two large-angle turns, whereas S-turns on final would require using four.

7

Remember the law of the roller-coaster: 9 feet per knot, per hundred knots.

8

On a training flight, you might want to let Tower know it is only a *simulated* problem, so they don't get unduly worried. When in doubt, ask the instructor, but usually you know it's a simulation because of the satanic grin on the instructor's face.

16 Flight Maneuvers

A small correction early is better than a large correction late.
— Aviation proverb

16.1 Fundamentals

During flight, you have quite a number of tasks and responsibilities:

- You are either accelerating, decelerating, or maintaining constant speed.
- You are either climbing, descending, or maintaining constant altitude.
- You are either turning left, turning right, or maintaining constant direction of motion.
- You are either slipping left, slipping right, or maintaining coordinated flight.
- You have control over the flaps, landing gear, various engine controls, et cetera.
- You must keep track of where you are, so you don't miss your destination, run into obstructions, or whatever.
- You need to keep track of weather conditions.
- You must keep watch at all times¹ to make sure you see and avoid other aircraft.
- Et cetera.

The first three items on this list are what I call the ``fundamentals" of maneuvering.² Simple maneuvers (including plain old straight and level flight) and even some quite complex maneuvers can be broken down into combinations of these three fundamental tasks. Of course, while you are maneuvering you still remain responsible for all the other items on the list.

Some of the maneuvers in this chapter are important parts of everyday flying. For instance, final approach requires lining up on a ``front window" ground reference. Flying the downwind leg of the airport traffic pattern requires paralleling a ``side window" ground reference. Oftentimes you or your passengers want to get a good view of some landmark, which requires turning around a point. If there is some wind (as there almost always is) you will need to correct for it.

The other maneuvers in this chapter, even though they are not directly practical, serve important pedagogical purposes. Chandelles and lazy eights are good illustrations of several of the points made in this book, including (a) the importance of angle of attack, (b) the relationship between angle of attack and pitch attitude, and (c) the behavior of the plane when its airspeed doesn't equal its trim speed. Some of these maneuvers may seem daunting at first, because they require doing several things at once. Fortunately, though, the ingredients are not particularly hard and can be learned separately.

16.2 Accelerating and Decelerating

This is a very important maneuver which has not always been sufficiently stressed during pilot training. The idea is to change speed while maintaining constant altitude, constant heading, et cetera.

Start from, say, cruise speed. Decelerate to V_Y . When you reach the new speed, set the engine controls and trim so that the plane will maintain that speed. After you have flown in this configuration long enough to convince yourself that everything is stable, decelerate to a speed well below V_Y (but with a reasonable margin above the stall). Again, stabilize the plane at the new speed. Then accelerate back to V_Y and stabilize. Then accelerate to cruise and stabilize. Iterate this a few times until you are sure you've got the hang of it.

You will have an easier time understanding how to use the throttle (especially at speeds below V_Y) if you keep in mind the concepts of *kinetic energy* and *power curve*. These are discussed at length in [section 7.2](#).

You will also want to keep in mind the relationship between trim and airspeed, as discussed in [section 2.6](#).

An interesting variation of this maneuver is to practice accelerating and decelerating with the flaps extended. (Make sure you observe the speed limit for flaps-extended operations, which is typically quite a bit lower than

for flaps-retracted operations.) This is interesting because on some planes, adding power with flaps extended causes a huge nose-up trim change; you will need to roll in some nose-down trim to compensate.

16.3 Phugoids

In flight it is fairly common for the airplane to find itself at an airspeed rather different from its trim speed. This situation will result in a *phugoid oscillation*, as discussed in [section 6.1.12](#). It is definitely worth seeing this behavior for yourself.

Start with an airspeed, say, halfway between V_Y and cruise. Pull back on the yoke until the airplane slows down about ten knots, and then let go. As discussed in [section 6.1.12](#), the airplane tries ``too hard" to return to its original airspeed, altitude, and attitude; it will overshoot and oscillate for several cycles.

From time to time during this maneuver, look at the airspeed indicator and altimeter. This will provide a good illustration of the law of the roller coaster (9 feet per knot, per hundred knots). See [section 1.2.1](#). This maneuver is also a good illustration of the principle of angle of attack stability, as discussed in [chapter 6](#).

Practice ``catching" the phugoid at various points in the cycle. That is, by pushing or pulling on the yoke, maintain constant altitude until the airspeed returns to normal. It is particularly interesting to catch it right when the airspeed equals the trim speed. By returning it to normal attitude at that moment, you can instantly end the oscillations.

If you use the wrong procedure (pushing on the yoke when the altitude is highest and pulling on the yoke when the altitude is lowest) you will just make the situation worse. This is an example of a *pilot-induced oscillation* (PIO). It is more common than you might think, and can cause serious trouble if it happens near the ground, as discussed in connection with evil zooms in [section 12.11.8](#) and [section 16.17.6](#).

16.4 Crabbing Along a Road

One of the most basic maneuvers involves choosing a ground reference such as a long, straight road and flying along it. The point of the maneuver is to practice perceiving and correcting for crosswinds, so choose the road so that there is a significant crosswind component.

Actually, correcting for the crosswind is the easy part. If the plane starts getting blown off to the left of the road, you will instinctively turn the plane a little to the right to compensate. The tricky part is to *notice* that you have done so. The situation shown in [figure 16.1](#) (crosswind from the left) seems quite normal. Similarly, the situation shown in [figure 16.2](#) (crosswind from the right) also seems quite normal. It is important to be able to perceive the difference. The outside world looks the same in both cases; the difference is that the alignment of the airplane has changed relative to the outside world.

[Figure 16.3](#) and [figure 16.4](#) show bird's eye views of the same situations.

You should always make a point of noting your direction of flight (which is aligned with the road in this case) relative to bolts on the cowling, marks on the windshield,³ and other parts of the airplane.

You should be especially alert to these perceptions during final approach, since you need information about the wind in order to prepare for a proper crosswind landing.


It also pays to notice the crosswind during the base leg. If the crosswind is trying to blow you toward the airport then you will have a tailwind on final and (most likely)⁴ a tailwind during landing. You might want to break off the approach and take a good look at the windsock before trying again.

These perceptions can give you rather precise information about the wind. The magnitude of the crosswind is proportional to the wind-correction angle and to your airspeed:


- At 60 knots one degree corresponds to 1 knot of crosswind.
- At 90 knots one degree corresponds to 1.5 knots of crosswind.
- At 120 knots one degree corresponds to 2 knots of crosswind.




[Figure 16.1](#): Crosswind from the Left



[Figure 16.2](#): Crosswind from the Right



[Figure 16.3](#): Crosswind from the Left — Bird's Eye View



[Figure 16.4](#): Crosswind from the Right — Bird's Eye View

16.5 Slipping Along a Road

Another useful maneuver is the following: Make sure you are at a safe airspeed. Line up on a road with a nice crosswind, as before. Now lower the upwind wing using the ailerons, and apply opposite rudder (i.e. push the rudder pedal on the downwind side). The idea is to establish a slip so that the airplane's axis *and* its direction of motion are both aligned with the road.

The slip will cause lots of drag. You will have to add power to maintain altitude. For goodness sake don't pull back on the yoke; you will be at a fairly low altitude (since this is a ground-reference maneuver) and you really don't want to stall in such a situation.

Make a note of how much bank angle and how much rudder pressure are needed for a given amount of crosswind. This varies considerably from one type of airplane to another. This knowledge comes in handy during crosswind landings; you don't want to wait until you are in the midst of a landing to figure it out.

16.6 Turns

The previous section discussed how to roll into and roll out of turns. Flying around in an established turn is relatively simple. You might need to deflect the rudder toward the inside of the turn (to compensate for the

long-tail slip effect) and deflect the ailerons toward the outside of the turn (to compensate for the overbanking tendency).

If you are turning to intercept a landmark, you need to think a little about how steep a turn to make and when/where to start the turn. It so happens that for any particular bankbank attitude angle, the turning radius depends on the *square* of your speed. A turn that consumes a tenth of a mile at 60 knots will consume nearly a mile at 180 knots.

speed	rate	radius	bank	load
(knots)	(°/sec)	(nm)	(degrees)	factor
60	10.5	0.09	30	1.15
75	8.4	0.14	30	1.15
90	7.0	0.20	30	1.15
105	6.0	0.28	30	1.15
120	5.3	0.36	30	1.15
135	4.7	0.46	30	1.15
150	4.2	0.57	30	1.15
165	3.8	0.69	30	1.15
180	3.5	0.82	30	1.15

[Table 16.1](#): Constant-Bank Turn

A *standard rate turn* is defined to be three degrees per second. This is what ATC expects when you're on an instrument clearance. It is also called a two-minute turn, because at that rate it takes two minutes to make a complete 360° turn. You can see from the following table that the bank angle required grows in proportion to the airspeed. Because of the changing bank, the radius of turn grows in proportion to the airspeed (not the square thereof).

You should figure out the bank angle that corresponds to a standard-rate turn for the airspeed(s) you normally use.

speed	rate	radius	bank	load
(knots)	(°/sec)	(nm)	(degrees)	factor
60	3	0.32	9.4	1.01
75	3	0.40	11.6	1.02
90	3	0.48	13.9	1.03
105	3	0.56	16.1	1.04
120	3	0.64	18.2	1.05
135	3	0.72	20.3	1.07
150	3	0.80	22.4	1.08
165	3	0.88	24.4	1.10
180	3	0.95	26.3	1.12

[Table 16.2](#): Standard-Rate Turn

[16.7](#) Coordination Exercises

Here is a good maneuver for learning about your plane's roll-axis inertia and adverse yaw, called ``coordinated wing rocking''. The procedure is: roll rather rapidly into a 45 degree bank to the left. Pause for a moment, then roll to wings level. Pause again, then roll 45 degrees to the right. Pause again, roll wings level, and repeat.

Refer to [chapter 11](#) for a discussion of various techniques for perceiving whether or not your maneuvers are accurately coordinated.

The rolls should be done sufficiently rapidly that significant aileron deflection is required. Do the maneuver at cruise airspeed, and then do it at approach speed and even slower speeds, so you can see how the amount of rudder required increases as the speed decreases. Do the maneuver while looking out the side (wings should go up and down like a flyswatter, with no slicing) and while looking out the front (rate of turn proportional to amount of bank, no backtracking on roll-in, no overshoot on roll-out). Pay attention to the seat of your pants.

You should do the maneuver two ways: once with large aileron deflection applied gradually, and once with large aileron deflection applied suddenly. The difference between the two demonstrates adverse yaw.

Here is another exercise. Unlike the previous one (which involved *coordinated* wing rocking) this one involves intentionally uncoordinated wing rocking. You bank the airplane but apply top rudder to keep it from turning. This is grossly uncoordinated, but it is amusing and educational because it lets you learn the feel of the controls and the response of the airplane. (Do not get in habit of starting turns in such a way.)

This uncoordinated wing-rocking exercise is related to (but not quite the same as) the Dutch roll oscillations discussed in [section 10.6.1](#). Both involve slipping to one side and then the other, like a Dutch kid on skates, making a series of slips (left, right, left, right) without much change in heading. The difference is that genuine Dutch roll oscillations involve a lot of yawing, while in the coordination exercise, you use the rudder to prevent any yaw. Let's call this the ``3/8ths hesitation roll'' since it resembles three eighths of an aerobatic 8-point roll.

Another uncoordinated exercise that is somewhat amusing and educational is as follows: keeping the wings level at all times, yaw the nose to the left with the rudder. Then raise the nose with the flippers. Then yaw the nose to the right with the rudder. Then lower the nose with the flippers, and repeat. Imagine you are drawing a rectangle on the sky in front of you, using the axis of the airplane as your pencil. Because of the slip-roll coupling described in [section 9.2](#), while pressing right rudder you will need to apply left aileron to keep the wings level. The purpose of this exercise is to illustrate yaw-axis inertia, yaw-axis stability, and yaw-axis damping. That is, you will notice that if you make a sudden change in rudder deflection, the nose will overshoot before settling on its steady-state heading. (Once again, the combination of controls used here is very different from proper turning procedure.)

[16.8](#) Familiarization Exercises; Configuration Changes

Imagine you are not completely familiar with the aircraft you are flying. You have just flown an instrument approach, and have broken out of the clouds about 150 feet above the runway. You are flying at 100 knots. Within the next 15 seconds or so, you need to slow down to 71 knots in preparation for landing. Therefore you take the following actions:

1. Pull the throttle to idle
2. Extend the flaps the rest of the way
3. Deploy the speed brakes⁵

Now imagine that those actions do not cause the airplane to slow down! You discover that on this airplane, each of those actions causes a nose-down trim change. The airplane pitches over and dives toward the ground at high speed. This is not good.

Therefore, in this airplane, a much better procedure would be to take the following actions:

1. Pull the throttle to idle and apply some nose-up trim to compensate.
2. Extend the flaps the rest of the way and apply some more nose-up trim to compensate.
3. Deploy the speed brakes and apply even more nose-up trim to compensate.
4. As you decelerate, apply yet more nose-up trim.

For any given airplane, you need to know how much trim it takes to compensate for each configuration change. This information is typically not provided by the Pilot's Operating Handbook. You need to obtain it empirically. Go to the practice area and do some experiments at a safe altitude.

First, just fly around for a while at normal cruise airspeed. This lets you see what the cruise angle of attack looks like; this information comes in handy on final approach, as discussed in [section 12.11.3](#).

You should also take this opportunity to learn how the airplane responds. Practice the basic maneuvers as

described in previous sections of this chapter. Acceleration/deceleration is worth practicing; some airplanes are much harder to slow down than others. Coordinated turns are worth practicing; different airplanes require different patterns of rudder usage. Nonturning slips are important for landings; you need to know how much yaw and how much drag is produced by a given amount of rudder pressure. Phugoids are definitely worth investigating; different airplanes respond differently.

Next, investigate the effect of the trim wheel. The wheel has bumps on it, which we can use as our unit of measurement. Move the wheel one bump, and see what effect that has on the airspeed. If you have electric trim, figure out how fast it moves (how many bumps per second).

Next, slow down to the airspeed you normally use in the traffic pattern. Again, get the airplane nicely trimmed and just fly around a while. Make a note of the angle of attack.

After the airplane is once again flying along, nicely trimmed at pattern speed, extend one notch of flaps. Maintain the same speed. Make careful note of how many bumps of trim it takes to maintain constant speed, compensating for the flap extension. Do not bother to maintain level flight. Leave the power setting along, and make a note of how much rate of descent is caused by the drag of the flaps. Also note how the pitch attitude changes; remember that extending the flaps changes the angle of incidence, as discussed in [section 2.4](#).

Do the same for each successive notch of flaps. In each case, make careful note of how much you have to move the trim wheel to maintain constant speed. Also observe the resulting rate of descent, and observe the change in incidence.

Do the same for other possible configuration changes (landing gear, speed brakes, et cetera).

After you have done that, investigate the effect of power changes. Determine how many RPM (or how many inches of manifold pressure) you need to remove in order to change from level flight to a 500 fpm descent. Also observe the effect that such a power change has on the trim speed.

Now, during the descent, check the effects of configuration changes again. You need two sets of observations: one using a power setting appropriate for level flight in the traffic pattern, and one using a power setting appropriate for final descent. In an ideal airplane, configuration changes would not affect the trim, but in a real airplane they do, by an amount that depends on the power setting.

At this point, you should be able to construct a crib card along the following lines:

- 300 RPM power reduction (clean), compensate with _____ bumps
- 300 RPM (approach configuration), compensate with _____ bumps
- first notch of flaps (level flight), compensate with _____ bumps
- first notch of flaps (descent power), compensate with _____ bumps
- second notch of flaps (descent power), compensate with _____ bumps
- third notch of flaps (descent power), compensate with _____ bumps
- extend gear, compensate with _____ bumps
- extend speed brakes, compensate with _____ bumps

where each of the blanks gets filled in with some positive number (for nose-up trim application) or negative number (for nose-down trim application). The exact values aren't important; the idea is to have enough information to prevent nasty surprises like the situation described at the beginning of this section.

Finally, fly around for a while slightly above minimum controllable airspeed, with flaps extended. See [section 16.16](#) for more discussion of slow flight procedures. Practice rocking the wings. Make sure you can bank the plane left or right, with reflexively correct use of ailerons and rudder.

Additional familiarization exercises are discussed in connection with landings in [section 12.11.4](#).


Familiarizing yourself with a new type of airplane can take a goodly amount of time, especially if you have modest total pilot experience. On the other hand, if you are just re-familiarizing yourself with the plane after a

period of inactivity, you can run through the maneuvers fairly quickly.

16.9 Turns around a Point

Turns are more challenging if you are trying to turn around a specific ground reference, maintaining a constant distance from it. If there is any significant wind (which there almost always is), this requires constantly changing bank angles.

The best way to analyze this situation is to begin by considering what happen if you do *not* make any correction for the wind. [Figure 16.5](#) shows three complete turns made using a constant bank angle.



[Figure 16.5](#): Turns Not Quite Around a Point

In the absence of wind, you would have performed three perfect circles around the southeasternmost tree in the orchard. However, since there is some wind, we can use the principle of relativity. Relative to the air, you have still made three perfect circles. However, the air itself has moved during the maneuver, carrying the whole pattern downwind. Therefore relative to the ground, we see the cycloid pattern shown in the figure.

To transform this pattern into one that is circular relative to the ground, you need a steeper bank at the points where you are headed downwind (e.g. point A and neighboring points), and a shallower bank at the points where you are headed upwind (e.g. point C and neighboring points). As you can see from [table 16.3](#), the effect can be fairly large.

speed	rate	radius	bank	load
(knots)	(°/sec)	(nm)	(degrees)	factor
60	2.9	0.33	9	1.0
75	3.6	0.33	14	1.0
90	4.3	0.33	19	1.1
105	5.0	0.33	26	1.1
120	5.7	0.33	32	1.2
135	6.4	0.33	39	1.3
150	7.2	0.33	45	1.4
165	7.9	0.33	50	1.6
180	8.6	0.33	55	1.7

[Table 16.3](#): Constant-Radius Turn

If you fly the maneuver at 90 K_{IAS}, your groundspeed will vary from 105 (downwind) to 75 (upwind). That's a ratio of 1.4 to 1. Let's assume you remain 1/3rd of a mile from the landmark, since that is the distance to which the table applies. The speed in the left-hand column of the table should be taken as a *ground* speed, since we want the radius to remain constant as seen from the *ground*. The table tells us the required bank angle will vary from 26 degrees at point A to 14 degrees at point C.

At points B and D in the figure, the bank angle will be the same as in the no-wind case — but you will need apply wind corrections to your heading, as discussed in [section 16.4](#).

16.10 Eights Around Pylons

Eights around pylons are performed by flying turns around a point clockwise around one pylon, and counterclockwise around another pylon, as shown in [figure 16.6](#).




Figure 16.6: Eights Around Pylons

You need to choose the right place to roll out of the turn and begin the straightaway section, so that the two circles will be the same size. It may help to visualize the desired figure-eight shaped ground track on the ground, and then just follow that track.

This maneuver is not to be confused with eights *on* pylons (which are discussed in [section 16.13.2](#)).

16.11 Chandelles

A *chandelle* is a stylized climbing turn. The key elements are:

- There is a total heading change of 180 degrees.
- During the first 90 degrees, there is a constant bank and smoothly increasing pitch attitude.
- During the second 90 degrees, there is a constant pitch attitude and smoothly decreasing bank.
- Climb power is used.
- At the 180 degree point, the wings are level and the airspeed is just above the stall.

Normally the maneuver is entered from level flight. If your airplane's manufacturer has specified a recommended entry speed, use that. Otherwise, cruise airspeed should do nicely.

The higher the entry speed, the greater the altitude gained during the maneuver. You can dive if necessary to achieve the chosen entry speed, but be careful not to overspeed the engine. A high entry speed is absolutely not required.

A chandelle is in some senses a "maximum performance" maneuver, but altitude gain is not one of the things that you are expected to maximize. (If people wanted absolute maximum altitude gain, they would use a rather different sequence of bank and pitch attitudes.) The maneuver is judged primarily on the accuracy and smoothness of the pitch and bank maneuvers.

The maneuver emphasizes headings and attitudes. You should use ground references to judge the correct headings, but you shouldn't bother to remain over a particular point or to correct headings for wind drift.

You have some discretion when selecting the initial bank angle. Usually 30 degrees works fine. If the bank is too shallow, during the second half of the maneuver you will find that the airplane has decelerated to its final speed before the turn is completed; ideally the final speed and the final heading should be reached simultaneously. Happily, since the airspeed is changing only rather slowly at the end, this is relatively easy to arrange.

The end of the maneuver depends on airplane performance. If your airplane has more than enough power to sustain level flight at stalling angle of attack, you are in luck. At the end of the maneuver you should lower the nose and accelerate at constant altitude.

If your airplane cannot sustain level flight at stalling angle of attack, you should arrange the timing so that at the end of the maneuver you are momentarily in level flight, at the top of the climb. Then you should lower the nose and dive gently to obtain an airspeed that will permit acceleration in level flight. Then level off and accelerate to a normal speed. You will need more skill and judgment than you would in a more powerful plane.

If you want to learn to do chandelles, it may help to divide the maneuver into separate "climb" and the "turn" components. It is sometimes useful to analyze and practice these components separately.

The second half of the climb contains an interesting lesson. The pitch attitude and power setting are constant, but the result is very far from being constant performance. The angle of attack is increasing, the airspeed is decreasing, and the rate of climb is decreasing.

This second part of the maneuver begins with the airplane climbing rapidly. The climb angle is, intentionally, unsustainable. The airplane will nevertheless climb in the short run. For a while, it can climb by cashing in airspeed, according to the law of the roller coaster.

As the airspeed decreases, the airplane must fly at an ever-higher angle of attack in order to support its weight. Since the pitch attitude is being held constant, this means that the direction of flight must be bending over. This is illustrated in [figure 2.11](#) in [section 2.10](#).

This should drive home the lesson that pitch attitude is not the same as angle of attack, and that angle of attack (not pitch attitude) is what directly determines performance.

You should not attempt to micro-manage the altitude during a chandelle. You should maintain the chosen pitch attitude and let the airplane's intrinsic vertical damping (and energy budget) take care of the vertical motion.

The choice of pitch attitude with which you begin the second half of the chandelle is obviously critical, since you will be stuck with it for the rest of the maneuver. If it is too nose-high, the airplane will decelerate too quickly and you will run out of airspeed before the turning part of the maneuver is completed. Conversely, if the pitch attitude is too low, you will have airspeed left over at the end of the turn. The right answer depends on the performance of the airplane (and on the timing of the turning part of the chandelle). The answer can be determined by trial and error. About 15 degrees is a good initial guess for typical training airplanes.


Now let's examine the turning component of the chandelle. Again, the second half is the interesting part. The second half, if properly performed, will take a certain amount of time. You have to roll the wings level, using a uniform roll rate over that time. If you roll too slowly, the airplane will turn through 90 degrees before the rollout is completed. Conversely, if you roll too quickly you will run out of bank before the 90 degree turn is completed. At each instant, you should estimate the amount of turn remaining and the amount of bank remaining, and fudge the roll-rate accordingly. As always, a small correction early is better than a large correction late. It is useful to practice this a couple of times in level flight, before combining it with the climbing component.

When performing the complete maneuver (climbing and turning together) there is one more wrinkle: Remember that rate of turn depends not only on bank angle but also (inversely) on airspeed. Since the airspeed is decreasing during the maneuver, you must take this into account when planning the roll rate for the complete maneuver.

Also, as the airspeed decreases you will need progressively more right rudder to compensate for the helical propwash, and progressively more right aileron to compensate for the rotational drag on the propeller blades. Furthermore, remember that adverse yaw and the effects of yaw-axis inertia become more pronounced at low airspeeds (as always). Maintain proper coordination (zero slip) at all times.

16.12 Lazy Eights

The *lazy eight* derives its name from the motion of the airplane's axis during the maneuver. In particular, imagine that the airplane is at a very high altitude, so we don't need to worry about the ground getting in the way. Further imagine that the airplane is centered in a cylinder of paper, 10 miles in diameter and 5 miles high. Also imagine that the airplane carries a very long pencil sticking out the front, aligned with airplane's axis. During the course of a lazy eight, the pencil will draw a giant figure eight, sideways, on the paper.



[Figure 16.7](#): Lazy Eight

[Figure 16.7](#) shows some of the details. Start at point *A*, in level flight. Pull the nose up. Gradually start banking to the right. At point *B*, stop pulling the nose up; let it start going down. Keep the bank; keep turning to the right. At point *C*, the pencil slices through the horizon. The body of the pencil is horizontal, while its tip is moving down and to the right. Start rolling out the bank. Point *D* is the lowest pitch attitude. The bank is about half gone; keep rolling it out. At point *E* the pitch attitude and the bank attitude should be level. Pull the pencil straight up through the horizon. Start rolling to the left. At point *F*, start letting the pitch attitude back down again. At point *G*, the pencil-point slices through the horizon again, this time moving down and to the left. Start rolling out the bank. Point *H* is the lowest point in the leftward stroke. By the time you return to point *A*, the pitch and bank attitudes should be level again. Pull the pencil straight up through the horizon again, and repeat the maneuver.

For the next level of refinement, arrange the timing and the bank angles so that point *B* is 45 degrees of heading away from point *A*; point *C* is at 90 degrees, point *D* is at 135 degrees, and point *E* is at 180 degrees.

For the next level of refinement, arrange the push/pull forces so that points *B* and *F* are about 20 degrees above the horizon, and points *H* and *D* are about 20 degrees below the horizon.

Note that up to this point we have not mentioned anything about altitude or airspeed. This is primarily an *attitude* maneuver, and you should learn it in terms of attitudes.

When learning the maneuver, it helps to separate the ``up/down'' part from the ``left/right'' part.

The left/right part of the maneuver is quite simple. You just very gradually roll into a turn to the right, then very gradually roll out. You continue the roll so it becomes a turn to the left, and then gradually roll out.

The up/down part of the maneuver is almost as simple. You just pull the nose above the horizon for a while, then lower it to the horizon; let it go below the horizon, then pull it back to the horizon and repeat.

One tricky part about combining the left/right part with the up/down part: the vertical motion goes through *two* cycles (ascending, descending, ascending, descending) while the horizontal motion is going through only one (rightward, leftward).

To get a deeper understanding of the maneuver, we must think a little about the altitudes and airspeeds.

During the whole quadrant from *A* to *C*, the nose is above the horizon. The airplane is climbing and decelerating. Therefore *C* is the point with the highest altitude and the lowest airspeed. Point *C* has a high altitude even though we (correctly) drew it in the figure on the same line as point *A*. That is because the maneuver is defined in terms of attitude, not altitude, and we imagine that the paper on which the lazy eight is drawn is so far away that the pencil has lots of leverage — the angle matters a lot, and the altitude matters hardly at all.

To you, the low airspeed at *C* is more immediately noticeable than anything else. The airplane is below its trim speed, so the nose wants to drop all by itself. At this point you will not need to push on the yoke; you just need to reduce the back pressure to let the nose go down at the desired rate.

During the whole quadrant from *C* to *E*, the nose is below the horizon. The airplane is descending and accelerating. Therefore point *E* has a much lower altitude than point *C*, and indeed should be level with point *A*.

The second ascending/descending cycle (from *E* back to *A*) should be pretty similar to the first.

The commercial-pilot Practical Test Standard requires that you return to your initial altitude and airspeed every time you pass point *A* and point *E*. You might hope that this would happen automatically if you leave the throttle setting alone, relying on the law of the roller coaster. But that hope is in vain, for the following reason.

Normally you start the maneuver at a speed well above V_Y , with a power setting appropriate for level flight at this speed. Now suppose you fly a nice smooth symmetric maneuver that returns to the original airspeed. The maneuver starts with a pull, and at all times you will have an airspeed at or below the initial airspeed. You will be flying the maneuver at more-efficient airspeeds, closer to V_Y .⁶ You will gain energy. You will gain altitude. If you try to fix the altitude by diving, you will end up with excess airspeed. The only way to make things come out even is to fly the maneuver using a slightly-reduced power setting. This is most noticeable in

airplanes with big engines and long wings, where the normal operating speeds are large compared to V_Y .

This maneuver contains a very nice lesson about the principles of flight. Much of the vertical part of the maneuver can be considered a "controlled phugoid". In particular, during the phase from *B* to *D* the nose is dropping but you are not pushing it down — indeed you are maintaining back pressure as you gently lower the nose. The feeling is sort of like the feeling you get when lowering a heavy object on a rope, and is quite striking.

This should drive home the message that the airplane is definitely not trimmed for a definite pitch attitude — it is trimmed for a definite angle of attack (or, approximately, a definite airspeed). At point *C*, among others, the airplane is well below its trim speed, so it wants to dive and rebuild its airspeed.

As the final level of refinement, you should make the altitudes at points *A* and *E* come out equal. You can do this by fudging the attitudes and/or power settings at strategic points in each half-cycle.

You have considerable discretion as to the steepness of the banks. Increasing it just speeds up the whole maneuver. A typical choice is to have 30 degrees of bank at points *C* and *G* (the points of maximum bank).

A lesser bank is also fine, but then you will want to choose a lesser nose-high attitude at points *B* and *F*. This is because you will be spending more time ascending, and you don't want to run out of airspeed. Make sure the airspeed at points *C* and *G* is 5 or 10 percent above the stall.

As with the chandelle, you will have to work a bit to maintain proper coordination. There is nothing surprising — just a wide range of roll rates and a wide range of airspeeds.

16.13 Eights on Pylons

The "eights on pylons" maneuver is required on the commercial and flight instructor practical tests. Being able to do this maneuver well, especially if there is a wind, definitely demonstrates that you can control the airplane around all axes at once.

This maneuver is not to be confused with eights "around" pylons (which are discussed in [section 16.10](#)). The ambiguous term "pylon eights" should be avoided.


16.13.1 Turns on a Pylon

Before we cover the "eights on pylons" maneuver ([section 16.13.2](#)), we need to discuss a little theory. We begin by considering turns on a (single) pylon.

The idea is simple: Imagine a pointer that pokes through the plane from wingtip to wingtip, parallel to the pitch axis; you want this pointer to remain pointed directly at the base of the pylon. This is quite a restriction; it means that at each point in the maneuver your bank and heading are completely determined by your altitude and position relative to the pylon. The only thing that makes the maneuver possible at all is that you are free to adjust your altitude.

*** No-Wind Case**

In the absence of wind, the maneuver will work at a particular altitude — the so-called pivotal altitude — and not otherwise. Interestingly, the pivotal altitude does not depend on what you choose as your distance from the pylon. As shown in [figure 16.8](#), if you start close to the pylon, you will have a large bank angle and therefore a lot of *G*s. But since you are close to the pylon, the circle will be small, and you will need a lot of *G*s in order to change the airplane's velocity (from northbound to southbound and back) in the small time available. In contrast, if you start out far from the pylon, the bank will be shallow, and you will pull a smaller number of *G*s for a longer time.



[Figure 16.8](#): Turns on a Pylon — Pivotal Altitude

The pivotal altitude is proportional to the square of the airspeed: 0.0885 feet per knot squared, or 885 feet

per (hundred knots) squared.

If you happen to be above the pivotal altitude, the airplane will be banked too steeply and will turn too quickly. Your sight-line past your wingtip, which is supposed to be pointed at the pylon, will be swept backward and will appear to fall behind the pylon. Or to say it the other way, the pylon will appear to be moving ahead of where you want it to be. The solution is to descend. At the lower altitude your bank will be less, and the problem will correct itself. Any airspeed you gain during the descent can only help you by further reducing the rate of turn.


Conversely, if you are too low, the bank will be too shallow and the pylon will appear to fall behind where you want it to be.

The rule is simple: go down to speed up and ``catch" the pylon; go up to slow down and ``wait for" the pylon.

You may be tempted to use the rudder to swing one wingtip a little bit forward or backward, but this defeats the purpose of the maneuver and is *not* the correct procedure.


* Windy Case

In the presence of wind, the pattern is no longer a perfect circle. In fact, the ground track is an ellipse with the pylon at one focus. You are nearest the pylon when the airplane is headed directly downwind. This gives max bank when flying downwind, which makes a certain amount of sense — you want to bank more steeply when the groundspeed is highest. This is shown in [figure 16.9](#).



[Figure 16.9](#): Turn On Pylon — Headings

The wind also prevents you from flying the pattern at constant altitude (for reasons that will be discussed below). The altitude is highest when the airplane is headed directly downwind. This is shown in [figure 16.10](#). Once again, this contributes to creating max bank when flying downwind, which makes sense.



[Figure 16.10](#): Turn On Pylon — Altitudes

There are two strategies, depending on how much the plane speeds up when it descends.

a

If you fly the pattern at high speed (i.e., well above V_Y), then tiny changes in airspeed will give you plenty of up-and-down action. I call this the constant-airspeed case.

b

If you fly the pattern at a speed near V_Y , then changing the airspeed has only a small effect on the long-term power required — all you are doing is making a one-time exchange of potential energy for kinetic energy according to the law of the roller-coaster. I call this the constant energy case.

The typical case will lie somewhere in between; fortunately the answers in the two cases are not very different.

a

In the constant-airspeed case, the ground track is a mathematically perfect ellipse. The altitude turns out to be inversely proportional to your distance from the pylon, which can be a surprisingly large

excursion even in moderate winds.


b

In the constant-energy case, the ground track deviates only imperceptibly from an ellipse (the distance deviation is less than 1%, even when the wind is 30% of your airspeed). The altitude variation (as a percentage) is about one-third as large as the variation in distance from the pylon.

When going upwind, you need to turn a lot slower. There are three factors at work:

1. you are farther away, so the bank angle is less (by geometry);
2. you are lower, so the bank angle is less (also by geometry); and
3. in the constant-energy case, you are going faster (making more forward progress per unit turn).

The first two factors are diagrammed in [figure 16.11](#). In the constant-airspeed case factor 1 does half the job and factor 2 does the other half. In the constant-energy case they all three divide the job, roughly in the ratio 50% : 20% : 30%.



[Figure 16.11](#): Turn On Pylon — Bank Geometry

By geometry, the angle of bank is inversely proportional to the distance r from the pylon. It is also proportional to height. In the constant-airspeed case, the height is itself inversely proportional to r . Combining these, you get that the airplane is ``attracted'' toward the pylon with an acceleration that goes like $1/r^2$.

You may recognize this situation from astronomy: an inverse-square central force. So it's not surprising you get a Keplerian ellipse. The airplane will sweep out equal areas in equal time, and its angular momentum about the pylon will be constant.

In the zero-wind case, the pivotal altitude is simply proportional to groundspeed squared. Several well-known books try to argue that on the upwind leg of the turn on pylon, the groundspeed is lower, so the altitude should be lower. That is a false explanation (even though the altitude is indeed lower there). The actual altitude change is much less than you would predict by the groundspeed argument (by a factor of 2 in the constant-airspeed case and by a factor of 4 or so in the constant-energy case).

You may wonder how this can be — how can the airplane keep the wing on the pylon if it is not at the pivotal altitude? The answer is simple: we are not trying to fly a circular pattern. Recall that if you are *above* the pivotal altitude, the airplane will spiral toward the pylon. This is exactly what is happening in half of the elliptical pattern — the airplane is above the pivotal altitude and flying gradually closer to the pylon.

Why is the center of the pattern shifted crosswind rather than downwind of the pylon? For sake of discussion, let's divide the pattern in half along the long axis (which includes the pylon). If the airplane is positioned to windward of this line, it is subject to a crosswind from outside the pattern, which tends to drift the plane sideways closer to the pylon, making the bank steeper. This effect occurs throughout the windward half, so the plane is *closest* and *steepest* when it crosses from the windward to the leeward half (at which point it is headed directly downwind).

For these turns *on* pylons (unlike turns *around* pylons), there is nothing you can do to prevent the plane from being blown sideways. Consider the point where the plane is directly upwind of the pylon. The heading is constrained to be directly across the wind. The pilot cannot crab into the wind. Therefore the plane will be blown toward the pylon.


By the same token, whenever the airplane is on the leeward side of dividing line, it is subject to a crosswind from inside the pattern, which tends to drift the plane sideways farther from the pylon and hence make the bank shallower. The effect is cumulative, so the plane is *farthest* and *shallowest* when it crosses from the leeward to windward half (at which point it is headed directly upwind).

Also, draw a line from the pylon to a generic point on the ellipse. The wings of the plane, at that point, will lie on that line; the heading of the plane will be perpendicular to that line. Except for the two special points at the ends of the ellipse, the heading will not be tangent to the ellipse; the angle between the heading and the

tangent is precisely the crosswind correction angle. You will note that the plane is always crabbed into the wind. This can be seen in [figure 16.9](#).

16.13.2 Eights on Pylons

The eights-on-pylon maneuver consists of a turn on one pylon followed by an opposite-direction turn on another pylon, as shown in [figure 16.12](#). The two-pylon maneuver adds the complexity of planning when to shift from one pylon to the other, but is actually *easier* to perform because you can use the straightaway between turns to recover from any small errors.



[Figure 16.12](#): Eights On Pylons

You don't want to pick pylons that are too close together. You do want pylons that are crosswind from each other, so that the pattern will be symmetric. It is good to enter on a downwind heading, as shown in the figure, so that your first turn will be your steepest turn. Maintain coordination; don't fudge things with the rudder.

In flight, you can follow these simple rules:

1. If the pointer is above or below the base of the pylon, it's easy to fix; just change your bank angle.
2. If the pointer is behind the pylon, go down to accelerate and ``catch" the pylon.
3. If the pointer is ahead of the pylon, go up to decelerate and ``wait for" the pylon.

In principle, these rules are all you need to know. However, the other information in this section makes your job 1000% easier. It allows you to anticipate the required altitude changes and the elliptical ground track. Anticipating the required actions is easier than waiting until there is an error and then making corrections.

16.14 Changing Headwinds and Tailwinds

In some ways, an airplane performs differently when going downwind as opposed to upwind — and in other ways it doesn't. There are a lot of misconceptions about both halves of this statement.

16.14.1 Steady Wind

Let us first consider the situation where there is a steady wind; that is, a wind that does not vary with time or with altitude.

Maneuvers relative to a ground reference will be different when headed downwind as opposed to upwind.

For instance, the airplane will climb and descend at a steeper *angle* (in terms of altitude per mile over the ground) when headed upwind.

Similarly, a constant-radius turn relative to a ground reference will require a steeper bank on downwind and a shallower bank on upwind.

Maneuvers that do not involve a ground reference will be unaffected by the wind.

For instance, the airplane will climb and descend at a *rate* (in terms of altitude per minute) that is independent of the wind.

Similarly, a constant-radius turn relative to a cloud will require the same angle of bank throughout the maneuver.

The point is that the airplane, the cloud, and the airmass are one big uniform moving system. By Galileo's principle of relativity, the

overall uniform motion doesn't matter.

Note that *obstacle clearance* is an important ground-reference maneuver. Your rate of climb is unaffected by the wind, but your *angle* of climb is affected. You can climb at a steeper angle on an upwind heading.

Finally, consider ground observers' perceptions. There are some maneuvers, such as an aerobatic loop, that *should not* be corrected for the wind. Imagine you are using a smoke generator. You want the smoke to form a nice round loop. Like the cloud mentioned above, the smoke is comoving with the air, so the overall wind speed shouldn't matter. However, especially if the smoke generator is turned off, the maneuver will *appear* different to an observer on the ground. This appearance does not (and should not) matter to the pilot in the cockpit, but it does matter if you are on the ground piloting a radio-controlled model, or judging an aerobatic contest.

There are several good reasons for being *aware* of your groundspeed, including:

- You need it for navigation, as discussed in [section 14.2](#).
- If you are flying cross-country and the groundspeed is lower than you planned for, recalculate your arrival time and re-appraise your fuel situation. All too many people run out of fuel because of unexpected headwinds.
- If you are about to land and the groundspeed seems abnormally high, you should consider the possibility that you have a tailwind. Go around, check the windsock, and try again.

On the other hand, during turns and other maneuvers, it would make absolutely no sense to try to maintain constant groundspeed.

16.14.2 Albatross Effect: Winds that Vary with Altitude

In the real world, the wind almost always changes with altitude. In particular, it is very common to find that the wind at ground level is blowing in the same general direction as the wind at 3000 feet AGL, but at a much lower speed. This is because of friction between the air and the surface.

Most of this frictional windshear is concentrated at the lowest altitudes. At low altitudes, it is common to see a windshear of several knots per hundred feet, while at enroute altitudes (several thousand feet AGL) it is more typical to see a windshear of a few knots per *thousand* feet.

Wooded areas, tall buildings, and/or steep hills upwind of your position can create particularly sharp shear layers.

On top of this, frontal activity (especially warm fronts) can cause very large windshears that are more complicated and less predictable than the normal, every-day frictional wind shear. I once was making an approach to a rather short, obstructed field. The windsock indicated that I had five or ten knots of headwind on the chosen runway, but the airplane acted as if I had at least 20 knots of tailwind on final. Even with zero engine power and full flaps I could not get the airplane to descend steeply enough to stay on the glide slope. Three approaches in a row ended in go-arounds (which allowed me to carefully check the windsock three times). On the fourth approach, by anticipating the windshear, I was able to make a reasonable landing. About half an hour later the surface wind shifted 180 degrees. This was consistent with the forecast warm front.

The situation in this anecdote (increasing headwind on final approach) is very atypical. The other 99.9% of the time there is a *decreasing* headwind as you descend on final. For the same reason, you expect to see an increasing headwind as you climb upwind on initial departure.

Let's analyze the effect of windshear. Suppose you start out at point *A*, and fly to point *B* where because of a windshear there is more headwind (or less tailwind). If the windshear is sudden, you will notice a sudden increase in airspeed. The windshear has added something to your energy [Z](#) budget. If the shear is more gradual, the airplane (because it is trimmed for a definite angle of attack) will probably convert the extra airspeed into extra altitude, but you will still wind up at point *B* with more energy than you would have without the windshear. During climb, this is great. It makes it look like you have a more powerful engine. (On short-field approach, as in the atypical anecdote above, this is not so good; the airplane behaves as if your engine were producing a fair amount of power even with the throttle at idle.)

We can apply the same line of reasoning to the opposite case: suppose you start out at point *C* and fly to a point *D* where (again because of a windshear) you have less headwind or more tailwind. This means you will

arrive at point *D* with less energy than you would have without the windshear. This commonly happens on approach, where you are descending into a decreasing headwind. This might be a good thing, allowing you to fly a steeper approach. On the other hand, a sudden decrease in headwind could rob you of energy at a critical time.

The energy that comes from an increasing headwind can be put to good use. I call it the *albatross effect*. The albatross is a huge bird that spends its life flying over the oceans of the world. It rarely needs to flap its wings, but it doesn't soar in updrafts the way hawks do. Instead, the albatross flies a figure-eight pattern in the shear zone near the surface, climbing into an increasing headwind on the upwind legs and descending into a decreasing tailwind on the downwind legs — gaining energy both ways.

Note that on a typical approach, the wind makes the angle of descent steeper in two ways:

1. the groundspeed is lower, due to the *average overall* headwind (as discussed in the previous section), and
2. the rate of descent is faster, due to the *decreasing* headwind (albatross effect).

16.14.3 Turning Downwind; Energy Budget

The previous section discussed how you could gain or lose energy *due to a windshear*. In this section, we return to considering only a steady wind, and discuss what happens if you convert a headwind into a tailwind simply by turning the airplane.

Let's consider the scenario described in [table 16.4](#).

true airspeed	100 knots
initial heading	north
final heading	south
time spent turning	1.2 min = .02 hour
mass of airplane	1 ton
wind speed	20 knots
wind direction	from the north

[Table 16.4](#): Downwind Turn Scenario

Let's calculate the energy and momentum twice, as shown in [table 16.5](#). In the ``balloon" column everything is measured relative to an observer in a balloon (comoving with the air mass), and in the ``ground" column everything is measured relative to an observer on the ground.

	balloon	ground
initial momentum	100	80
final momentum	-100	-120
change in momentum	-200	-200
average N-S force	10000	10000
initial energy	5000	3200
final energy	5000	7200
change in energy required	0	4000
N-S distance during turn	0	.4
energy provided by wind	0	4000

[Table 16.5](#): Downwind Turn Analysis

Here's what the first four rows mean: The momentum is calculated using the usual formula: mass times velocity. (The units here are rather strange, tons time knots, but it's OK as long as consistent units are used throughout the calculation.) The North-South component of the average force is just the change in momentum divided by the time. We see that although the initial and final momenta appear different in the two columns, the change in momentum is the same. This upholds Galileo's principle of relativity: the force required to turn the airplane is independent of the frame of reference.

Here's what the last five rows mean: The energy is calculated using the usual formula: one half of the mass

times velocity squared. According to the ground observer, the airplane needs to gain quite a lot of energy during the turn. You may be wondering where this energy comes from. Obviously it does not come from the airplane's engine. Actually it gains energy the same way a baseball gains energy when it is struck by a bat. You know that although a ball does not gain any energy when it bounces off a stationary wall, it does gain energy when it bounces off a fast-moving bat. The energy gain is force times distance (counting only distance in the same direction as the force). According to the observer in the balloon, the force of the turn is (at every instant) perpendicular to the direction of the force, so there is no energy gain. Meanwhile, according to the observer on the ground, the *wind* moves the airplane 0.4 miles in the North-South direction during the turn, and turning the airplane requires a huge force in this direction. This effect — the airplane being batted by the wind — supplies exactly the needed energy. Again, we see that the principle of relativity is upheld: the energy budget works out OK no matter what frame of reference is used.

Note that if you overlooked the bat effect you would fool yourself into thinking that turning downwind caused a huge energy deficit. It doesn't. Don't worry about it.

16.14.4 Summary: Changing Headwinds and Tailwinds

- For ground-reference maneuvers, a steady wind has a direct effect.
- For other maneuvers, a steady wind has no effect on the airplane or on the pilot in the cockpit. However, the maneuvers will appear different to ground-based observers.
- In the presence of windshears, you can gain or lose energy due to the albatross effect. In real life, this means for instance that you will get slightly better performance climbing into the wind. This gives you a reason to turn downwind a little later than you otherwise would.
- To a ground-based observer, the airplane actually does have more energy on downwind. It picks up this energy by being batted by the wind during the turn.

16.15 Remarks: Ground Reference Maneuvers

16.15.1 Accounting for the Wind

Throughout each flight — and certainly before starting any ground reference maneuvers — you should have in mind a good estimate of the speed and direction of the wind. There are various ways you can figure this out

- Remember the ``winds aloft" forecast. Sometimes it's even right.
- ATIS and AWOS broadcasts give the surface winds.
- The airport windsocks give information about surface winds.
- Ordinary flags provide similar information.
- The smoke or vapor from smokestacks is an excellent indicator of the winds near the ground and sometimes winds aloft.
- If you see ripples on a pond at one side and not the other, the wind is very likely blowing from the unrippled side toward the rippled side. Also, the texture of the ripples generally runs crosswise to the wind.
- Last but not least, you can note the amount of wind correction needed to perform ground-reference maneuvers.

It is a good idea to know the wind *before* starting a maneuver (rather than trying to figure it out ``on the fly"). It really helps to be able to plan the maneuver and anticipate the necessary wind corrections.

16.15.2 Entry Strategy

It is a good idea to begin ground-reference maneuvers such (as turns around a point) a downwind heading, as shown in [figure 16.5](#), so that your first bank will be your steepest bank. You don't want to be a position where (late in the maneuver) you must choose between abandoning the effort or using an excessive bank angle.

16.15.3 Visual Reference

It really helps to have a precise visual reference for pitch and yaw, as discussed in [section 11.5.2](#).

You can use your finger and/or a mark on the windshield, as illustrated in [figure 11.3](#). If you can't find a suitable mark on the windshield, you can make one.

The reference should be directly in front of your dominant eye. It is a common mistake to choose a mark on the cowling. Such a mark is below where it should be, and tempts you to use too much rudder when rolling into right turns, and too little rudder when rolling into left turns. It is another common mistake to choose a reference point that is on the centerline of the airplane. Assuming your eye is quite a bit to the left of the centerline, your sight line through this point is very far from being parallel to the axis of the airplane. This tempts you to make diving left turns and climbing right turns.

As you become more experienced, you won't need to use your finger or an explicit mark on the windshield; you can just *imagine* where the reference point must be. Just make sure you use a point directly in front of your dominant eye.

16.15.4 Checklist

You want to take a systematic approach to all maneuvers. John Beck teaches a ``mental checklist" for ground reference maneuvers:

- Pick a mark on the windshield; trace a line along the horizon.
- Check for traffic.
- Check your ground reference.
- Check your instruments.

Repeat this list to yourself over and over again as you do the maneuver. Chant it aloud if you wish. Doing each thing as you say it not only keeps you from overlooking something, but also gives a nice rhythm to the work.

16.16 Slow Flight

If you are not proficient in handling the plane at low speeds, you have no business trying to land the plane.

To begin a practice session, go up to a safe altitude and make sure there are no other aircraft nearby. Decelerate to a speed, say, 15 knots above the stall speed. Once you are comfortable with this, reduce the speed another 5 knots. Again, once you are comfortable, reduce the speed another 5 knots.

During the maneuver, you should

- Maintain coordination — keep the ball in the center.
- Maintain a definite altitude.
- Watch out for other traffic. Your pitch attitude will be so high that it will be difficult or impossible to see over the nose, so you should change heading every so often and look around.
- Between turns, maintain a definite heading — don't let the nose wander willy-nilly.
- Keep an eye on the engine gauges — there are some aircraft that will overheat if you spend too much time in a low-air-speed, high-power configuration.

16.16.1 Airspeed and Altitude

As discussed in [section 7.3](#) and elsewhere, it would be OK to use the yoke to control altitude *if* you were on the front side of the power curve *and* you were willing to accept an airspeed excursion. However, during this slow flight maneuver, you definitely are not on the front side of the power curve and you definitely cannot tolerate airspeed excursions. Therefore you will need to use the yoke (and trim) to control airspeed, and once you've got the desired airspeed, you will need to use the throttle to control altitude. (To adjust airspeed at constant altitude, you will need to use the throttle and yoke together, as discussed in [section 16.2](#).)

16.16.2 Yaw and Roll

Remember that the airplane is optimized for cruise flight. During cruise, you can fly straight and level with little or no control force, and you can make gentle turns with little or no use of the rudders, using ailerons alone.

In contrast, during slow flight

- You will need steady rudder deflection to overcome the helical propwash effect.
- You will need steady aileron deflection to overcome the rotational drag of the propeller.
- You will need considerable rudder deflection whenever the ailerons are deflected, to deal with adverse yaw and roll-axis inertia.

Because (as discussed in [section 5](#)) there will be very little roll damping, you will need to apply lots of little aileron deflections to maintain wings-level flight, especially in the presence of turbulence.

16.16.3 Procedures and Perceptions

Make a note of the pitch attitude that corresponds to level flight at minimum controllable airspeed (with and without flaps). Note the pitch attitude of the nose against the forward horizon, and the wingtip against the lateral horizon. This information will come in very handy during landing, as discussed in [section 12.11.3](#).

Practice rocking the wings. Make sure you can bank the plane left or right, with reflexively correct use of ailerons and rudder. Practice making turns to a precise heading.

Practice diving 50 feet. That is, push the nose down a few degrees (not so much that you experience negative *G* loads), dive for a few seconds, and then pull back and level out. Make a note of how much airspeed you gain by diving 50 feet. This information will come in handy during stall recoveries, as discussed in the next section.

16.17 Stall Practice

16.17.1 Preliminaries

- It should go without saying, but here goes: Make absolutely sure there are no other airplanes near you during stall practice. In particular, you will need to make frequent clearing turns to rule out the possibility that there are some folks behind and below you, who might be very surprised and annoyed if your drop down onto them.
- Make sure you practice stalls at an altitude that gives a generous margin of safety. An intentional stall can easily lead to an unintentional spin, and a spin recovery can eat up a lot of altitude.
- Finally, a word about the philosophy of stall recovery: Try to recover with minimum loss of altitude. Imagine that you were flying at 100 feet AGL and then did something stupid that led to a stall. The idea is to recover from the stall and climb back to a safe altitude, without ever losing more than 100 feet. Therefore the emphasis is on *recognition* and *recovery*: prompt recognition that the stall has occurred, and proper technique during the recovery.

There are many variations on the stall maneuver. You can stall the airplane with or without flaps extended, with or without power, during straight or turning flight, while pulling one or multiple *G*s, and during level, climbing, or descending flight.

To keep the discussion simple, let's first go through one specific scenario, and discuss the possible variations later.

Scenario #1: Start out in level flight at a typical traffic-pattern speed, in the landing configuration (full flaps extended,⁸ landing gear extended, carb heat on, et cetera). Then reduce the power to idle. As the airplane decelerates, pull back on the yoke at a steady rate, cashing in airspeed to pay for drag, maintaining altitude. Maintain constant heading. Maintain coordination. When the airspeed gets low enough, you may observe a sudden, distinct stall. The nose will drop, even though you are pulling back on the yoke. Obviously it is time to begin your stall recovery, as discussed below.

16.17.2 Provoking a Distinct Stall

However, it is quite possible you will not always observe a sudden, distinct stall. In particular, if your airplane is loaded so that its center of mass is right at the forward edge of the weight and balance envelope, you may be unable to deflect the elevator enough to cause a stall using the procedure described above.⁹ At this point you are at a very low airspeed, unable to stall the airplane, and unable maintain altitude by pulling back on the yoke. At this point you should declare an end to the attempted stall and begin your stall recovery

procedure. The ability to recognize the low-speed limit of performance in this situation is valuable, and should be practiced, but you should practice full-blown stalls also.

The most elegant way to improve your chances of observing a full-blown stall is to move the center of mass farther aft, using ballast. As described in [section 6.1.9](#), 100 pounds of water stowed securely in the back of the airplane¹⁰ should make it a whole lot easier to raise the nose.

Another trick that might increase your control authority is to use a little bit of engine power, perhaps 1500 RPM. On many airplanes the propwash flowing over the elevator increases the control authority by just enough to permit a quite distinct stall. On other airplanes (including those with high T-tails, and others) this trick doesn't work at all — the propwash over the wings lowers the stalling speed more than the propwash over the tail improves the control authority.

A third way to provoke a distinct stall is to zoom a little bit. That is, you maintain constant altitude while you decelerate *most* of the way. Keep track of how far back you have pulled back on the yoke. When you have used up most of the available backward motion, use the last inch or so to pull back faster than would be needed to maintain 100% level flight. The airplane will rotate to a more nose-high attitude, climb a few feet, then stall.

[16.17.3](#) Stall Recovery

Stall recovery, especially for poorly-trained pilots, poses psychological problems. In particular, if you are laboring under the dangerous misconception that the yoke is the up/down control, your instincts will be all wrong: the nose is dropping and the airplane is losing altitude, so you will be tempted to pull back on the yoke. This makes a bad situation much worse.

The correct way to think about the stall is to realize that the shortage of airspeed is your biggest problem. You need to push on the yoke and dive to regain airspeed.

In addition to the airspeed problem, you also have an energy problem. Therefore, while you are pushing on the yoke with one hand, you should be pushing on the throttle with the other hand.

As a further step to improve the energy situation, remove unnecessary drag. On most airplanes with N notches of flaps, the first several notches are somewhat helpful, because they allow you to fly slowly without stalling. The N th notch, however, typically doesn't contribute much to lowering the stall speed, and just adds a lot of drag. This would be useful if you were trying to descend, but since we are trying to climb at the moment, you should retract the N th notch of flaps as early as possible during the stall recovery. If the maneuver began with less than full flaps extended, leave the flaps alone, dive to regain airspeed, and then gradually retract the flaps.

While all this is going on, you should use the rudder and ailerons to keep the wings level and maintain a more-or-less constant heading.

You don't need to dive very far to regain a reasonable flying speed. According to the law of the roller coaster (as discussed in [section 1.2.1](#)), if you start out at 45 knots and dive 45 feet, you will wind up at 55 knots. If you start out at 50 knots and dive 80 feet, you will wind up at 65 knots.¹¹

At the bottom of the dive, perform a nice gentle pull-out. If you pull too rapidly, you put a big G load on the wings, which will cause them to stall at a speed that would otherwise have been just fine.

After you have leveled out at the bottom of the dive, accelerate horizontally to best-climb airspeed. Retract any remaining flaps as you accelerate. Then climb at V_Y to a safe altitude.

To summarize: the key elements of stall recovery include

- Dive to regain airspeed.
- Apply power.
- Reduce drag.
- Maintain wings level.
- Climb back to a safe altitude.

16.17.4 Power-On Stalls

A non-pilot might have thought that it would be hard to stall an airplane with the engine at full power, but in fact it is quite possible, and the accident statistics show that it happens fairly frequently. Therefore let's consider another scenario:

At a safe altitude in the practice area, set up for a power-off descent in the landing configuration. In particular, let this be a short-field approach, with the airplane trimmed to fly at the lowest practical airspeed. Then apply full power, as if for a go-around. In some airplanes (including the widely-used C-152, C-172, and C-182), and depending on where the center of mass is, this combination of trim, flaps, and power will cause the nose to pitch up quite dramatically. The airplane will climb very steeply and then stall. You don't need to pull back at all. Indeed, you may want to push a little bit so that the stall won't be too extreme.

In airplanes with better go-around characteristics (including a C-172 with the flaps retracted) you will need to work a little harder to perform a power-on stall. A possible — but not very stylish — way to perform this maneuver would be to start from cruising flight, add full power, and pull back until you get a stall. This is perhaps worth doing once, but it is not the recommended way of demonstrating a power-on stall, because results in climbing an unnecessarily long way. That is, it just isn't logical to apply full power while you are trying to decelerate. Therefore the conventional procedure is this: At a safe altitude, *reduce* power and decelerate in level flight to a speed a few knots above the stall. Then add power. (Use partial power the first time, and then use progressively more power as you learn how the airplane behaves.) Then gradually pull back some more.

As the airspeed bleeds off, you will need to apply more and more right¹² rudder to maintain coordination (i.e. to compensate for the helical propwash). Coordination is very important, because even a slight slip angle will cause one wing to stall before the other. This could easily result in a spin, and even if you don't get a full-blown spin, the sudden change in bank angle is pretty unpleasant.

Also, in this high-power low-airspeed situation, you will need to apply steady right aileron (to compensate for the rotational drag of the propeller). Note that (as discussed in [section 5.4.2](#)) the roll damping goes to zero at about the same point where the stall occurs, so you will need to intervene rather actively to keep the wings level. The standard advice applies: make sure you use the ailerons and rudder together. Because the airspeed is low, you will need a whole lot of rudder deflection to coordinate with a small amount of aileron deflection, and indeed right near the stall you can quite nicely control the bank angle using the rudder alone. Imagine that the left wing is about to stall. By stepping on the right rudder pedal, you can swing the nose to the right, causing the left wing to speed up and become unstalled. During this maneuver, you might want to lower the nose a tiny bit, so the right wing, which is swinging backwards, doesn't stall.

If you manage to maintain perfect coordination and perfectly level wings right up to the point of the power-on stall, you can still expect that the airplane will want to yaw and roll to the left just after the stall. There are several factors at work:

1. As discussed above, you are holding steady right aileron. This increases the effective angle of attack of the left wing, so it will stall first. The airplane will roll to the left.
2. The helical propwash causes the airflow to hit the left wing root area at a higher angle, and the right wing root area at a lower angle. This also causes the left wing to stall first. The airplane will roll to the left.
3. There is gyroscopic precession. That is, when the lift of the wings is suddenly reduced (while the lift at the tail is unchanged), it produces a torque — a nose-down pitching moment. In the absence of gyroscopic effects, this would cause the nose to drop, but as discussed in [section 19.9.2](#), the rule for gyroscopes is that the force is 90 degrees ahead of the motion. This means that the torque around the pitch axis will cause a motion around the yaw axis — in this case, a yaw to the left. See [figure 19.15](#). The resulting yaw will swing the left wing backwards, making it more stalled. The airplane will yaw to the left and roll to the left.

Of course, you can anticipate this, and apply additional right rudder as the nose drops. With a little experience, you can arrange that the wings stay level and the nose drops straight ahead.

The recovery from a power-on stall is basically the same: dive to regain airspeed, add power (if you were not already at full power), maintain wings level, reduce drag, and climb back to a safe altitude.

The practical test standard calls for performing power-on stalls with the flaps in the takeoff configuration and gear down (the takeoff configuration) or gear retracted (departure configuration) which simulates a stall

happening shortly after takeoff. It is well worth practicing other configurations, too — particularly the approach configuration, which simulates what might happen if you mishandle a go-around.

16.17.5 Accelerated Stalls

The stall occurs at a definite angle of attack. This is not quite the same as a definite airspeed, for reasons discussed in [section 2.12.5](#). As an example: if you are in a 45 degree bank, the load factor is 1.4, and the stalling speed will be 20% higher than it would be in ordinary one-G flight. Therefore, if you are relying on the airspeed indicator to warn you of an impending stall, you will be fooled.

Ironically, in some ways recovery from an accelerated stall is easier, because of the extra airspeed. You may not have to dive any significant distance. Otherwise, the procedure is the same: reduce the back pressure (and dive if necessary to regain airspeed), roll the wings level, add power, reduce drag, and climb back to a safe altitude.

It is the bank itself that makes a turning stall exciting. If you start out banked 45 degrees to the left, and then (due to a lapse in coordination, and/or any of the other factors discussed in [section 16.17.4](#)) the left wing drops another 45 degrees, you could wind up in a knife-edge attitude. If this happens, push forward on the yoke to the position that corresponds to zero angle of attack, so there is no load on the wings. Then just sit there for a second or two. The rudder will cause a rotation around the now-horizontal yaw axis, so that the airplane is soon pointing somewhat nose-down. After the airplane has dived a couple dozen feet, you will have enough airspeed that the ailerons are effective. Use the ailerons to roll the wings level, then pull out of the dive in the usual way.

Another thing that makes accelerated stalls a bit more challenging has to do with perception of the stall. Imagine an airplane where (due to a lack of nose-up elevator authority or whatever) the stall doesn't "break" suddenly. Then as you approach an ordinary, straight-ahead stall, you have a constant heading and everything looks fairly normal. You can devote a lot of attention to looking for subtle signs of the stall. Now contrast that with a stall during a 60-degree bank. The pitch axis is not horizontal, so any pitch change will move the nose mostly along the horizon, not perpendicular to the horizon. Also, no single outside reference stays in front of the nose for very long, so subtle pitch changes are quite a bit harder to notice. A student once complained ``I can't get this thing to stall''. I replied ``We're going down more than 2000 feet per minute. This is stalled enough for me''.

Finally, steep turns are not the only way that an accelerated stall can arise. A sharp pull-up with wings level will do it. In an aerobatic loop, for instance, you are pulling about 4 Gs at the bottom, so the stalling speed is about twice what it would be in ordinary unaccelerated flight. Especially since you might be rapidly approaching the ground at this point, there is a strong temptation to pull back, but be careful, because this would be a really inopportune time to stall. Make sure you have plenty of altitude and plenty of airspeed before attempting any high-G maneuvers.

16.17.6 Evil Zooms

As discussed in [section 12.11.8](#), it is fairly easy to get into a situation where you have a nose-high pitch attitude, very little airspeed, and very little altitude. In this situation, the usual stall-recognition and stall-recovery techniques will do you no good whatsoever. You need to recover *before* the airplane stalls, and you need to recover with zero loss of altitude.

Therefore it is a good idea to practice recovering from this situation. The procedure is:

1. Go up to a safe altitude.
2. Set up for a power-off glide in the landing configuration.
3. Gradually pull back on the yoke until you are a few knots above the stall speed.
4. Then pull back on the yoke quite a bit more. Observe that the airplane rotates to a very high nose-up attitude and begins to climb.
5. Before the airplane has climbed more than a few feet, and *before* it stalls, push the nose back down to the attitude that corresponds to level flight at a very low airspeed.
6. At the same time, apply full power.
7. Fly level until you regain airspeed, using the usual go-around procedures.

Practice this over and over, until you are confident that you can recover from a pitch excursion with zero loss of altitude.

1

... unless you are inside a cloud, in which case you hope everybody in that cloud is on an IFR clearance so that ATC can provide separation.

2

You may have seen some books that refer to the ``four fundamentals". Here's how they get from three to four:

- They list straight-and-level as a separate item, whereas I consider it the natural consequence of zero change in altitude, zero change in airspeed, and zero turn.
- They treat climbs as different from descents.
- They treat left turns the same as right turns.
- They entirely disregard acceleration and deceleration, whereas I consider airspeed control to be quite fundamental.

3

If you can't find a suitable scratch or bug corpse on the windshield, it may be instructive to make a mark, as discussed in [section 11.5.2](#).

4

The other possibility is a windshear from tailwind to headwind, which is the *opposite* of what you normally encounter. This will give you excess energy, which could be a problem, especially if the field is short.

5

These are flat plates that pop up from the top of each wing. The air hits them broadside. They approximately double the airplane's coefficient of parasite drag.

6

...unless you spend a lot of time on the back side of the power curve, which is usually not practical.

7

The physics works like this: Your kinetic energy relative to the new air is greater than your kinetic energy relative to the old air. Your airspeed relative to the ground has not changed, or may even have decreased slightly, but that is irrelevant. The airplane doesn't care about the ground. The local air is the only thing that matters.

8

It is a little hard to explain why, in everyday flying, you would be flying level with full flaps extended, but don't worry about that. This maneuver (a) is a good training exercise, and (b) is an important part of the FAA practical test.

9

There is, after all, a physical limit to the amount of force any finite-sized elevator can produce, and this typically explains why the forward edge of the envelope is where it is.

10

This works fine in a four-seat aircraft with two people aboard, or a two-seater with one person aboard, but it may not be possible in a two-seater with two people aboard, because of limits on the total weight.

11

You can practice most elements of the stall-recovery maneuver without actually stalling the plane. That is, starting from level flight a few knots above the stalling speed, push the nose over, dive 50 feet or so to gain airspeed, and then level off. Don't forget to apply power, reduce drag, and maintain wings level. This sort of practice often helps students overcome their fear of stalls, by building up their confidence in their recovery procedures.

12

Assuming a standard American engine that rotates clockwise as seen from behind.

17 Multi-Engine Flying

Q: In an underpowered twin, what is the role of the second engine?

A: It doubles your chance of engine failure, and it will fly you to the scene of the accident.

17.1 Normal Operations

In normal conditions, operating a twin is not very different from operating a heavy, high-powered, complex single. Even so, the weight, power, and complexity may be more than you are accustomed to. In particular:

- A typical twin has a much higher stalling speed than a typical single.¹ Other critical speeds are increased in proportion. One consequence is that you will notice that normal turns in the traffic pattern require considerably more room.
- The rate of climb will be considerably higher, so you will reach pattern altitude sooner. The pitch attitude during climb will be higher.
- A higher-performance airplane will typically operate over a wider range of speeds. For example: in a Cessna 172, you might climb at 78 knots and cruise at 105 knots (a ratio of 1.3), whereas in a Seneca II you might climb at 90 knots and cruise at 170 knots (a ratio of 1.9). When you level off following a climb, the Seneca will take much longer to accelerate to its cruising speed. After the level-off, you let it accelerate for a few seconds, and then trim it — but then it will accelerate some more and you will have to trim it some more. You should plan on prolonged acceleration and repeated trimming.

Normal multi-engine takeoff procedures are discussed in [section 13.6](#). Most of the rest of this chapter is devoted to the procedures and principles of engine-out flying.

17.2 Engine Out Scenarios

This section discusses some of the things you might observe when an engine fails, and what you can do in response.

17.2.1 Takeoff

If an engine fails during the takeoff roll, you have a decision to make: In some cases you must close the throttles and try to stop on the remaining runway, whereas in other cases you must try to fly away using the remaining engine. You won't have a lot of time to think about it, so you want to be 100% prepared to do the right thing instantly. This is discussed in [section 17.3.1](#) and [section 17.3.2](#).

17.2.2 Climb

For our next scenario, suppose you are at a reasonable altitude, climbing with full power on both engines. Then one engine fails. One of the first things that you will notice is that the single-engine rate of climb is not half of the two-engine rate of climb. No, indeed! The reason is simple: as shown in [figure 17.1](#), when an engine is shut down, you are not splitting the difference between two-engine performance and level flight; you are splitting the difference between two-engine performance and a zero-power descent.




Figure 17.1: Two-, One-, and Zero-Engine Performance

The power curves in the figure are roughly representative of a Piper Apache, a well-known light twin trainer. Point A corresponds to the two-engine best rate of climb, 1150 fpm at 86 knots. Point B corresponds to the single-engine best rate of climb, 160 fpm at 82 knots. (These numbers apply to a fully-loaded aircraft at sea level in the clean configuration.) We see that the single-engine rate of climb is less than 15% of the two-engine rate of climb.

At density altitudes above 5000 feet the Apache cannot climb at all on one engine. Also, if the engines, propellers, and paint job are not quite factory-new, the performance will be even less than these book values suggest.

You must not allow yourself to think that just because airliners can climb with an engine out, your favorite light twin can climb with an engine out.

It is legal to operate a light twin with anemic or nonexistent single-engine climb performance. In such cases, engine failure at low altitude is perhaps the most critical situation that arises in general aviation with any appreciable frequency. Like a single-engine aircraft with partial power failure, you need to make a forced landing. The problem with the twin is that (because of the asymmetric thrust) if you mishandle the situation, your chance of getting into a spin is much higher than it would be in a single.

On the other hand, even if you are not climbing, you are probably not descending very fast. You can treat it as another "noisy glider" situation as mentioned in [section 15.1](#). If you start out several thousand feet above the ground, you can probably travel dozens of miles while gradually descending. Look around and find a nice place for a forced landing.

[17.2.3](#) Slip String

Generally, the best way to fly any airplane is to keep the airflow aligned with the fuselage. Alas, in a multi-engine airplane with asymmetric thrust, it can be rather tricky to perceive just what the slip angle is. The most direct way to get information about this angle is to use a *slip string*.

To create a slip string, tape a piece of yarn to the nose of the airplane, in front of the center of the window where you can see it. Leave a foot or two of yarn dangling free, so it will align itself with the airflow.

In typical aircraft, the fuselage accelerates the sideways flow, more than it accelerates the fore-to-aft flow. This increases the sensitivity of the string. That is, the angle of the string is larger than the actual slip angle.


Even in ordinary single-engine airplane, slip strings work surprisingly well. The straight-back component of propwash decreases the sensitivity somewhat, and the helical component of propwash biases the string slightly to one side.

The slip string is commonly referred to as a "yaw string", even though it measures the slip angle, not the yaw angle (i.e. heading) or yaw rate. As discussed in [section 19.6.3](#), the slip angle measures the angle between the fuselage and the relative wind, whereas yaw is defined relative to some fixed spatial direction. Heading and heading change (i.e. yaw rate) are easy to perceive by looking out the window, while it is not easy to perceive slip angle except by reference to a slip string. Heading can also be perceived using the directional gyro (heading indicator), and yaw rate can be perceived using the rate-of-turn gyro or by observing the rate of motion of the directional gyro.

The inclinometer ball is often referred to as the slip/skid instrument, but that is a misnomer. It measures inclination. As we shall see in [section 17.2.5](#), it is quite possible for the airplane to be inclined but not slipping. Again: there is no good way to determine the slip angle without a slip string.


[17.2.4](#) Coordination

For our next scenario, imagine that you are in level flight at cruise airspeed at a comfortable altitude. Let's also suppose that your airplane has a slip string installed. Then (surprise!) an engine fails. To simplify the discussion, let's suppose the right-hand engine² is the one that quits. You will immediately notice that the airplane will develop a slip angle. In this case, the airplane will yaw to the right, shown in [figure 17.2](#). This is because the drag of the airplane is more or less symmetrically disposed, while the thrust is now quite asymmetric. As always, two forces with a lever arm between them create a pure torque.




[Figure 17.2](#): Engine Out — Torque due to Asymmetric Thrust

This torque will produce an initial heading change. This will be a pure yaw; that is, it will change the direction the airplane is *pointing* without any immediate change in the direction the airplane is *going*. Then, after a short time (a second or so), the torques will come back into equilibrium, because of the airplane's natural yaw-axis stability (as discussed in [section 8.2](#)). That is, the uncoordinated airflow hitting the rudder will create a torque that opposes the asymmetric thrust. If you managed to keep the wings level, you will be in a boat turn to the right. The slip string will off-center to the right, indicating an asymmetric airflow, and indicating that you need to apply some left rudder.



[Figure 17.3](#): Engine Out — No Pilot Action

At this point, you could either (i) sit there and be a spectator, as shown in [figure 17.3](#), or (ii) press on the left rudder pedal to center the slip string, as shown in [figure 17.4](#). From the point of view of directional control, your choice doesn't matter very much. That is, in case (i) the airflow strikes the whole airplane (rudder included) at a nonzero angle, while in case (ii) the airflow strikes the airplane at a zero angle, with the rudder deflected. In either case, the amount of tail force produced is approximately the same. The most important difference is that the airplane will climb better in case (ii), because the airflow will be aligned with the fuselage.



[Figure 17.4](#): Engine Out — Correct Compensation

Let's see what happens after the yaw-axis torques have returned to equilibrium. Let's assume for now that you are keeping the wings level. In case (i), the airplane will make a steady turn toward the dead engine. This is obviously a boat turn, due to the uncoordinated airflow striking the fuselage, as discussed in [section 8.10](#). This will be a genuine CM-turn (as defined in [section 8.8](#)), changing the direction of motion of the center of mass; the heading will follow the CM-turn in order to maintain a constant slip angle.

What is perhaps more surprising is that even in case (ii), if you keep the wings level the plane will make a CM-turn toward the dead engine (toward the right in this case). It will not turn as rapidly as in case (i), but it will turn nonetheless. The reason is that the rudder force, in addition to creating a torque, is creating an unbalanced force. This force is changing the direction of motion of the center of mass. A possible (but non-optimal) way to stop this turn would be to apply even more pressure on the left rudder pedal, which would create a wings-level non-turning slip, as will be discussed below in conjunction with [figure 17.5](#). For now, though, let's consider the correct strategy, which is to keep the slip string centered and apply a *bank* (to the left, in this case) to stop the turn. This is shown in [figure 17.4](#). This uses a leftward component of lift³ paired with the rightward rudder force. Once again we have a pair of forces with a lever arm between them, i.e. a pure torque. This lift/rudder pair cancels the thrust/drag pair discussed above.

To reiterate: when engine trouble develops, the first result of the asymmetric thrust is to make the airplane yaw toward the dead engine. The airplane changes its *heading* immediately (whereas only later does it gradually change the direction it is *going*). That is, a slip angle develops immediately. If you don't deflect the rudder, the slip angle will grow until the uncoordinated airflow striking the rudder develops enough torque to stop further yawing. This is the basic yaw-axis stability mechanism as discussed in [section 8.2](#). The result is that the airplane does *not* spin around and around like a Frisbee — it just develops a few degrees of slip angle and then stabilizes.

[17.2.5](#) Perception and Initial Response

If an engine fails suddenly, the initial yaw toward the dead engine is usually quite noticeable. On the other hand, if it fails gradually, the initial yaw (heading change) may be harder to perceive than you might have

guessed. It is especially hard to perceive if it occurs during a turn — the turn just proceeds a little faster or slower than normal. The subsequent boat turn may not be super-easy to perceive, either.⁴ There are three basic ways to perceive and deal with the slip and yaw.

1. If you happen to have a slip string installed, the procedure is simple. If the string is deflected to one side, step on the rudder pedal on the opposite side, until the string is centered. The mnemonic is: ``Step away from the string".
2. A less-elegant, less-accurate technique is to use the inclinometer ball. If the ball is deflected to one side, step on the rudder pedal on the same side. The mnemonic is: ``Step on the ball". Center the ball, then (to establish zero slip, as discussed below) relax the pedal force to let the ball go off center about one-third to one-half of its width.
3. A third technique (which is really the most common technique) is to roll the wings level and then apply the rudder as needed to stop the boat turn. The advantage of this procedure is that it can be done without reference to instruments. The main disadvantage is that it doesn't help you regain or retain control in a turn.⁵ (There are rare situations where even though an engine has failed you might *want* to be turning.) Then, once you've got the wings level and the turn stopped, you should establish the optimal zero-slip condition, by raising the dead engine a few degrees and releasing some of the rudder pressure.⁶

At this point, you will find yourself maintaining a rudder deflection and a bank angle, both toward the side with the working engine. Use the rudder deflection (not the bank!) to identify which engine has failed. The mnemonic is: *working foot, working engine; dead foot, dead engine*. Specifically, if the right foot is *not* being used to deflect the rudder, bend your right knee. Raise that knee an inch or two, pat it a couple of times, and say ``right engine has failed". (More on this later.)

To maintain zero slip, you will need to bank the plane very slightly toward the working engine. The mnemonic is: *raise the dead*. This also implies that the inclinometer ball will be slightly displaced toward the working engine. This correct procedure ([figure 17.4](#)) requires slightly *more* aileron and slightly *less* rudder than you would need for wings-level, ball-centered, non-turning, uncoordinated flight ([figure 17.5](#)).

Having the slip string centered but the inclinometer ball not centered may seem a bit counterintuitive, so let's examine the aerodynamics of the situation a little more closely.

The asymmetric thrust produces a yaw-axis torque, which cannot remain unopposed. The rudder is part of the solution, but remember that while the rudder is producing the desired torque it is also producing a force. We need the forces to be in balance, as well as the torques.

Suppose you try to maintain zero bank instead of raising the dead. Initially the airplane is not in equilibrium, because the rudder is producing an unbalanced force toward the dead-engine side. There are then two possibilities:

1. Suppose the sideways force remains unbalanced. This will cause the airplane to turn. This will be a wings-level, coordinated turn. I call this a pseudo boat turn. Unlike an ordinary boat turn, the airflow is coordinated along the fuselage, but unlike a regular turn, the horizontal force is not coming from the wings.
2. Suppose you push a little harder on the rudder pedal, establishing a slip toward the dead engine. The ball is centered, the wings are level, and the rate of turn is zero. (Any two of those things implies the third, regardless of engine status.) The forces are in balance because there is enough uncoordinated airflow over the fuselage to create a sideways force that balances the rudder force. The slip string is offset to the left, indicating that you are applying too much left rudder. In this situation, as shown in [figure 17.5](#), you are using the fuselage as an airfoil. The problem is that the fuselage has a really poor lift-to-drag ratio. The sideways fuselage lift force is accompanied by a huge drag force, which steals energy from you. It would be much better to use the wings, as previously discussed in conjunction with [figure 17.4](#).




Figure 17.5: Engine Out — Wings-Level Nonturning Slip

The proper technique is counterintuitive, because in any normal situation, proper coordination implies that the rate of turn will be proportional to the amount of bank. (See [section 11.5.2](#), for example.) With an engine out, though, proper coordination requires a slight bank when you are not turning.

The amount of bank for a typical airplane can be estimated using the following argument: The lift-to-drag ratio of the airplane is roughly ten-to-one. In level flight the thrust must therefore be one tenth of the lift. The lever arm between the wings and the rudder is typically about three times the lever arm between the thrust and the drag. Since the torques must cancel, the rudder force (and the horizontal component of lift) must be one third of the thrust, and therefore one thirtieth of the lift. We conclude that the horizontal component of the lift is one thirtieth of the total lift. One thirtieth of a radian is two degrees — not exactly a huge bank.

To find out exactly how much bank you need to maintain coordinated airflow over the fuselage, it helps to use a slip string. At a safe altitude, set up for single-engine flight at V_{YSE} . Apply enough rudder pressure so that the slip string indicates zero slip angle. Bank as required to maintain nonturning flight. Experiment with slightly greater or lesser rudder pressure, to see what produces the best climb performance.

You will discover that in optimal single-engine flight, the inclinometer ball is *not* centered, but the slip string is centered. The airplane is inclined, but it has zero slip angle. Make a note of how much inclination is indicated by the inclinometer ball; typically it will be off-center by one-half or one-third of its diameter. You can use this information to set up a good approximation of engine-out coordinated flight during subsequent flights when you don't have a slip string installed.

The inclinometer ball measures the inclination of the wings relative to the E-down⁷ direction. The inclinometer is sometimes referred to as a slip/skid ball, but that is a misnomer⁸ because the slip string (as discussed in [section 17.2.3](#)) provides your only direct information about the slip angle.

Achieving zero slip is the key to optimal climb performance.⁹ The idea is to have the airflow aligned with the fuselage. Centering the inclinometer ball is not what determines performance. Practice with the slip string until you learn how much inclination is required for a given amount of asymmetric thrust.

17.2.6 Yaw Control at Reduced Speeds

So far, we have discussed engine-out climb rate ([section 17.2.2](#)) and discussed the value of maintaining coordinated flight ([section 17.2.4](#)). We now begin a discussion of airspeed. As you might imagine, this is rather important.

So, let's consider what happens when *airspeed variations* are added to the previous scenario. Let's assume you start out at cruise airspeed and gradually decelerate. Again, to simplify the discussion, let's assume the right engine has failed.

The amount of asymmetric thrust does not depend on airspeed; it depends only on the power output of the engine. In contrast, the amount of force the rudder produces depends on the airspeed squared, and on the rudder's angle of attack. Therefore as you decelerate you will need progressively more rudder deflection in order to maintain zero slip. If you do it right, the sideways force developed by the rudder will remain unchanged, and the bank angle will remain unchanged (for now).

At some point you will run out of rudder deflection. The pedal (or the rudder itself) will hit the stops. You will be unable to maintain zero slip.

Now, suppose you continue to decelerate beyond this point. As a slip develops, the airflow hits the tail and rudder at an angle. This gives the tail/rudder an angle of attack over and above whatever angle of attack you

created by deflecting the rudder. You are using the slip angle as a substitute for additional rudder deflection. Up to a point, this higher angle of attack allows the tail/rudder to produce a higher coefficient of sideways lift, allowing it to produce the required force in spite of the lower airspeed.

In addition to the air hitting the rudder, you now have the uncoordinated airflow hitting the fuselage. You are relying on the rudder to produce at least 100% of the torque needed to oppose the asymmetric thrust. The air hitting the fuselage makes a small unhelpful contribution to the torque budget, and (more noticeably) contributes to the sideways force budget, producing an undesirable boat turn. This boat turn is in addition to the pseudo boat turn that the rudder is producing, so you will need to increase the bank angle to maintain nonturning flight.

Obviously there is a limit to this process. If you keep increasing the rudder's angle of attack, at some point the rudder will stall. Remember, the amount of asymmetric thrust does not depend on airspeed, whereas the *absolute maximum* amount of force the rudder can produce depends on the airspeed squared. Therefore, for any nonzero amount of asymmetric thrust, there must be some airspeed below which the rudder cannot develop enough torque. At that point there will be an uncontrollable yaw toward the dead engine. The airplane will spin like a Frisbee.

You might think you could improve the situation by releasing the rudder pedal, thinking this would reduce the rudder angle of attack. Alas, it won't work. It will just cause the airplane to establish a greater slip angle. Remember the rudder needs to produce a certain amount of force to oppose the asymmetric thrust, and the airplane's natural yaw-axis stability will adjust the tail/rudder's angle of attack, trying to create the necessary force.¹⁰

If the rudder stalls, it will be about as unpleasant as anything you can imagine. There will be a sudden uncontrollable yawing motion. Because of the yawing motion, the wingtip on the side with the good engine will have a higher airspeed than the wingtip on the other side. Because of the difference in airspeed (plus the difference in propwash patterns) the good-side wing will produce much more lift, so you will get an uncontrollable roll. As the inside wing drops, it will probably stall (since you were already at a low airspeed). You are now in a spin. There is no guarantee that it will be possible to recover from the spin; multi-engine airplane certification regulations do not require spin recoveries.

On some planes (such as an Apache, a common trainer) low-speed engine-out performance is limited by the rudder, as described above. On some other planes (such as a Seneca, another common trainer) you don't need to worry about the rudder because the wings will stall first.¹¹ This is not much of an improvement, because a stall with asymmetric power is also rather likely to result in a spin.

To prevent such nasty things from happening, you need to maintain a safe airspeed. The manufacturer gives you some guidance in this regard, as is discussed in the following section.

17.2.7 Minimum Control Speed — Definitions

The symbol V_{MC} denotes "minimum control airspeed". There are at least four different definitions of this term, including:

I)

FAR 23 (the certification requirements for typical general-aviation¹² airplanes) gives a very specific definition of V_{MC} , namely:

FAR 23.149 Minimum control speed.

(a) VMC is the calibrated airspeed at which, when the critical engine is suddenly made inoperative, it is possible to maintain control of the airplane with that engine still inoperative, and thereafter maintain straight flight at the same speed with an angle of bank of not more than 5 degrees. The method used to simulate critical engine failure must represent the most critical mode of powerplant failure expected in service with respect to controllability.

(b) VMC for takeoff must not exceed $1.2 V_{S1}$, where V_{S1} is determined at the maximum takeoff weight. VMC must be determined with the most unfavorable weight and center of gravity position and with the airplane airborne

and the ground effect negligible, for the takeoff configuration(s) with--

- (1) Maximum available takeoff power initially on each engine;
- (2) The airplane trimmed for takeoff;
- (3) Flaps in the takeoff position(s);
- (4) Landing gear retracted; and
- (5) All propeller controls in the recommended takeoff position throughout.

[...]

II)

FAR 1 (the "definitions" section) defines V_{MC} as "minimum control speed with the critical engine inoperative". It does not specify any restrictions as to weight, configuration, altitude, et cetera.

III)

The FAA Practical Test Standards for the multi-engine rating call for demonstrating " V_{MC} " in a particular way that emphasizes losing yaw control *without* stalling the wing or rudder, even though (as discussed below) for many airplanes V_{MC} (under definition I or II) is limited by wing stall and/or rudder stall.

IV)

In common parlance, pilots apply the term V_{MC} to the airspeed where the airplane (multi-engine or otherwise) becomes uncontrollable, no matter what the reason, no matter what the configuration, and no matter whether any engine is inoperative.

Note that none of these definitions require that the airplane exhibit a positive rate of climb at V_{MC} .¹⁴ Also note that during a V_{MC} demonstration, the pilot is not required to optimize the climb rate or to maintain zero slip — although zero slip may be an advantage if it can be achieved.

The V_{MC} number in the Pilot's Operating Handbook is determined according to the FAR 23.149 definition. This airspeed is marked with a red radial line on the airspeed indicator, and is sometimes called the *FAR 23.149 redline* airspeed.¹⁵

There are various ways to lose control; whichever happens first determines where the manufacturer sets the redline:

a)

In some airplanes, under some conditions, you can maintain control, even with an engine out, right down to the point where wing stalls. This is discussed below in conjunction with [figure 17.6](#). A stall with asymmetric thrust could be rather sudden and nasty.

b)

In others (with a smaller rudder, larger wing, and/or higher-thrust engines), there will be conditions under which the rudder will stall before the wings do, as is discussed below in conjunction with [figure 17.8](#). A rudder stall could be very sudden and very nasty.

c)

In yet others (larger rudder: area but a shorter tail-boom, so that the rudder is closer to the wings), there will be situations where neither the wing nor the rudder is stalled, but the boat-turn forces are so large that it requires more than 5 degrees of bank to counteract them and maintain nonturning flight. The airplane would be perfectly controllable if the bank were not limited to 5 degrees. Since a bank of 15 or 20 degrees is not particularly dangerous, the 5 degree limitation must be considered arbitrary. If your airplane, at a given weight and altitude, does run up against this limitation, the resulting "loss of control" is neither sudden nor nasty. The airplane will just make a gentle boat turn toward the dead engine, as is discussed below in conjunction with [figure 17.7](#).

Possibility (c) is in some ways attractive, but you have no guarantee that this is what will happen. Rudder stall depends on slip angle, so you may be wondering why FAR 23.149 should mention a *bank* angle as opposed to a *slip* angle. Bank does not cause slip.¹⁶ If you want to establish any connection between bank and slip, you must consider:

1. bank angle (i.e. the angle between wings and horizon)
2. slip angle (as indicated by the slip string)
3. rate of turn
4. asymmetric thrust

If any three of these are zero, the fourth is guaranteed to be zero. More generally, other things being roughly

equal, given any three of these you can estimate the fourth. The problem is that other things are generally *not* equal — depending on weight, airspeed, airplane design, et cetera, five degrees of bank could correspond to a large slip angle or perhaps no slip angle at all. So this regulation is not 100% logical.

Some people seem to assign a near-religious significance to the "5 degree bank" mentioned in FAR 23.149. However, the real significance is quite limited:

- This regulation applies to the manufacturer during certain tests. It does not apply to you in your ordinary flying. If you have a real engine failure, you are limited only by the laws of aerodynamics.
- This regulation does not even apply to you during the checkride for a multi-engine rating. In particular, the FAA Practical Test Standard says you should bank for "best performance and controllability". Alas, that's inconsistent; best controllability requires a lot more bank than best performance, and the PTS doesn't tell you how to make the tradeoff.
- Five degrees is no guarantee of optimum performance. The optimal bank could be five degrees, or more, or less (usually less).
- There is nothing in FAR 23.149 or anywhere else that guarantees the airplane is well behaved at the "official" 5-degree bank angle. The maximum bank you can safely use in nonturning engine-out flight could be five degrees, or more,¹⁷ or less.¹⁸ In particular, there is no guarantee that by limiting yourself to 5 degrees you will always get aerodynamic warning (in the form of a nice, gentle boat turn) before you get a nasty rudder stall or wing stall. If you want to demonstrate a gentle warning, you might need to limit yourself to much less than 5 degrees.

One thing we learn from this is that you should not use bank angle or anything else as a substitute for proper airspeed control.

For that matter, airspeed control requires a little thought, too. Perhaps because FAR 23.149 uses words like "most critical" and "most unfavorable", people commonly assume that it is always possible to control the airplane at redline airspeed, no matter what. This assumption is wrong — dangerously wrong — in many airplanes. For example, there are some airplanes where the certified takeoff configuration¹⁹ calls for the flaps to be extended, and the FAR 23.149 redline is essentially equal to the stalling speed in the takeoff configuration. Then if you operate with the flaps retracted, you will lose control of the airplane at an airspeed well above redline.²⁰

Specific procedures for dealing with engine failure are discussed below, in [section 17.3](#).

[17.2.8](#) Effect of Altitude, Weight, etc.

FAR 23 tells us that the airplane, when operated under a particular set of circumstances, can maintain directional control at redline airspeed. The question is, what happens under other circumstances?

Let's discuss an example. Consider a non-turbocharged airplane for which the handbook calls for flaps retracted during takeoff. Then, under standard conditions (takeoff configuration, maximum weight, etc.), the situation is shown in [figure 17.6](#). The single-engine stall speed for the example airplane is shown by a black vertical line in the middle of the figure. The FAR 23 redline is shown as a bright red tick mark on the airspeed axis. The manufacturer had to set it a knot or two above the stalling speed, since that is what limits the low-speed handling for this airplane in this configuration.

Also, in this figure, the magenta curve shows the airspeed below which the rudder cannot develop enough force to oppose the asymmetric thrust. Thirdly, the dotted cyan curve shows the airspeed below which the boat turn forces are so large that it would require more than 5 degrees of bank to maintain nonturning flight.

Since the example airplane is not turbocharged, as altitude increases there is less thrust available on the good engine. The required rudder force declines accordingly. This is why the magenta and cyan curves trend to the left as they go up. Note that in this configuration, for this airplane, rudder performance is not a limitation — the wing stall is the only relevant limitation.




Figure 17.6: Speed & Altitude affect Directional Control (basic)

Now, suppose that several things change:

- You fly at reduced weight, about half of the legal maximum.
- You extend the flaps.²¹
- You limit yourself to 3 (not 5) degrees of bank.
- You limit yourself to less than full rudder deflection.

Let's consider what happens under these new circumstances. The reduced weight will lower the stalling speed. Similarly, extending the flaps lowers the stalling speed. This is indicated by the black line, which moves to the left as we go from [figure 17.6](#) to [figure 17.7](#).

The amount of torque developed by the engine depends on altitude in the same way as before, and is unaffected by the weight, flaps, and other variations. The amount of force the rudder can produce is also unaffected. Therefore the magenta curve is the same in the two figures.

At the reduced weight, less lift is needed for supporting the weight of the airplane. As always, the horizontal component of lift, at any particular bank angle, is proportional to the weight of the airplane. Therefore, at any particular bank angle, you have less ability to oppose a boat turn. This is one reason why the cyan dotted curve moves to the right as we go from [figure 17.6](#) to [figure 17.7](#).





Figure 17.7: More Flaps, Less Weight, etc.

Limiting yourself to less than full rudder deflection does *not* reduce the amount of torque that must be produced in order to oppose the asymmetric thrust; it just means that the airplane will establish a slip to create the necessary force.²² In this slipping condition, the fuselage produces a boat turn on top of whatever pseudo boat turn the rudder is producing, so you will need more bank to oppose the turn, and you will run up against the bank limitation sooner. This is the second reason that the dotted cyan curve (the bank limit) moves to the right.


And of course, if you limit yourself to a smaller bank, you will run up against the bank limit sooner. This is the third reason that the cyan curve moves to the right as we go from [figure 17.6](#) to [figure 17.7](#). Conversely, if you allow yourself a large bank (15 or 20 degrees) you can push the dotted cyan curve very far to the left, as indicated in [figure 17.8](#).



[Figure 17.8](#): More Bank

Finally, let's consider what happens in different airplanes. For example #2, let's take an airplane where the wing has a very low stalling speed. For such a plane, [figure 17.6](#) never applies; [figure 17.7](#) (or [figure 17.8](#), depending on bank angle) applies even at max weight with the flaps retracted.

For example #3, let's go to the opposite extreme and consider an airplane that has somewhat smaller wings. To compensate, the manufacturer specifies that flaps are to be extended in the certified takeoff configuration. The result is that the certified performance of the new plane is identical to the performance of example airplane #1, as shown in [figure 17.6](#). The interesting wrinkle is this: if you fly the new airplane with flaps *retracted*, the performance is as shown in [figure 17.9](#). Note the higher wing stall speed. The airplane will become uncontrollable at an airspeed well above the FAR 23.149 redline.



[Figure 17.9](#): Certified Flaps Not Used

Let's summarize this information into a form that is perhaps more directly useful when you are actually in the cockpit, planning or performing engine-out maneuvers.

- Maintain a safe airspeed. This speed should be above the wing's stalling speed *in the current configuration* and above the FAR 23.149 redline, whichever is higher. Leave yourself a reasonable margin of safety.
- The best procedure involves establishing zero slip (or minimum slip, if full rudder deflection isn't enough to establish zero slip), and banking to stop the turn. The amount of bank increases as the weight decreases; use whatever bank angle does the job. Remember, though, that maintaining a safe airspeed is more important than getting exactly the right slip angle or bank angle.

- In your multi-engine training, you were probably given the chance to demonstrate "loss of directional control" under conditions where the "loss" resulted in a gentle boat turn toward the dead engine. You absolutely must not assume that the airplane will always behave this way. In other circumstances, you might get a sudden rudder stall or wing stall, either of which could result in a spin.
- If you want to demonstrate the gentle boat turn, you can *arrange* that it occurs before any of the nastier alternatives, as suggested in [figure 17.7](#). You just have to put sufficiently strict limits on the bank and rudder deflection. Reduced weight helps, too. Turbocharging makes it easier to perform the demonstration at a safe altitude.
- If you go exploring speeds below redline, things get dicey. If you allow yourself unlimited [23](#) bank, there is no doubt that you can maintain directional control right down to the point where the wing stalls and/or the rudder stalls. You can get a good estimate [24](#) of the wing's stalling speed, but I can't think of a safe way for you to find out whether or not the rudder will stall before the wings. [25](#) Please do *not* try to find this out experimentally!

For more information on engine-out procedures, see [section 17.3](#).

[17.2.9](#) Effect of Center of Mass

We know that we have to pay careful attention to the location of the airplane's center of mass, since it has a big effect on the angle of attack stability; see for example [section 6.1.1](#).

This leads us to wonder what effect center-of-mass position has on V_{MC} . There are two possible answers:

1. CM location has no effect whatsoever if you use the unwise wings-level technique depicted in [figure 17.5](#).
2. CM location does matter if you use the recommended procedure depicted in [figure 17.4](#). As the CM (or, more precisely, the center of lift) moves aft, V_{MC} increases.

In both cases, you need to create a torque to oppose the asymmetric thrust. You create it using a pair of forces with a lever arm between them. One force comes from the rudder.

In case (1), the rudder force is paired with a horizontal force due to air hitting the side of the fuselage. This fuselage horizontal lift depends on the shape of the airplane, but does not at all depend on the CM location.

There is a deep theorem of physics that says that for any two axes parallel to each other, the torque around one is the same as the torque around the other (provided there are no overall unbalanced forces on the system). In the zero-bank case, it means that V_{MC} can't depend on center of mass location (unless the airplane is actually turning, i.e. being accelerated sideways).

To understand the basis of this theorem, refer again to [figure 17.5](#). Let's pick two pivot points *A* and *B* somewhere along the rudder/wing lever arm, as shown in the figure. (You can, if you wish, imagine them to be two possible locations of the center of mass; the CM is no better or worse than any other pivot point.)

When we calculate the total torque around each pivot point:

- The lever arm from *A* to the rudder is long, but the lever arm from *A* to the other horizontal force is short.
- The lever arm from *B* to the rudder is short, but the lever arm from *B* to the other horizontal force is long.

The total torque around *A* is exactly the same as the total torque around *B*. The total torque is the only thing that affects V_{MC} , and that is the same no matter what pivot point is used.

In case (2), the story is slightly different. The rudder force is paired with the horizontal component of lift from the wings, tail, et cetera. This component arises because you are in a slight bank, as illustrated in [figure 17.4](#). The location of this force depends *indirectly* on the CM location, according to the following chain of reasoning:

a)

The large vertical component of lift must be located very close to the center of mass, to oppose the force of gravity; otherwise the airplane would be out of equilibrium in pitch.

b)

The small horizontal component of lift is located at the same place as the large vertical component. Here's another way of saying the same thing: the location of the lift vector depends directly on the shape of the airplane, but you have to adjust the shape of the airplane in order to keep the center of lift located very close to the center of mass. Note that we are not talking about the lift of the wings alone, but the lift of the entire airplane including the tail. In the particular example illustrated in [figure 17.4](#), the center of mass is located rather far forward. The tail has been adjusted to produce a negative amount of lift in order to maintain equilibrium in pitch. The horizontal component of lift depends directly on this contribution from the tail, which in turn depends on CM location.


As the center of lift moves aft, the lever arm between it and the rudder gets shorter. This means you need more rudder deflection and more bank to oppose any given amount of asymmetric thrust.

17.2.10 Effect of Drag (e.g. Landing Gear)

To reiterate: in engine-out flight you have two problems: impaired rate of climb, and asymmetric thrust which can lead to uncontrollable yaw if you're not careful.


You may be thinking that it is possible to counteract the asymmetric thrust using asymmetric drag. Technically, that's true, but as we shall see, it isn't particularly practical.

An unrealistically good type of asymmetric drag is shown in [figure 17.10](#). A source of additional drag (a small parachute) is attached far out on the wing (on the working-engine side). Because it has a long lever arm, a modest amount of drag force will create a significant amount of yaw-axis torque. This will help you maintain directional control. Of course, the drag will exacerbate your rate-of-climb problems.



[Figure 17.10](#): Asymmetric Drag (Useful, but Unrealistic)

If the parachute is attached at a different point, the results will be different. If it is attached near the working engine, as shown in [figure 17.11](#), its contribution to the yaw budget will be exactly the same as if you had throttled back the working engine; added drag is the same as reduced thrust. The effect on climb performance is also the same as if you had throttled back. Obviously, using the throttle is more convenient and practical than adding asymmetric drag.



[Figure 17.11](#): Asymmetric Drag (Useless)

Now we can do a more detailed analysis of how the landing gear contributes to the yaw-axis stability and equilibrium. Let's take the gear-up situation as a starting point, and see what *differences* arise when you put the gear down.

With the gear up, the forces are in equilibrium: thrust balances drag. With the gear down, there is extra drag. Eventually equilibrium will be restored somehow. Let's assume²⁶ the airplane just slows down, so that the extra drag of the gear will be balanced by reduced drag on the rest of the airplane.

So we have two new forces: a rearward contribution from the gear, and a forward contribution from the reduced drag on the rest of the airplane.

First, let's see what happens when the slip angle is zero. In that case the two new forces are oriented right along the line between them. This contributes nothing to the yaw-axis torque budget, because the forces

have no component perpendicular to the lever-arm between them.




Figure 17.12: Extended Landing Gear (Worse than Useless)

Next, let's see what happens when (as shown in [figure 17.12](#)) a slip angle has developed. Once again, the new force on the wheel will be mostly a drag force, rearward in the direction of the relative wind. The other new force (the reduced drag on the rest of the airplane) will act in the opposite direction, centered at a place called the center of lateral effort.

Now we have a pair of forces with a component perpendicular to the lever arm. This will create a yaw-axis torque. The torque will grow in proportion to the slip angle. On most airplanes the nose wheel is far ahead of the center of lateral effort, so this will make a *negative* contribution to the yaw-axis stability.

As a final refinement, we consider the fact that when the wheel meets the air at an angle (as shown in [figure 17.12](#)), it acts a little bit like an airfoil and produces a force *perpendicular* to the relative wind, i.e. a sideways lift force. This force grows in proportion to the slip angle and makes another negative contribution to the yaw-axis stability.

To summarize this section:

- Extending the landing gear always creates drag, which impairs the rate of climb.
- To the extent that the landing gear creates *symmetric* drag, it contributes nothing to yaw-axis equilibrium.
- The landing gear typically makes a negative contribution to yaw-axis stability.[27](#)
- Usually the only contribution that is even theoretically helpful comes from the asymmetry of having one landing gear in the propwash of the working engine. However, this is not a practical advantage since you could achieve better rate of climb, the same equilibrium, and better stability by keeping the landing gear retracted and slightly reducing power on the working engine.

Of course, during the descent and landing phases, there are some obvious advantages to extending the landing gear.

* Other Scenarios

The foregoing assumed that you maintained level flight at constant engine power, with non-constant airspeed. But this is just one of several possibilities:

- You might compensate for the added drag of the landing gear by a reduction in other drag, by way of a speed change (or otherwise). As discussed in detail above, this causes a torque that depends on the lever arm between the added drag and the reduction in other drag.
- If/when the airplane is decelerating, while things are out of equilibrium, there will be a torque that depends on the lever arm between the added drag and the center of mass.
- You might decide to prevent any speed-change by lowering the nose, compensating for the added drag by a new component of weight along the direction of flight. Einstein's principle of equivalence tells us this is just like the previous case; gravity is equivalent to an acceleration. There will be a torque that depends on the lever arm between the added drag and the center of mass.
- You might decide to prevent any speed-change by adding power on the working engine. This creates a torque that depends on the lever arm between the added drag and the engine.

The first three cases usually make a negative contribution to yaw-axis stability; they create a torque when the slip angle is nonzero. The fourth case is even worse: it creates an unhelpful torque even when the slip angle is zero.

17.2.11 Roll Control

Whenever one or more engines are producing power, propeller drag will cause a rolling moment, as discussed in [section 9.5](#). You will need to deflect the ailerons to the right to compensate.

Losing an engine will cause additional roll-axis problems on top of all your other problems. That's because the working engine creates more propwash over its wing, producing more lift on that side. You need to deflect the ailerons toward the working engine to compensate. Many airplanes have *aileron trim* to help you deal with this.

17.2.12 Critical Engine


On a typical twin, you will notice that the left engine causes more yaw trouble than the right engine does. There are several reasons for this, including helical propwash, twisted lift, and possibly P-factor.

First: Helical propwash was discussed in [section 8.4](#) in connection with single-engine airplanes. The multi-engine story is partly the same and partly different. To be specific, let's consider a plane where the engines rotate clockwise as seen from behind.


Typically, in normal flight, most of the propwash misses the vertical tail, as shown in [figure 17.13](#). However, because it spreads out on its way from the engine to the tail area, some fraction of the propwash does manage to hit the tail. The effect may be large or small, depending on the size and shape of the airplane.

You need to apply right rudder to compensate, just like in single-engine planes. If your airplane requires right rudder during the initial takeoff roll, it must be due to propwash; it can't be due to twisted lift (because there is no lift yet), and it can't be due to P-factor (because the prop disks are not inclined).

With one engine out, as long as you are able to maintain zero slip, the effect will be roughly half as large, because only one engine's propwash is acting on the vertical tail, as shown in [figure 17.14](#).



[Figure 17.13](#): Normal: Helical Propwash Effect May Be Small or Large



[Figure 17.14](#): Engine Out: Helical Propwash Has Half As Much Effect



[Figure 17.15](#): Slip: Helical Propwash Has More Effect

If you don't apply enough rudder to maintain zero slip, more of the tail will move into the propwash, as shown in [figure 17.15](#). (At low airspeeds, you could easily have a situation where you *can't* apply enough rudder to prevent this.) Since the vertical tail sticks up, not down, the propwash from the right engine will be rotating in such a way as to reduce rudder effectiveness.²⁸ If possible, you should apply additional right rudder to compensate.

It is a bit ironic that propwash affects the yaw-axis torque budget *more* when you already have a big slip angle. Normally you don't allow that to happen unless you are forced to, so this effect is usually only

noticeable at low airspeeds — such as a V_{MC} demonstration, or a crosswind takeoff (especially a crosswind from the left).

In a plane with *four* propellers, the tail will be much more affected by the propwash from the inboard engines than from the outboard engines. By using the engines one at a time, and in various pairings, you can shed a lot of light on the effects discussed in this section.

Secondly: As mentioned in [section 17.2.11](#), propeller drag creates a rolling moment and requires right aileron no matter which engine is running. This aileron deflection will produce a certain amount of twisted lift, even though the magnitude of the lift vector is the same on both sides, as discussed in [section 8.8.4](#). You will need to apply right rudder to compensate. This will be most noticeable in high-power low-air-speed situations.

Thirdly: P-factor (asymmetric disk loading) makes a small contribution to the yaw-axis torque budget. I measured this in a light twin, as discussed in [section 8.5.4](#), using both engines. The effect was small, but could be observed if you looked closely. With only one engine, the effect would be half as large.

I also calculated from theory that when the airspeed decreases from cruise to V_{MC} , the corresponding increase in angle of attack causes the center of effort of the propeller disk to move to the right by about one inch. That's not zero, but it's not very much, either.

Most of the effects that people blame on P-factor are really mainly due to a combination of adverse yaw and helical propwash.

To summarize: Some yaw contributions are unbiased, requiring rudder deflection depending on which engine is out. Working foot, working engine. These include the asymmetric thrust (as diagrammed in [figure 17.4](#)) and the increased lift over the working engine's wing (as mentioned in [section 17.2.11](#)).

Some other contributions are biased to the right, requiring right rudder no matter which engine is out. These include helical propwash acting on the tail, propeller drag acting via twisted lift, and P-factor. These are what make one engine more critical than the other.

Terminology: The engine you most regret losing is called the *critical engine*. In a twin where both engines rotate clockwise, that will be the left engine. With the left engine out, you will run out of rudder authority sooner, because the biased contributions add to the unbiased contributions. (If the right engine were out, the biased contributions would work in your favor, reducing the amount of left rudder required.)

Some twins have counter-rotating propellers. (That is, one engine rotates clockwise while the other rotates counterclockwise.) In that case both engines cause equally much yaw trouble, and either (or neither) can be considered the critical engine.

[17.3](#) Engine Out Procedures

Engine failure is an emergency. You might want to review the general discussion of emergencies in [chapter 15](#).

Make sure you know the emergency checklist for your airplane. Not all airplanes are the same. The following discussion applies to a ``generic" airplane, and serves to illustrate some important concepts, but should not be taken as a substitute for airplane-specific knowledge.

[17.3.1](#) Basic Takeoff Considerations

During takeoff, it is important to be able to detect any problems promptly. Early in the takeoff roll, you should glance at the gauges (RPM, manifold pressure, and fuel flow) to make sure the readings are normal — and that both engines are the same. Make sure the airplane ``feels" like it is pulling straight, i.e. no unusual steering effort is required to keep it going straight.

If anything funny happens while there is adequate runway remaining ahead of you, close both throttles immediately and stop straight ahead. In a high-powered airplane, such as an airliner, there will be a point where it is not possible to stop on the runway but it is possible to accelerate and fly away safely on one engine. See [section 17.3.2](#).

A light twin taking off on the same runway will use a smaller fraction of the runway for a normal takeoff, but will have worse single-engine performance. As a consequence, there will typically be a time even after liftoff when it is better to close the throttles and re-land on the remaining runway. Indeed, even if the remaining

runway is not quite enough, you might want to land on it: Suppose that because of density altitude or whatever, your aircraft has poor single-engine climb performance. You will sustain vastly less damage if you land and slide off the end of the runway at low speed, rather than making an unsuccessful attempt to climb out on one engine.

In many light twins, the climb performance is OK with the landing gear retracted but very poor with it extended. Therefore a common rule is the following: when there is no more useful runway ahead, retract the gear. If an engine fails before that point, you know you are committed to landing; if it fails after that point, you know you are committed to climbing.

17.3.2 Balanced Field Length; Takeoff Decision Speed

Sometimes you need to make a more sophisticated stop-versus-go decision. This requires a bit more pre-flight planning. The result will be expressed in terms of a *takeoff decision speed*, denoted V_1 . During the takeoff roll, note the point where the airspeed crosses V_1 . If you lose an engine before that point, stop. If you lose an engine after that point, continue the takeoff.




Figure 17.16: Accelerate-Stop and Accelerate-Go Distances

To understand V_1 , refer to [figure 17.16](#). The magenta curve is the *accelerate-go* distance, i.e. the runway length required to accelerate up to a given speed, lose an engine, and then go ahead with the takeoff. The required distance is plotted as a function of the speed at which the engine failure occurs. The curve has three parts:

- The rightmost part of the curve, where it runs horizontally, represents the runway usage in the normal case with all engines running; obviously takeoff won't be successful if the runway is shorter than this.
- If the engine failure occurs during the takeoff roll, the runway requirement is increased, because of the acceleration will be impaired. This is represented by the gently sloping section in the middle of the curve.
- If the engine failure occurs at a very low airspeed, you might have uncontrollable yaw problems. Continuing the takeoff under such circumstances is not recommended. V_{MCG} denotes the minimum control speed for ground operations.

The cyan curve in [figure 17.16](#) is the *accelerate-stop* distance, i.e. the runway length required to accelerate up to a given speed and then decide to close the throttles and stop straight ahead. It is plotted as a function of the speed at which the decision is made. If the decision is made at a low speed, very little runway is used.

This curve continues well to the right of V_R , representing the case where you actually have become airborne,²⁹ but decide to close the throttles and re-land straight ahead. If you have plenty of runway, this might be a very sensible thing to do.

The horizontal black line near the top of [figure 17.16](#) represents the actual runway length. You are allowed to use less distance than this, but you can't use more.

If you lose an engine at a low airspeed, anywhere to the left of the cyan curve, you can stop. If you lose an engine at a high airspeed, anywhere to the right of the magenta curve, you can continue the takeoff and fly away. If you are in the cross-hatched region below both curves, you are in big trouble: Given this combination

of speed and distance, you can neither continue nor stop safely. In the region above both curves, you can either stop or continue; either choice will work.

Of particular interest is the point where the accelerate-go curve crosses the accelerate-stop curve. This determines the *balanced-field length*, i.e. minimum length to guarantee that at every point during the takeoff there will be at least one good option.

If the actual runway is barely longer than the balanced-field length, there is only one value of V_1 that makes sense, namely $V_1 = V_{1bbf}$. In general, V_1 represents a decision speed, and the suffix ``bbf'' refers to ``barely balanced field''. On such a runway you need to be very aware of the point where you cross V_1 . You mustn't try to continue the takeoff when you should be stopping, and you mustn't try to stop when you should be continuing.

If there is lots of extra runway, you have some freedom to choose a V_1 value higher or lower than V_{1bbf} . The lower limit is where the accelerate-go curve (magenta) crosses the runway-length line (black). The upper limit is where the accelerate-stop curve (cyan) crosses the runway-length curve. Given such a choice, most people choose a relatively high value, since accelerate-stop is usually preferable to accelerate-go. It is better to be on the ground wishing you were in the air, instead of being in the air wishing you were on the ground. Whatever you do, don't choose an extremely-high or extremely-low value; you should distribute some of your safety margin to the accelerate-stop maneuver and some to the accelerate-go maneuver.

The details of the curves in [figure 17.16](#) will depend on factors such as braking conditions, wind, temperature, weight, type of airplane, et cetera. If you have four powerful engines and ice on the runway, V_{1bbf} could be quite low compared to V_R . On the other hand, if you have two smallish engines and really good braking, V_{1bbf} will be a larger fraction of V_R . In the extreme case where you can't climb with an engine out, the concept of balanced field loses its meaning.

- Some Pilot's Operating Handbooks will give you charts and tables for V_{1bbf} for various conditions. They call it V_1 , on the unjustifiable assumption that you will choose $V_1 = V_{1bbf}$ even if you have extra runway.
- Sometimes they don't mention V_1 at all; they just suggest a value for V_R and tabulate the accelerate-stop and accelerate-go distances for various conditions. Check both tables; whichever number is larger tells you how much runway you want to have when choosing $V_1 = V_R$.
- You can gain additional insight as to the shape of the accelerate-stop curve by looking at other tables: add the ``takeoff ground roll'' distance to the ``landing ground roll'' distance. Similarly add the ``takeoff over obstacle'' distance to the ``landing over obstacle'' distance.
- Obviously you should adjust the black line (runway length) according to the runway you will be using.

Airliners are not allowed to take off unless the available runway exceeds the balanced field length. In contrast, in general aviation, you may use a shorter runway if you want. In that case, there will be a period during the middle of the takeoff roll where you can neither stop nor continue safely on one engine. In such a case you must shut down the good engine and apply the brakes. It is much better to hit the trees at the end of the runway when you are ``almost'' stopped than to hit them when you are ``almost'' at full flying speed. This seems obvious on paper, but when you are in the cockpit it takes a lot of willpower to actually shut down the good engine. Think about this. Promise yourself that you will do it right.

17.3.3 Procedure: Low Altitude

Once you are airborne and assured of single-engine climb performance, the following checklist applies to our generic airplane at low altitudes: three things, five things, four things. Specifically:

- Three things: airspeed, ball and needle.
- Five things: mixtures, propellers, throttles, gear, flaps.
- Four things: identify, verify, feather, secure.

Now let's spell each item out in more detail, for the case where your initial speed is above V_{MC} :

- Three things: airspeed, ball, and needle. That is, pitch to maintain best-climb speed. Then the easiest thing to do is apply enough rudder pressure to center the inclinometer ball. ``Step on the ball''. This involves more rudder pressure than you need to establish zero slip, but at this stage of the game you

are in a hurry and centering the ball is a rough-and-ready approximation. With the ball centered, nonturning flight will require a slight bank toward the working engine. (Wings-level non-turning flight is really overdoing it. It involves a slip toward the dead engine, which puts an unnecessary burden on the rudder and degrades climb performance.)

- Five things: mixture controls forward (rich), propeller controls forward (fine pitch), throttles forward (maximum allowable power), landing gear retracted, flaps retracted.
- At this point you might want to check airspeed, ball, and needle again. The airplane has probably decelerated quite a bit, so you may need to make a pitch adjustment and retrim.
- Four things: identify, verify, feather, secure. Let's suppose the right engine has failed.
 - Identify: raise your right knee (dead foot, dead engine) and say aloud, "the right engine has failed".
 - Verify: retard the right throttle. There should be no change in the situation. If you retard the wrong throttle you will notice immediately; push it forward again, go back to step 1 (airspeed, ball, and needle), and try again.
 - Feather: grab the correct propeller control, pull it back a little ways and listen to make sure you've got the right one, then pull it all the way into the feather position. (If this is a simulated emergency, just pull it back half an inch or so and tell your instructor that you are simulating the feather.)
 - Secure: when the engine has stopped spinning, shut off its mixture control, its fuel supply, its boost pumps, its alternator, its magnetos, et cetera. Close its cowl flaps and open the cowl flaps on the working engine.
- Finally, check airspeed, ball, and needle again. Make sure you are trimmed for best-climb speed. Establish zero slip by applying somewhat less rudder pressure than would be necessary to center the ball. In zero-slip flight the ball will be off-center by one-third to one-half of its diameter. Use the rudder trim to hold this arrangement. In this condition, nonturning flight will require banking a few degrees toward the good engine: "raise the dead".

Here are the same items again, for the where you have a fair bit of initial altitude, but your initial speed is below V_{MC} :[30](#)

- Three things: airspeed, ball, and needle. You urgently need to dive to regain V_{MC} . As always, "step on the ball" to get the airflow approximately coordinated. You may be unable to establish zero slip, even with full rudder deflection, in which case you should apply full rudder and let the airplane establish whatever slip is necessary to oppose the asymmetric thrust. To establish nonturning flight, the wings should be almost horizontal, with the dead engine raised slightly.
- If you are in danger of losing directional control, you will have to reduce power on the good engine. Don't worry about which is which; retard *both* throttles. Skip the mixture and propeller controls for now. Retract the gear and flaps.
- Five things: After you have gotten back to V_{MC} , advance both throttles, both propeller controls, and both mixture controls. Confirm gear retracted and flaps retracted.
- Continue accelerating to best-climb speed. This will probably require cashing in additional altitude.
- Four things: identify, verify, feather, secure — the same as before.

In the case where your initial altitude and initial airspeed are both rather low, it may not be possible to regain V_{MC} by diving. In this case the procedure is rather simple: you have to close the throttles and make an immediate landing.

Reading about these things is good, but not sufficient. You really should to up with an instructor and practice these things. Practice until the right actions become routine. Review it at least once every six months.

17.3.4 Procedure: Higher Altitude

Finally, here is the procedure for the case where you have a reasonable airspeed and a reasonable altitude, say 1000 feet AGL or more. You should *not* be in any big hurry to feather the offending engine. If the problem

is minor, restarting will be a lot easier if the engine is not feathered. The checklist should be:

- Three things: airspeed, ball, and needle.
- Five things: mixtures, propellers, throttles, gear, flaps.
- Four things: identify, verify, debug, think.

Take a systematic approach to debugging. Start somewhere on the panel and then check everything you come to, systematically.

It doesn't hurt to be logical, but remember that in an actual emergency, you will be much less logical than you normally are. Unless it is obvious what the problem is, check everything, in order. Don't just check the things that come to mind. Systematic habits are more likely to stay with you.

After you've checked everything once, then try applying logic. What was the last thing you fiddled with before the failure? Did you just shut off the fuel boost pumps? Maybe you should switch them back on; look at the fuel pressure... or did you miss the boost pumps and turn off the magnetos instead? Did you just switch from the inboard to the outboard tanks? Maybe you should switch back, or switch to crossfeed.

Remember that you may be unable to climb or even maintain altitude on one engine. See [section 17.2.2](#) for a discussion of this.

17.3.5 Airspeed Management

The airspeed that gives the best single-engine rate of climb is referred to as V_{YSE} . The value of V_{YSE} for standard conditions (max weight, sea level, etc.) is marked on the airspeed indicator by a blue radial line, and is commonly called *blueline airspeed*.

If an engine fails, you should (except in certain special situations) maintain a speed at or above V_{YSE} . Maintain thine airspeed lest the ground arise and smite thee.

One exception to the foregoing rule: If you need altitude to avoid an obstacle, you'll be better off at V_{XSE} (best *angle* of climb) as opposed to V_{YSE} (best *rate* of climb). In typical trainers, the single-engine performance is so anemic that V_{XSE} will be only slightly slower than V_{YSE} , for reasons illustrated in [figure 7.8](#). Indeed, if you are above the single-engine absolute ceiling, the climb rate is negative and V_{XSE} is slightly *faster* than V_{YSE} .

Another exception: The optimal airspeed on final approach is typically less than V_{YSE} . You're not climbing, so you don't need to worry about climb performance (unless you need to go around, as discussed in [section 17.3.6](#)).

Yet another exception: Suppose your airplane has enough single-engine climb performance that the zero-climb speed V_{ZSE} is significantly slower than V_{YSE} . (See [figure 7.7](#).) Further suppose you lose an engine at night at low altitude over a dark forest, at a very low airspeed. You don't want to dive all the way to V_{YSE} , because that could take you into the trees. A much more modest dive will produce a speed above V_{ZSE} . Thereafter you can accelerate in level flight, or climb at constant airspeed. In this scenario you don't need *best* rate of climb as long as you have *some* rate of climb.

Another relevant airspeed is the minimum control airspeed, V_{MC} . As discussed in [section 17.2.7](#), you could get into big trouble if the airspeed gets too much below V_{MC} . At any speed above V_{MC} you should apply full power on the good engine and accelerate to best-climb speed. Don't be shy about diving to get to best-climb speed; remember the airplane might not be able to climb or accelerate at all at lower speeds.

At speeds below V_{MC} , you will be forced to use *less* than full power on the good engine, to keep the yaw from getting out of hand while you accelerate to V_{MC} . Losing an engine at an airspeed below V_{MC} is a really nasty situation which (for obvious reasons) most people don't practice during training. To recover, you have to partially close the throttle on the good engine, which takes a lot of willpower. You don't have much time to think. Then you have to dive, cashing in quite a lot of altitude to get the needed airspeed. The usual procedure calls for accelerating to V_{MC} plus a few knots, to give yourself a little margin, before returning the good engine to full power.

17.3.6 Engine-Out Go-Arounds

The first thing to be said about engine-out go-arounds is that you should make every possible effort to make sure that you do not ever need to perform one. The most common reason for a go-around is that you are about to land long and run off the end of the runway. Therefore, if at all possible, fly to somewhere that has a really long runway before attempting any engine-out landing.

The second thing to be said is that for typical airplanes there is a certain height above the ground — often a surprisingly great height — below which an engine-out go-around is simply not possible. The reason for this is simple: the typical approach speed is quite slow — not only below V_{YSE} (best-climb speed) but near or even below V_{ZSE} (zero-climb speed, as defined in [figure 7.7](#)). If you try to climb out at low airspeed, the rate of climb may well be negative. In order to accelerate from approach speed to any reasonable climb speed, you will need to cash in quite a lot of altitude. You will also consume time (and altitude) while you retract the landing gear, et cetera. In a Seneca, the decision to go around must be made above 400 feet AGL; below that altitude, you *are* going to touch down. If the runway is obstructed, land on the taxiway, or the infield, or whatever. If you have enough runway to touch down but not stop, consider doing a touch and go (which works better if you leave the gear down). Also consider landing anyway, with the expectation of going off the end at low speed; this is vastly preferable to hitting obstructions at high speed during an unsuccessful go-around.

17.3.7 Low-Speed Engine-Out Demonstrations

There are several key ideas I want my students to know about low-speed engine-out performance, including:

#0)

If you are ever in a position where you can descend to a safe landing *without* using high power on the good engine, by all means do so. This is not a game where you get extra points for climbing when you don't have to.

The rest of the discussion assumes you need the maximum achievable power from the good engine.

#1)

Starting from moderate speeds, as you slow down you will need more and more rudder to maintain coordinated flight. This is the coordinated regime. The amount of bank needed to maintain nonturning flight is basically constant.

#2)

There comes a point where you run out of rudder authority and cannot maintain coordinated flight. As the speed decreases further, the slip angle automatically increases, and more boat turn gets added to the pseudo boat turn. This is the uncoordinated regime. The bank angle must increase as airspeed decreases if you want to maintain nonturning flight.

#3)

You can maintain control down to V_{MC} (i.e. FAR 23.149 redline) in the takeoff configuration.[31](#)

#4)

If you persist in engine-out flight down to sufficiently low airspeed, at some point the wings and/or rudder will stall and you will be very sorry.

#5)

If you are below V_{MC} , you should reduce power on the good engine, dive to regain V_{MC} , and then re-open the throttle on the good engine.

#6)

If you are below V_{YSE} , you should dive to regain V_{YSE} , obstructions permitting.

#7)

To clear distant obstacles, you want to dive to achieve V_{XSE} as soon as possible. To clear nearby obstacles, you don't want to dive below their altitude, obviously. For a combination of obstacles, you face some tricky tradeoffs. The best solution is to make sure you never get into a low-altitude low-airspeed situation.

The aircraft manufacturer is supposed to specify a minimum safe speed for intentional engine cuts, denoted V_{SSE} , which is typically quite a bit higher than V_{MC} .

To demonstrate these key ideas, you should start in the takeoff configuration at a speed at or above V_{SSE} . Then cut one engine, and gradually decelerate. This will demonstrate idea #1 immediately. If there is a chance you will reach V_{MC} before you have a chance to demonstrate idea #2, it is a good idea to artificially limit the available rudder deflection, perhaps by blocking the pedal with the toe of your other shoe. We do not wish to demonstrate idea #4. After demonstrating flight slightly above V_{MC} (idea #3), return to V_{YSE} (idea #6) and then resume normal flight.

To demonstrate a portion of idea #5, we use a separate maneuver. Starting with both engines at idle, perform a power-off stall. Recover to V_{MC} , then using only one engine, recover to V_{YSE} .

The FAA commercial pilot multi-engine practical test standard ("PTS") contains a task called "ENGINE INOPERATIVE — LOSS OF DIRECTIONAL CONTROL DEMONSTRATION". The requirements are a bit confusing. For one thing, the PTS speaks of banking "for best performance and controllability" but doesn't say how to trade off performance versus controllability. Best climb performance typically requires less bank than best ultra-low-speed controllability.

Among many examiners, the traditions concerning this task are as follows:

1. Start at a safe altitude and safe airspeed. The PTS calls for V_{YSE} plus ten knots.
2. The PTS calls for flaps set for takeoff. However, there are some planes (such as a Seneca) where the certified takeoff checklist calls for zero flaps, and where redline is essentially equal to the stalling speed. In such planes you can lower the stalling speed by extending the flaps, which will make the demonstration safer and easier. Most examiners are happy to permit this. Call it "short field" takeoff configuration if you like. In other planes (such as an Apache) where the stalling speed is already well below redline, don't bother extending the flaps.
3. Reduce power on one engine to idle. Do not actually stop or feather the engine.
4. Depress the rudder to establish zero slip. This gives best performance.
5. Bank to establish nonturning flight. This will be a very shallow bank.
6. Block the rudder motion to ensure that you run out of rudder before the airspeed gets close to the edge of the envelope. Stay away from redline and wing-stall speed, whichever is higher. The PTS calls for staying 20 knots above the wing-stall speed.[32](#)
7. Gradually decelerate.
8. After you run out of rudder deflection, the unwritten rule is that you should not increase the bank. That means the airplane will start to turn. The turn is your signal that it is time to begin the recovery phase. You are *not* being asked to demonstrate key idea #2.

Before the checkride, you should discuss these unwritten rules with your examiner, to make sure you are both singing the same tune.

The airspeed limit is needed to ensure safety. The artificial limits on rudder deflection and bank are needed so that you can demonstrate a nice gentle boat turn, by pretending to run out of control authority; otherwise the airplane would be controllable at all safe airspeeds and there would be nothing to demonstrate.

Note that in everyday (non-checkride) flight, if you run out of rudder authority at a speed above redline (and if you are sure you want to be flying so slowly) you would just smoothly enter the uncoordinated regime and increase the bank.

You should *not* do demonstrations the way FAR 23.149 seems to suggest:

- Do not suddenly shut down one engine at a low airspeed. Shut it down at or above V_{YSE} and then decelerate. (Alternatively, I suppose it would be safe to fly below V_{YSE} and *gradually* reduce power on one engine, but I can't think of a reason why you would want to.)
- Do not explore low-speed performance at low altitude. The practical test standard calls for a minimum of 3000 feet AGL.

Full-blown FAR 23 V_{MC} determinations should be left to professional test pilots. For that matter, not even test pilots dare to experiment with loss of control at low altitude. They are not crazy; they experiment at a series of safe altitudes and then extrapolate.

1

A single-engine aircraft is required (by FAR 23.49) to have a stalling speed of 61 knots or less, and in certain models it is quite a bit less. This is important for safety in case of a forced landing. In contrast, a twin with sufficiently good engine-out performance is exempt from this restriction. The theory is that a twin that can climb on one engine should never need to make an off-airport landing.

2 We won't discuss aircraft that have centerline thrust, e.g. the Cessna 337 Mixmaster.

3 This means total lift, including the contributions of the wings, horizontal tail, et cetera. The center of lift will be located quite close to the center of mass, for reasons discussed in [section 6.1.3](#).

4 The effects are reduced if the working engine is developing less power (because of small engine size, high altitude, and/or reduced throttle setting), the airspeed is high, and the rudder is large. Conversely, in a twin with a small rudder, a large engine, full power, and a low airspeed, a sudden failure will definitely get your attention.

5 If you forget to roll the wings level before using the rudder to stop the heading change, you could easily find yourself stepping on the wrong rudder. For instance, if you are in a turn to the right and the left engine fails, you might be tempted to stop the turn by stepping on the left (wrong!) rudder.

6 To know how much bank and/or how much inclinometer ball deflection corresponds to zero slip, you can (a) recall from your training flights what configuration corresponds to best performance, (b) recall from flights with a slip string what configuration corresponds to zero slip, or (c) let the inclinometer ball go off-center by half its width, which is usually ``close enough" to the right answer.

7 See [section 19.4](#) for a discussion of E-down and related concepts.

8 The confusion is understandable, since asymmetric thrust is about the only way you can maintain an inclination without being in a slip of some kind.

9 ... or cruise performance, for that matter — engine-out or otherwise.

10 In fact, an undeflected rudder produces a less stall-resistant shape, which will probably stall at a higher airspeed.

11 If redline (as defined in the following section) is down near the bottom of the green arc, it is a good guess that wing stall is what limits the airplane's low-speed controllability. Conversely, if redline is much higher than the bottom of the green arc, you can guess that rudder stall is what limits the low-speed controllability (unless the redline is artificially high because of the arbitrary 5 degree bank limit in FAR 23.149).

12 A very similar regulation, FAR 25.149, applies to transport-category aircraft (e.g. airliners).

13 The notion of ``critical engine" is discussed in [section 17.2.12](#).

14 If the airplane weighs more than 6000 pounds, FAR 23.66 requires the airplane to be able to climb with one engine inoperative, at an airspeed ``equal to that achieved at 50 feet" after takeoff. Even this does not require climb at V_{MC} .

15 There are of course other redlines: at the high end of the airspeed indicator, on the tachometers, etc., but they are not relevant to the present discussion.

16 ... except perhaps in a minor way that is not relevant here. See [section 11.5.6](#) for a general discussion of slip angle versus bank angle.

17 There are many airplanes that are quite nicely behaved even under conditions that require more than 5 degrees of bank in order to maintain non-turning engine-out flight.

18 That's right: there is no guarantee that 5 degrees is safe. It is commonly assumed that ``the manufacturer must have tested a 5 degree bank because that is the maximum allowed". But in fact, best control might be achieved at 2 or 3 degrees, and there is no reason to assume that the manufacturer ever tried more than that. Remember, in non-slipping flight the bank required is quite modest (and independent of airspeed), and there is not much that the manufacturer can achieve by slipping that could not be better achieved by more rudder deflection.

19 This means the takeoff configuration as specified in the Pilot's Operating Handbook or Airplane Flight Manual. Remember that these documents are legally part of the airplane. You can't have a certified V_{MC} without a certified takeoff checklist.

[20](#)

In this example, we are assuming that wing stall (not rudder stall) is what limits low-speed handling.

[21](#)

This option is available to us because, for this airplane, the certified takeoff configuration did not call for flaps to be extended.

[22](#)

If there were an unbalanced torque, the airplane would not only turn, it would accelerate around the yaw axis.

[23](#)

We are talking about rather modest bank angles, perhaps 15 or 20 degrees, that would not get you into trouble in other circumstances.

[24](#)

It may be a good idea to check the wing's stalling speed by performing a stall with both engines at zero thrust. The zero-engine stalling speed won't be quite the same as the one-engine stalling speed, but it should be a useful estimate.

[25](#)

I have done calculations that indicate that for certain light trainers, the wings will almost always stall before the rudder, but you absolutely should not assume that this is true for all airplanes in all circumstances.

[26](#)

At the end of this section, other possibilities will be considered.

[27](#)

One could imagine designing an airplane with the landing gear so far aft that they were behind the lateral center of effort, in which case they would increase yaw-axis stability.

[28](#)

The propwash from the left engine is actually helpful.

[29](#)

When we consider the possibility of taking off and re-landing straight ahead, we need to slightly relabel the horizontal axis in [figure 17.16](#). In all cases, what matters is the amount of mechanical energy you must get rid of in order to stop the plane. Before liftoff, speed is the only contribution to the energy, so we can label the axis either as speed or as square root of energy; they are equivalent. After liftoff, the latter (Öenergy) remains correct but the former (airspeed) does not.

[30](#)

It's a little hard to see how this situation could arise in the course of normal flying. However, (a) such a situation is sometimes created as part of a training exercise, and (b) it could arise if the pilot mishandles an engine-out situation, squandering the initial airspeed.

[31](#)

In other configurations, you can maintain control down to redline or V_S , whichever is higher.

[32](#)

In airplanes where V_{MC} is at or near V_S , $V_S + 20$ may seem like a generous or even excessive margin of safety. In other airplanes, however, $V_S + 20$ is not nearly enough. You need to be careful, since there are plenty of airplanes where $V_S + 20$ is near (or even below) V_{MC} . A better criterion might be to stay above $V_S + 10$ and above redline + 10, whichever is higher.

18 Stalls and Spins

Caution: Cape does not enable user to fly.

—
warning label on Superman costume sold at Walmart

Spins are tricky. After reading several aerodynamics texts and hundreds of pages of NASA spin-tunnel research reports, I find it striking how much remains unknown about what happens in a spin.

18.1 Stalls: Causes and Effects

Here's a basic yet important fact: if you don't stall the airplane, it won't spin. Therefore, let's begin by reviewing stalls.

As discussed in [section 5.3](#), the stall occurs at the critical angle of attack, which is defined to be the point where a further increase in angle of attack does not produce a further increase in coefficient of lift.

Nothing magical happens at the critical angle of attack. Lift does not go to zero; indeed the coefficient of lift is at its maximum there. Vertical damping goes smoothly through zero as the airplane goes through the critical angle of attack, and roll damping goes through zero shortly thereafter. An airplane flying 0.1 degree beyond the critical angle of attack will behave itself only very slightly worse than it would 0.1 degree below.

If we go far beyond the critical angle of attack (the ``deeply stalled" regime) the coefficient of lift is greatly reduced, and the coefficient of drag is greatly increased. The airplane will descend rapidly, perhaps at thousands of feet per minute. Remember, though: the wing is *still supporting the weight of the airplane*. If it were not, then there would be an unbalanced vertical force, and by Newton's law the airplane would be not only descending but *accelerating* downward. If the wings were really producing zero force (for instance, if you snapped the wings off the airplane) the fuselage would accelerate downward until it reached a vertical velocity (several hundred knots) such that weight was balanced by fuselage drag.

18.2 Stalling Part vs. All of the Wing

We can arbitrarily divide the wing into sections; each section contributes something to the lift of the whole wing. It is highly desirable (as discussed in [section 5.4.3](#)) to have the coefficient of lift for sections near the wing-root reach its maximum early, and start decreasing, while the coefficient of lift for sections near the tips continues increasing¹ (as a function of angle of attack).

Therefore it makes perfect sense to say that the sections near the roots are stalled while the sections near the tips are not stalled. If only a small region near the root is stalled, the wing *as a whole* will still have an increasing coefficient of lift — and will therefore not be stalled.

We see that the wing will continue to produce lots of lift well beyond the point where part of it is stalling. This is the extreme slow-flight regime — you can fly around all day with half of each wing stalled (although it takes a bit of skill and might overheat the engine).

18.3 Boundary Layers

There is a very simple rule in aerodynamics that says the velocity of the fluid right next to the wing (or any other surface) is *zero*. This is called the *no-slip boundary condition*. Next to the surface there is a thin layer, called the *boundary layer*, in which the velocity increases from zero to its full value.

18.3.1 Separated versus Attached Flow

The wing works best when the airflow is *attached* to the wing surface by a simple boundary layer. The opposite of attached flow is *separated flow*.

For attached flow, as we move through the boundary layer from the wing surface out to the full-speed flow, there is practically no pressure change. Sometimes it helps to think about attached flow in the following way: Imagine removing the boundary layer and replacing it with a layer of putty that redefines the shape of the wing. Then imagine ``lubricating" the new wing so that the air slides freely past it; the no-slip condition no longer applies. Bernoulli's principle can be used to calculate the pressure on the surface of the putty; obviously it could never be applied inside the boundary layer. The putty-covered wing may not be the most desirable shape, but it won't necessarily be terrible.

For separated flow, the putty model does not work. Suppose I want to pick up a piece of lint from the floor using a high-powered vacuum cleaner. If I keep the hose 3 feet away from the floor, it will never work; I could have absolute zero pressure at the mouth of the hose, but the low pressure region would be ``separated" from the floor and the lint. If I move the hose closer to the floor, eventually it will develop low pressure near the floor. This is part of the problem with separated flow: there is low pressure somewhere, but not where you need it. Separation can have multiple evil effects:

- Separation means the air doesn't follow the contour of the wing. This is somewhat like having a really thick boundary layer. The wing can't force the air into the optimal flow pattern, so not as much low pressure is produced.
- Whatever low pressure is produced isn't all attached to the wing surface. This is a new problem that an attached flow would not have, no matter how thick the boundary layer.
- On a non-streamlined object such as a golf ball, there is a lot of drag (specifically: form drag, as discussed in [section 4.3](#)) because separation disrupts a desirable high-pressure area behind the ball.

18.3.2 Laminar versus Turbulent Flow

In the simplest case, there is *laminar flow*, in which every small parcel of air has a definite velocity, and the velocity varies smoothly from place to place. The other possibility is called *turbulent flow*, in which:

- at any given point the velocity fluctuates as a function of time, and
- at any given time the velocity changes rapidly as we move from point to point, even for nearby points.

The closer we look, the more fluctuations we see.

Attached turbulent flow produces a lot of mixing. Some bits of air move up, down, left, right, faster, and slower relative to the average rearward flow.

For separated laminar flow, there will be some reverse flow (noseward, opposing the overall rearward flow) but the pattern in space will be much smoother than it would be for turbulent flow, and it will not fluctuate in time.

You can tell whether a situation is likely to be turbulent if you know the *Reynolds number*. You don't need to know the details, but roughly speaking small objects moving slowly through viscous fluids (like honey) have low Reynolds numbers, while large objects moving quickly through thin fluids (like air) have high Reynolds numbers. Any system with a Reynolds number less than about 10 is expected to have laminar flow everywhere. If you drop your FAA ``Pilot Proficiency Award" wings into a jar of honey, they will settle to the bottom very slowly. The flow will be laminar everywhere, since the Reynolds is slightly less than 1. There will be no separation, no turbulence, and no form drag — just lots of skin-friction drag.

Systems with Reynolds numbers greater than 10 or so are expected to create at least some turbulence. Airplanes operate at Reynolds numbers in the millions. The wing will have a laminar boundary layer near the leading edge, but as the air moves back over the wing, at some point the boundary layer will become turbulent. This is called the *transition to turbulence* or simply the *boundary layer transition*. Also at some point (before or after the transition to turbulence) the airflow will become separated. The designers try to keep the region of separation rather small and near the trailing edge. In order to make a wing develop a lot of lift without stalling, it helps to minimize the amount of separation.

18.3.3 Boundary Layer Control

One scheme² for controlling separation involves the use of vortex generators (VGs). The VGs are the little blades you see on the top of some wings, sticking up into the airstream at funny angles. Each blade works like a turnplow, reaching out into the high-velocity airstream and turning the layers over — plowing energy into the inner layers.

Re-energizing the boundary layer allows the wing to fly at higher angles of attack (and therefore higher coefficients of lift) without stalling. This improves your ability to operate out of short and/or obstructed fields.

The vorticity created by these little VGs should not be confused with the bound vortex, the big vortex that generates the circulation that supports the weight of the airplane. As discussed in [section 3.12](#), to create lift you have to make the air circulate *around* the wing; that is, there must be vortex line running along the span.

VGs don't do that; their vortex lines run chordwise, not spanwise.³

Boundary-layer turbulence (whether created by VGs or otherwise) also helps prevent separation, once again by stirring additional energy into the inner sublayers of the boundary layer.

On a golf ball, 99% of the drag is form drag, and only 1% is skin-friction drag. The dimples in the golf ball provoke turbulence, adding energy to the boundary layer. This allows the flow to stay attached longer, maintaining the high-pressure region behind the ball, thereby decreasing the amount of form drag. The turbulence of course increases the amount of skin-friction drag, but it is worth it.⁴

Bernoulli's principle does not apply inside the boundary layer, separated or otherwise. As discussed in [section 3.4](#), Bernoulli's principle applies in situations where pressure (potential energy) and airspeed squared (kinetic energy) add up to a constant. This is not the case in the boundary layer, because friction there converts a significant amount of the energy into heat.

Do VGs play the same role as dimples on a golf ball? Not exactly. Unlike a golf ball, a wing is supposed to produce lift. Also unlike a golf ball, a wing is highly streamlined; consequently, its form drag is not predominant over skin-friction drag. VGs are typically used to improve lift at high angles of attack (by fending off loss of lift due to separation). They may or may not improve performance at low angle of attack (by decreasing form drag at the expense of skin-friction drag).

If you want ultra-low drag, and don't care about short-field performance, you want a wing with as much laminar flow as possible. Designing a ``laminar flow wing'' is exquisitely difficult, especially in the real world where the laminar flow could be disturbed by rain, ice, mud, and splattered bugs on the leading edge.

There is always *some* separation on every airfoil section. The separation grows as the angle of attack increases. If there is too much separation, it cuts into the wing's ability to produce lift. If there were no separation, the wing could continue producing lift up to very high angles of attack (thereby achieving very high coefficients of lift).

Having lots of separation is the dominant cause (but not the definition) of stalling.⁵ Remember: the stall occurs at the critical angle of attack, i.e. the point where max coefficient of lift is attained.

[18.3.4](#) Summary

A full discussion of turbulence and/or separated flow is beyond the scope of this book; indeed, trying to *really* understand and control these phenomena is a topic of current research. There is nothing simple about it. But there are a few things we can say.

- The opposite of *separated flow* is *attached flow*.
- The opposite of *turbulent flow* is *laminar flow*.
- Separated flow need not be very turbulent, nor vice versa.
- Laminar flow need not be attached, nor vice versa.
- Turbulence doesn't cause separation (and indeed often helps prevent it).

For more information, see e.g. [reference 17](#).

[18.4](#) Coanda Effect, etc.

The name *Coanda effect* is generally applied to any situation where a thin, high-speed jet of fluid meets a solid surface and follows the surface around a curve. Depending on the situation, one or more of several different physical processes might be involved in making the jet follow the surface.

As a pilot, you absolutely do *not* need to know about the Coanda effect or what causes it. Indeed, many professional aerodynamicists get along just fine without really understanding such things. The main purpose of this section is to dispel the notion that a normal wing produces lift ``because'' of some type of Coanda effect.

Using the Coanda effect to explain the operation of a normal wing makes about as much sense as using bowling to explain walking. To be sure, bowling and walking use some of the same muscle groups, and both at some level depend on Newton's laws, but if you don't already know how to walk you won't learn much by

considering the additional complexity of the bowling situation.

18.4.1 Tissue-paper Demonstration

You can demonstrate one type of Coanda effect for yourself using a piece of paper. Limp paper, such as tissue paper, works better than stiff paper. Drape the paper over your fingers, and then blow horizontally, as shown in the following figures.⁶



Figure 18.1: Tissue Paper; No Coanda Effect

If the jet passes just above the paper, as shown in [figure 18.1](#), nothing very interesting happens. The jet just keeps on going. The paper is undisturbed.




Figure 18.2: Tissue Paper; Coanda Effect

On the other hand, if the jet actually hits the paper as shown at point C in [figure 18.2](#), the downstream part of the paper will rise up. This is because the air follows the curved surface; as it does so, it creates enough low pressure to lift the weight of the paper.

The air in your lungs, at point A, is at a pressure somewhat above atmospheric. At point B, after emerging from the nozzle, the air in the jet is at atmospheric pressure.

As discussed in [section 3.3](#), the fact that the fluid follows a curved path proves that there is a force on it. This force must be due to a pressure difference. In this case, the pressure on the lower edge of the jet (where it follows the curve of the tissue paper near point D) is less than atmospheric, while the pressure on the upper edge of the jet (near point E) remains more-or-less atmospheric. This pressure difference pulls down on the jet, making it curve. By the same token it also pulls up on the paper, creating lift.

People who only half-understand Bernoulli's principle will be surprised to hear that the jet leaves the nozzle at high speed at atmospheric pressure. It's true, though. In particular, the crude statement that "high velocity means low pressure" is an oversimplification that cannot be used in this situation. The correct basis of Bernoulli's principle is that *for a particular parcel of air* the mechanical energy (pressure plus kinetic energy per unit volume) remains more-or-less constant. If you want to compare two *different* parcels of air, you'd better make sure that they started out with the same mechanical energy.

In this case, the air in the jet leaves the nozzle with a higher mechanical energy than the ambient air. Your lung-muscles are the source of the extra energy.

When this high-velocity, atmospheric-pressure air smacks into the paper at point C, it actually creates above-atmospheric pressure there. Indeed, we can use the streamline-curvature argument again: if the air turns a sharp corner, there must be a very large pressure difference.


In order to make this sharp turn, the air needs something to push against. A good bit of the required momentum comes from the air that splatters backward, as suggested by the squiggles just below and upstream of the point of contact. This process is extremely messy. It is much more complicated than anything that happens near a wing in normal flight. To visualize this splatter, blow a jet of air onto a dusty surface.⁷ Even if you blow at a very low angle, some of the dust particles blow away in the direction *opposite* to the main flow.

18.4.2 Blowing the Boundary Layer

Since we saw in [section 18.3](#) that de-energizing the boundary layer is bad, you might think adding energy to

the boundary layer should be good... and indeed it is. One way of doing so uses vortex generators, as discussed in [section 18.3](#). [Figure 18.3](#) shows an even more direct approach.

- We use a pump to create a supply of air at very high pressure.
- The air comes out a nozzle. The result is a jet of high-velocity air at the same pressure as the local air.⁸
- The jet shoots out of a slot in the top of the wing, adding energy to the boundary layer at a place where this could be very helpful.



[Figure 18.3](#): Blowing the Boundary Layer

Once again, the Coanda effect cannot explain how the wing works; you have to understand how the wing works before you consider the added complexity of the blower.

In this case we expect one spectacular added complexity, namely *curvature-enhanced turbulent mixing*. This phenomenon will not be discussed in this book, except to say that it does not occur near a normal wing, while it is likely to be quite significant in the situation shown in [figure 18.3](#).

Curving flows with lots of shear can be put to a number of other fascinating uses, but a discussion is beyond the scope of this book. See [reference 9](#).

[18.4.3](#) Teaspoon Demonstration

Another example of a jet following a curved surface uses a jet of water. You can easily perform the following experiment: let a thin stream of water come out of the kitchen faucet. Then touch the left side of the stream with the convex back side of a spoon. The stream will not be pushed to the right, but instead it will follow the curve of the spoon and be pulled to the left. The stream can be deflected by quite a large amount. In accordance with Newton's third law of motion, the spoon will be pulled to the right.

I don't understand everything I know about this situation, but it is safe to say the following:

1. This water-in-air jet differs in fundamental ways from the air-in-air jet situation described above.
2. This effect has practically nothing to do with the way a normal wing produces lift.

To convince yourself of these facts, it helps to have a higher velocity and/or a larger diameter than you can conveniently get from a kitchen faucet. A garden hose will give you a bigger diameter, and if you add a nozzle you can get a higher velocity. You can easily observe:

- The amount of lift⁹ you can produce is pathetically small, compared to the dynamic pressure and area of the water jet.
- The lift-to-drag ratio is terrible. Indeed this makes it very hard to measure the lift; if you get the angle slightly wrong you will inadvertently measure a drag component instead.
- The water spreads out when it hits the surface, making a thin coating over a wide area of the surface. This is in marked contrast to what happens in the air-in-air jet, as you can demonstrate by placing thin strips of tissue paper side by side. You can easily blow on one strip and lift it without disturbing its neighbors.
- Some of the spreading layer flows backwards, ahead of the point of contact of the jet, corresponding to a negative amount of upwash. This is grossly different from what happens near a real wing.
- The effect does not depend on curvature-enhanced turbulent mixing with the ambient air. This is quite unlike what happens in a real airplane with boundary-layer blowing.

It appears that surface tension plays two very important roles:

1. At the water/air interface it prevents mixing of the air and water.

2. At the water/wing interface it plays a dominant role in making the water stick to the surface.

In both respects this is quite unlike the air-in-air jet, where the air/wing surface tension has no effect and there is no such thing as air/air surface tension.

To convince yourself of this: Take a thin sheet of plastic. Get it wet on both sides, and drape it over a cylinder. You will not be able to lift it off the cylinder using a tangential water jet. The surface tension holding the wet plastic to the cylinder is just as strong as the tension between the plastic and the jet. In contrast, when the same piece of plastic has air on both sides, you can easily lift it off the cylinder using an air jet.

18.4.4 Fallacious Model of Lift Production

You may have heard stories saying that the Coanda effect explains how a wing works. Alas, these are just fairy tales. They are worse than useless.

1. For starters, these fairy tales often claim that blowing on tissue paper (as described just above) proves that "high velocity means low pressure" which is absolutely not what is being demonstrated. The high-velocity air coming out of your mouth is at atmospheric pressure. If you blow across the top of a *flat* piece of paper, it will not rise, no matter what you do. There is no low pressure in the jet (unless and until it gets pulled around a corner). Therefore the Coanda stories give a wrong explanation of normal wings and basic aerodynamics. And by the way, such stories cannot even begin to explain the operation of flat wings — yet we have seen in [section 3.10.1](#) that a barn door doesn't behave very differently from other airfoils.
2. The Coanda-like notion of airflow following a curved surface cannot possibly explain why there is upwash in front of the wing. In [figure 18.2](#) there must be a stagnation point on the upper surface of the paper near point *C*. This is completely different from the situation near a normal wing, where the stagnation line must be somewhere below the leading edge of the wing. Upwash is important, since it contributes to lift while creating a *negative* amount of induced drag. A further consequence, by the way, is that these Coanda-like stories cannot possibly explain the operation of stall-warning devices, as discussed in [section 3.5](#).
3. As mentioned above, the distribution of velocities necessary to create curvature-enhanced turbulent mixing is produced by a high-speed jet but is not produced by a normal wing.
4. Sometimes the fairy tales say that the jet "sticks" to the surface because of viscosity. This implies that if the viscosity of the fluid changes, the amount of lift an airfoil produces should change in proportion. In fact, though, the amount of lift produced by a real wing is independent of viscosity over a wide range. Also, many of the processes responsible for the real Coanda effect require the production of turbulence, so they only work if the viscosity is sufficiently *low*.[10](#)
5. In the real Coanda effect, we know where the high-velocity air comes from. It comes from a nozzle. Upstream of the nozzle is a pump (or a rocket engine, or some other device) to supply the necessary energy. The jet makes high-velocity air above the wing, not below, because that's where we aim the nozzle. An ordinary wing is completely different. It is wonderfully effective at creating high-velocity air above itself, without nozzles, without pumps, and without transferring energy[11](#) to the air.
6. The fairy tales generally neglect the fact that the wing speeds up the air in its vicinity, and just assume that the relative wind meets the wing at the free-stream velocity and follows the curve in a Coanda-like way. As a consequence, they miscalculate the pressure gradients by a factor of ten or so.
7. Finally, in the real Coanda effect we know how big the jet is. Its initial size is determined by the nozzle. The amount of mixing depends on the speed of the jet, the speed of the ambient air, the curvature of the surface, and other known quantities. Awareness of the Coanda effect is a small part of — not a replacement for — a full analysis of the wing in [figure 18.3](#). In contrast, (a) the typical fairy tales imply that the entire flow pattern of a normal wing can be explained by mentioning the magic words "Coanda effect", yet (b) they cannot explain how thick a chunk of air is deflected by the wing. One inch? Six inches? A chord-length? A span-length? Some amount proportional to the viscosity of the air? It would be very hard to calculate how much. [12](#)

18.4.5 Summary

Don't let anybody tell you that squirting a spoon or blowing on tissue paper is a good model of how a wing works.

If you want to ``get the feel" of lift production, the obvious methods are the best. These include holding a model airfoil¹³ downstream of a household fan, or sticking it out of a car window.


18.5 Spin Entry

Case 1: In normal flight, rolling motions are very heavily damped, as discussed in [section 5.4](#). Even though the static stability of the bank *angle* is small or even negative, you cannot get a large roll *rate* without a large roll-inducing force; when you take away the force the roll rate goes away.

Case 2: Near the critical angle of attack, the roll damping goes away. Suppose you start the aircraft rolling to the right. The roll rate will just continue all by itself. The right wing will be stalled (beyond max lift angle of attack) and the left wing will be unstalled (below max lift angle of attack).

Case 3: At a sufficiently high initial angle of attack (somewhat greater than the critical angle of attack), the roll will not just continue but accelerate, all by itself. This is an example of the ``departure"¹⁴ that constitutes the beginning of a snap roll or spin. The resulting undamped rolling motion is called *autorotation*.

At a high enough angle of attack, the ailerons lose effectiveness, and at some point they start working in reverse.¹⁵ [Figure 18.4](#) shows how this reversal occurs. Suppose you deflect the ailerons to the left. This raises the angle of attack at the right wingtip and lowers it at the left wingtip. Normally, this would increase the lift on the right wing (and lower it on the left), creating a rolling moment toward the left. Near the critical angle of attack, though (as seen in the left panel of the figure), raising or lowering the angle of attack has about the same effect on the coefficient of lift, so no rolling moment is produced (for now, at least).



[Figure 18.4](#): Lift and Drag at Departure

We see that at this angle of attack, anything that creates a rolling mo-ment will cause the aircraft to roll like crazy, and indeed to keep accelerating around the roll axis. There will be no natural roll damping, and you will be unable to oppose the roll with the ailerons.

There are two main ways of provoking a spin at this point:

1. Suppose the airplane is in a steady slip to the left. That is, you are steadily pushing on the right rudder pedal. Then the slip/roll coupling (as discussed in [section 9.1](#) and [section 9.2](#)) will cause it to spin to the right.
2. Suppose the airplane is not in much of a slip, but you suddenly cause it to yaw to the right. The left wingtip will temporarily be moving faster, and the right wingtip will temporarily be moving slower. This difference in airspeeds will create a difference in lift, causing a spin to the right. The initial yawing motion could come from a sudden application of rudder, or from adverse yaw, or what-ever. Note that in the right panel of [figure 18.4](#), the aileron deflection has a tremendous effect on the drag. This means that ailerons deflected to the left cause a yaw to the right which in turn provokes a roll to the just the opposite of what ailerons normally do.

18.6 Types of Spin


18.6.1 Spin Modes

The word ``spin" can be used in several different ways, which we will discuss below. The spin family tree includes:

- ``departure", i.e. onset of undamped rolling;
- incipient spin — i.e. one that has just gotten started; or
- well-developed spin, which could be


- a steep spin, or
- a flat spin.

[Figure 18.5](#) shows an airplane in a steady spin. You can see that the direction of flight has two components: a vertical component (down, parallel to the spin axis) and a horizontal component (forward and around).




[Figure 18.5](#): Airplane in a Steady Spin

[Figure 18.6](#) is a close-up of a wing in a steep spin. We have welded a pointer to each wingtip, indicating the direction from which the relative wind would come if the wing were producing zero lift; we call this the Zero-Lift Direction (ZLD). (For a symmetric airfoil, the ZLD would be aligned with the chord line of the wing.) Remember that the angle between the direction of flight and the ZLD pointer is the angle of attack.



[Figure 18.6](#): Steep Spin — Geometry



[Figure 18.7](#): Steep Spin — Coefficient of Lift

In this situation, both wingtips have the same vertical speed, but they have significantly different horizontal speeds — because of the rotation. Consequently they have different directions of flight, as shown in the figure. This in turn means that the two wingtips have significantly different angles of attack, as shown in [figure 18.7](#). The two wings are producing equal amounts of lift, even though one is in the stalled regime and one in the unstalled regime.

[Figure 18.8](#) shows another spin mode. This time the rotation rate is higher than previously. The spin axis is very close to the right wingtip. The outside wing is still unstalled, while the inside wing is very, very deeply stalled, as shown in [figure 18.9](#).




Figure 18.8: Flat Spin — Geometry




Figure 18.9: Flat Spin — Coefficient of Lift




Figure 18.10: Doubly Stalled Flat Spin — Coefficient of Lift

Figure 18.10 shows yet another possible spin mode. In this case, the outside wing is stalled, while the inside wing is, of course, much more deeply stalled. Whether this spin mode, or the one shown in figure 18.9 (or both or neither) is stable depends on dozens of details (aircraft shape, weight distribution, et cetera). There is a common misconception that in a spin, one wing is stalled and the other wing is always unstalled. This is true of ``most'' spins but it is not a defining property. It would be safer to use the following definition:

In a spin, at least one wing is stalled,
and the two wings are operating at very different angles of attack.

18.6.2 Samaras, Flat Spins, and Centrifugal Force

A samara is a winged seed. Maples are a particularly well known and interesting example.

Maple samaras have only one wing, with the seed all the way at one end. Its mode of flight is analogous to an airplane in a flat spin. In an airplane, the inside wing is deeply stalled, while in the samara the inside wing is missing entirely.

In a non-spinning airplane, if one wing were producing more lift than the other, that wing would rise. So the question is, why is a flat spin stable? Why doesn't the outside wing continue to roll to ever-higher bank angles? The secret is *centrifugal force*.¹⁶ Suppose you hold a broomstick by one end while you spin around and around; the broomstick will be centrifuged outward and toward the horizontal.

In an airplane spinning about a vertical axis, the high (outside) wing will be centrifuged outward and downward (toward the horizontal), while the low (inside) wing will be centrifuged outward and upward (again toward the horizontal). In a steady flat spin, these centrifugal forces cancel the rolling moment that results from one wing producing a lot more lift than the other. This is the only example I can imagine where an airplane is in a steady regime of flight but one wing is producing more lift than the other.

As discussed in [reference 6](#), an aircraft with a lot of mass in the wings will have a stronger centrifugal force than one with all the mass near the centerline of the fuselage. In particular, an aircraft with one pilot and lots of fuel in the wing tanks could have completely different spin characteristics than the same aircraft with two pilots and less fuel aboard.

18.6.3 NASA Spin Studies

In the 1970s, NASA conducted a series of experiments on the spin behavior of general-aviation aircraft; see [reference 8](#) and [reference 7](#) and other papers cited therein. They noted that there was "considerable confusion" surrounding the definition of steep versus flat spin modes, and offered the classification scheme shown in [table 18.1](#).

spin mode	Steep	Mod'ly Steep	Mod'ly Flat	Flat
angle of attack	20 to 30	30 to 45	45 to 65	65 to 90
nose attitude	extreme nose-down		less nose-down	
rate of descent	very rapid		less rapid	
rate of roll	extreme		moderate	
rate of yaw	moderate		extreme	
wingtip-to-wingtip difference in angle of attack	modest		large	
nose-to-tail difference in slip	large		large	

[Table 18.1](#): Spin Mode Classification

The angle of attack that appears in this table is measured in the aircraft's plane of symmetry; the actual angle of attack at other positions along the span will depend on position.

The NASA tests demonstrated that general aviation aircraft not approved for intentional spins commonly had unrecoverable flat spin modes.

18.6.4 Effects of Spin-Axis Changes

In all cases NASA studied, the flat spin had a *faster* rate of rotation (and a slower rate of descent) than the steep spin. Meanwhile, [reference 15](#) reports experiments in which the flatter pitch attitudes were associated with the *slower* rates of rotation. This is not a contradiction, because the latter dealt with an unsteady spin (with frequent changes in pitch attitude), rather than a fully stabilized flat spin. A *sudden change* to a flatter pitch attitude will cause a *temporary* reduction in spin rate, for the following reason.

In any system where angular momentum is not changing, the system will spin faster when the mass is more concentrated near the axis of rotation. The general concept is discussed in [section 19.8](#). By the same token, if the mass of a spinning object is redistributed farther from the axis, the rotation will slow down.

When the spinning airplane pitches up into a flatter attitude, whatever mass is in the nose and tail will move farther from the axis of rotation. Angular momentum doesn't change in the short run, so the rotation will slow down in the short run.

In the longer run — in a steady flat spin — the aerodynamics of the spin will pump more angular momentum into the system, and the rotation rate will increase quite a lot. The rotation rate of the established flat spin is typically twice that of the steep spin.

Recovering from an established flat spin requires forcing the nose down. This brings the mass in the nose and tail closer to the axis of rotation. Once again using the principle of conservation of angular momentum, you can see that the rotation rate will increase (at least in the short run) as you do so — which can be disconcerting.

18.7 Recovering from a Spin

If you find yourself in an unusual turning, descending situation, the first thing to do is decide whether you are in a spiral dive or in a spin. In a spiral dive, the airspeed will be high and increasing; in a spin the airspeed will be low. You should be able to *hear* the difference. Also, the rate of rotation in a spiral is much less; the high speed means the airplane has lots of momentum and can't turn on a dime. In a spin, the aircraft will be turning a couple hundred degrees per second.

To get out of a spin,^{[17](#)} follow the spin-recovery procedures given in the Pilot's Operating Handbook for your airplane. The literature is full of home-brew spin recovery procedures that probably work most of the time in most airplanes, but if you want a procedure that works for sure, follow the handbook for your airplane.

For typical airplanes, the spin recovery procedure contains the following items:

- Retard the throttle to idle
- Retract the flaps
- Neutralize the ailerons
- Apply full rudder in the direction opposing the spin
- Briskly move the yoke to select zero angle of attack.

Now let's discuss each of these items in a little more detail.

Retarding the throttle is a moderately good idea for a couple of reasons. For one thing (especially if you have a fixed-pitch prop) it keeps the engine from overspeeding during the later stages of the spin recovery. More importantly, gyroscopic precession of the rotating engine and propeller can hold the nose up, flattening the spin and interfering with the recovery (depending on the direction of spin).

Propwash might increase the effectiveness of the horizontal tail and therefore *assist* in the spin recovery, but (especially in a flat spin) the propwash could be blown somewhere else by the abnormal airflow — so you may not be able to count on this.

Retracting the flaps is a moderately good idea because you might exceed the ``max flaps-extended speed" if you mishandle the later stages of the spin recovery and you don't want to damage the flaps.

Retracting the flaps may help with the spin recovery itself. Recall from [section 5.4.3](#) that the flaps effectively increase the washout of the wings. Washout ensures that the airplane will stall before it runs out of roll damping. (This produces a nice straight-ahead stall.) In the spin, though, when you have lost all vertical damping and roll damping, the washout doesn't help. The early stages of spin recovery are not like the early stages of stall entry.

Neutralizing the ailerons is usually a good idea for the simple reason that it is hard to think of anything better to do with them. Deflecting the ailerons effectively increases the angle of attack of one wingtip and decreases the angle of attack of the other wingtip. In a spin, the part of the wing where the ailerons are may (or may not) be in the stalled regime — so deflecting the ailerons to the left may (or may not) produce a paradoxical rolling moment to the right.

Depressing the rudder to oppose the spin is obviously a good thing to do.

Finally, you want to move the yoke to select zero angle of attack. In typical trainers, this means shoving the yoke all the way forward, but in other aircraft, especially aerobatic aircraft, all the way forward might select a large *negative* angle of attack. Shoving the yoke all the way forward in such a plane would likely convert the spin to an inverted spin — hardly an improvement. This is just one example of why you want to know and follow the spin recovery procedure for your specific airplane.

The relative significance of the rudder compared to the flippers in breaking the spin depends radically on the design of the airplane, the loading of the airplane, and on the spin mode, as discussed in [reference 6](#).

In normal non-spinning flight, you should apply smooth pressures to the controls. Spin recovery is the exception: it calls for brisk, mechanical *motions* of the controls, almost without regard to the pressures involved.

If you get into a spin in instrument conditions, you should rely primarily on the airspeed indicator and the rate-of-turn gyro. The inclinometer ball cannot be trusted; it is likely to be centrifuged away from the center of the airplane — giving an indication that depends on where the instrument is installed, telling you nothing about the direction of spin. The artificial horizon (attitude indicator) cannot be trusted since it may have tumbled. The rate-of-turn gyro is more trustworthy, since it is a rate gyro, not a free gyro; that is, it has no gimbals and cannot possibly suffer from gimbal lock.

Recovery from a so-called incipient spin (one that has just gotten started) is easier than from a well-developed spin. Normal-category single-engine¹⁸ certification requirements say that an airplane must be able to recover from a one-turn spin (or a 3-second spin, whichever takes longer) in not more than one additional turn. If you let the spin go on for several turns, you might progress from a steep spin to a flat spin. Recovery could take a lot longer — if it is possible at all.

If you load the airplane beyond the aft limit of the weight and balance envelope, even the incipient spin may be unrecoverable; see [section 6.1.1](#). Imperfect repairs to the wing, or slack in the control cables, could also impede spin recovery.

Finally, the spin is yet another reason why it is *NOT SAFE* to think of the yoke as simply the up/down control.¹⁹ In a spin you have a low airspeed and a high rate of descent. If you think of the yoke as the up/down control, you will be tempted to pull back on the yoke, which is exactly the wrong thing to do. On the other hand, if you think of the yoke as (primarily) the fast/slow control, you will realize that you need to push forward on the yoke, to solve the airspeed problem.

18.8 Don't Mess With Spins

It is quite impressive how well a samara works. A maple seed descends very slowly, riding the wind much better than a parachute of similar size and weight ever could. Flat spins can be extremely stable; a wing by itself *loves* to spin. That's why spins (and flat spins in particular) are so dangerous: it takes a lot of rudder force to persuade a wing to stop spinning.

Spins are extremely complex. Even designers and top-notch test pilots are routinely surprised by the behavior of spinning airplanes. Spin-test airplanes are equipped with cannon-powered spin-recovery parachutes on the airframe, and quick-release doors in view of the distinct possibility that the pilot will have to bail out. Tests are conducted at high altitude over absolutely unpopulated areas. Therefore please don't experiment with spinning a plane except exactly as approved by the manufacturer — one unrecoverable spin mode can ruin your whole day.

¹

This happens naturally on a rectangular wing; it can be enhanced by washout and other designers' tricks.

²

An even more direct method of adding energy to the boundary layer uses a jet of high-velocity air, as discussed in [section 18.4.2](#).

³

Of course, the VGs contribute *indirectly* to maintaining the health of the big bound vortex, since they help maintain attachment and therefore help create lots of circulation.

⁴

See [reference 17](#) for a nice discussion of golf balls, cricket balls, and boundary layers in general.

⁵

... just as having lots of water is the cause, but not the definition, of drowning — you can get very wet without drowning.

⁶

You can blow directly from your lips, but it's better to use a flexible straw or a thin piece of tubing, so that you can get a better view of what's happening. If you put a nozzle at the end of the tube, the jet will keep its shape better.

⁷

Ground pepper is a convenient source of suitable dust.

⁸

This won't be exactly atmospheric, since the local pressure has been affected by the wing.

⁹

Remember, lift is the force perpendicular to the flow and perpendicular to the surface.

¹⁰

Indeed, as long as the viscosity is not *exactly* zero, the smaller the viscosity, the greater the turbulence.

¹¹

Of course *some* energy is transferred, in the form of friction and induced drag, but this is very small, out of all proportion to the energy that the air parcel transfers from its own speed to pressure and back again.

¹²

Nonsensical things are often rather hard to calculate.

[13](#)

If you don't have a good model airfoil, start with a flat piece of cardboard.

[14](#)

This refers to ``departure from normal flight". It has nothing to do with takeoff or with a ``departure stall" which merely refers to a stall in the takeoff configuration.

[15](#)

Under present-day certification rules, the ailerons are required to work normally up to at least stalling angle of attack. However, some older airplanes were built under older rules. These planes, including many aerobatic aircraft, have much less washout, and therefore lose aileron effectiveness earlier. All planes lose effectiveness eventually. For simplicity, this section ignores washout.

[16](#)

See [section 19.4](#) for a discussion of the nature of centrifugal fields.

[17](#)

Recovery from a spiral dive is discussed in [section 6.2.4](#).

[18](#)

Multi-engine aircraft are not required to be recoverable from any sort of spin, incipient or otherwise.

[19](#)

This point is discussed in [chapter 7](#).

19 The Laws of Motion

There is no gravity.
The earth sucks.
— Physicist's bumper sticker

This chapter pulls together some basic physics ideas that are used in several places in the book.

We will pay special attention to rotary motion, since it is less familiar to most people than ordinary straight-line motion. Gyroscopes, in particular, behave very differently from ordinary non-spinning objects. It is amazing how strong the gyroscopic effects can be.

19.1 Straight-Line Motion

First, let's review the physical laws that govern straight-line motion. Although the main ideas go back to Galileo, we speak of Newton's laws, because he generalized the ideas and codified the laws.

The *first law of motion* states: "A body at rest tends to remain at rest, while a body in motion tends to remain in motion in a straight line unless it is subjected to an outside force". Although that may not sound like a very deep idea, it is one of the most revolutionary statements in the history of science. Before Galileo's time, people omitted frictional forces from their calculations. They considered friction "natural" and ubiquitous, not requiring explanation; if an object continued in steady motion, the force required to overcome friction was the only thing that required explanation. Galileo and Newton changed the viewpoint. Absence of friction is now considered the "natural" state, and frictional forces must be explained and accounted for just like any others.

The *second law of motion* says that if there is any change in the velocity of an object, the force required (F) is proportional to the mass (m) of the object, and proportional to the acceleration (a). (Acceleration is defined to be the rate-of-change of velocity.) In symbols,

$$F = m a \quad (19.1)$$

The following restatement of the second law is often useful: since *momentum* is defined to be mass times velocity, and since the mass is not supposed to be changing, we conclude that the force is equal to the rate-of-change of the momentum. To put it the other way, change in momentum is force times time.

The *third law of motion* says that if a force is applied to an object, an equal and opposite force must be applied somewhere else. This, too, can be restated in terms of momentum: if we impart a certain momentum to an object, we must impart an equal and opposite amount of momentum to something else. [1](#) This means that the total momentum of the world cannot change. This principle — conservation of momentum — is one of the most fundamental principles of physics, on a par with the conservation of energy discussed in [chapter 1](#).

* Force is Not Motion

As simple as these laws are, they are widely misunderstood. For example, there is a widespread misconception that an airplane in a steady climb requires increased upward force and a steady descent requires reduced upward force. [2](#) Remember, lift is a force, and any unbalanced force would cause an *acceleration*, not steady flight.

In unaccelerated flight (including steady climbs and steady descents), the upward forces (mainly lift) must balance the downward forces (mainly gravity). If the airplane had an unbalanced upward force, it would not climb at a steady rate — it would *accelerate* upwards with an ever-increasing vertical speed.

Of course, during the transition from level flight to a steady climb an unbalanced vertical force must be applied momentarily, but the force is rather small. A climb rate of 500 fpm corresponds to a vertical velocity component of only 5 knots, so there is not much momentum in the vertical direction. The kinetic energy involved in the vertical motion is negligible.

In any case, once a steady climb is established, all the forces are in balance.

19.2 Sitting in a Rotating Frame

Suppose Moe is riding on a turntable; that is, a large, flat, smooth, horizontal rotating disk. Moe has painted an X, Y grid on the turntable, so he can easily measure positions, velocities, and accelerations relative to the rotating coordinate system. His friend Joe is nearby, observing Moe's adventures and measuring things relative to a nonrotating coordinate system.

Moe quickly notices that Newton's first law does not apply in a rotating reference frame. A hockey puck initially at rest on the floor (anywhere except right at the center of rotation) will be centrifuged outward.

In Joe's nonrotating frame, Newton's laws do apply. In a nonrotating frame, there is no such thing as centrifugal acceleration.

Moe holds the puck in place with a rubber band (one end attached to the puck, the other end attached to the floor).

Joe sees the same rubber band and the same force. The elongation of the rubber band indicates the magnitude of the force.

Moe says the puck is not moving relative to his reference frame. The rubber band compensates for the centrifugal force.

Joe says that the puck's momentum is constantly changing due to the rotation. The rubber band provides the necessary force.

19.3 Moving in a Rotating Frame

We now consider what happens to an object that is *moving* relative to a rotating reference frame.

Suppose Moe has another hockey puck, which he attaches by means of a rubber band to a tiny tractor. He drives the tractor in some arbitrary way. We watch as the puck passes various marks (A, B , etc.) on the floor.

Moe sees the puck move from mark A to mark B . The marks obviously are not moving relative to his reference frame.

Joe agrees that the puck moves from mark A to mark B , but he must account for the fact that the marks themselves are moving.

So let's see what happens when Joe analyzes the compound motion, including both the motion of the marks and the motion of the puck relative to the marks. He actually sees *four* contributions:

1. If the puck is accelerating relative to Moe's rotating frame, Joe agrees and counts that as a contribution to the acceleration. Both observers agree on how much force this requires.
2. From Joe's point of view, mark A is not only moving; its velocity is changing. Changing this component of the puck's velocity requires a force. (From Moe's point of view, this is the force needed to oppose centrifugal acceleration, as discussed previously.)
3. The velocity of mark B is different from the velocity of mark A . As the puck is towed along the path from point A to point B , the rubber band must provide a force in order to change the velocity so the puck can ``keep up with the Joneses".
4. The velocity of the puck *relative* to the marks is also a velocity, and it must also rotate as the system rotates. This change in velocity also requires a force.

We can say a few words about each of these contributions from Moe's point of view:

1. This `` $F=ma$ " contribution is completely unsurprising. It is independent of position, independent of velocity, and independent of the frame's rotation rate.
2. The centrifugal contribution depends on position, but is independent of the velocity that Moe measures relative to his rotating reference frame. It is also independent of any acceleration created by Moe's tractor. It is proportional to the square of the frame's rotation rate.
3. This contribution is independent of position. It is proportional to the velocity that Moe measures, and is always perpendicular to that velocity. It is also proportional to the first power of the frame's rotation rate.
4. This contribution is also independent of position. It is also proportional to the velocity relative to the rotating frame, and is perpendicular to that velocity, and is proportional to the first power of the frame's rotation rate.

Contribution #3 is numerically equal to contribution #4. The total effect is just twice what you would get from either contribution separately. Together these two contributions are called the *Coriolis effect*.³

The Coriolis effect can be described as an *acceleration* (proportional to the object's velocity), and equivalently it can be described as a *force* (proportional to the object's momentum).

Let's consider a reference frame attached to an eastward-rotating planet, such as the earth. Near the north pole, the Coriolis acceleration is always toward your right, if you are facing forward along the direction of motion. Northward motion produces a Coriolis acceleration to the east; a very real westward force is necessary to oppose it if you want to follow a straight line painted on the earth. Eastward motion produces a Coriolis acceleration to the south; a very real northward force is necessary to oppose it.

The Coriolis argument only applies to motion in the plane of rotation. Momentum parallel to the axis is unaffected. In all cases the Coriolis acceleration lies in the plane of rotation and perpendicular to the motion.

Near the equator, we have to be careful, because the plane of rotation is not horizontal. In this region, eastward motion produces a Coriolis acceleration in the upward direction, while westward motion produces a Coriolis acceleration in the downward direction. In this region, north/south motions are parallel to the axis of rotation and produce no Coriolis effects.

To reiterate: The Coriolis effect and the centrifugal field are two separate contributions to the story. The Coriolis effect applies only to objects that are moving relative to the rotating reference frame. The centrifugal field affects all objects in the rotating frame, whether they are moving or not.

*** Magnitude of the Effect**

Suppose you are in an airplane, flying straight ahead at 120 knots along the shortest path between two points on the earth's surface. Because of the rotation of the earth, the airplane will be subject to a Coriolis acceleration of about 0.001G. This is too small to be noticeable.

Now suppose you and a friend are standing 60 feet apart, playing catch in the back of a cargo plane while it is performing a standard-rate turn (three degrees per second). If your friend throws you the ball at 60 mph, it will be subject to a horizontal Coriolis acceleration of more than a quarter G. That means the ball will be deflected sideways about 2 ½ feet before it gets to you — which is enough to be quite noticeable. In normal flying, though, we don't often throw things far enough to produce large Coriolis effects.

The wind, moving relative to the rotating earth, is subject to a Coriolis acceleration that is small but steady; the cumulative effect is tremendously important, as discussed in [section 20.1](#).

[19.4](#) Centrifuges with and without Gravity

[19.4.1](#) The Centrifugal Field is Real

An airplane in a turn, especially a steep turn, behaves like a centrifuge. There are profound analogies between centrifugal and gravitational fields:

The gravitational field at any given point is an acceleration. It acts on objects, producing a force in proportion to the object's mass.

The centrifugal field at any given point is also an acceleration. It, too, acts on objects, producing a force in proportion to the object's mass.

Strictly speaking, neither gravity nor centrifugity is a "force" field. Each is really an *acceleration* field. Of course there is a force involved, but it is always a force per unit mass, which is properly called an acceleration.

Einstein's principle of equivalence states that at any given point, the gravitational field is indistinguishable from an acceleration of the reference frame.⁴ In a freely-falling reference frame, such as a freely-orbiting space station, everything is weightless.

My laboratory is *not* a free-falling inertial frame. It is being shoved skyward as the earth pushes on its foundations. If you measure things relative to the laboratory walls, you will observe gravitational accelerations.

Similarly, the cabin of a centrifuge is clearly not an inertial frame. If you measure things relative to the cabin, you will observe centrifugal accelerations.

From a modern-physics point of view, both gravity and centrifugity emerge as consequences of working in an accelerated frame. There is nothing wrong with doing so, provided the work is done carefully. Accounting for centrifugal effects is not much trickier than accounting for gravitational effects. When people think this can't be done, it is just because they don't know how to do it.

For a ground-bound observer analyzing the flight of an airplane, it may be convenient to use a reference frame where gravity exists and centrifugity does not. However, the pilot and passengers usually find it convenient to use a frame that includes both gravity and centrifugity.

The centrifugal field is not crude or informal or magical. (The problem with magic is that it can explain false things just as easily as true things.) Like the gravitational field, it is a precise way of accounting for what happens when you work in a non-freely-falling reference frame.

19.4.2 Centrifuge

To get a better understanding of the balance of forces in a turning and/or slipping airplane, consider the centrifuge shown in [figure 19.1](#). For the moment we will neglect the effects of gravity; imagine this centrifuge is operating in the weightless environment of a space station. We are riding inside the centrifuge cabin, which is shown in red. We have a supply of green tennis balls. At point A (the southernmost point of our path) we drop a tennis ball, whereupon it flies off as a free particle. Our centrifuge continues to follow its circular path.



[Figure 19.1](#): An Object Departing a Centrifuge

Case 1a: Consider the point of view of a bystander (not riding in the centrifuge). The dropped tennis ball moves in a straight line, according to the first law of motion. Contrary to a common misconception, the bystander does *not* see the ball fly radially away from the center of the centrifuge. It just continues with the purely eastward velocity it had at point A, moving tangentially.

Case 1b: Consider our point of view as we ride in the centrifuge. At point A, the tennis ball has no velocity relative to us. For the first instant, it moves along with us, but then gradually it starts moving away. We *do* see the ball accelerate away in the purely radial direction. The tennis ball — like everything else in or near the centrifuge — seems to be subjected to a centrifugal acceleration field.

Einstein's principle of equivalence guarantees that our viewpoint and the bystander's viewpoint are equally valid. The bystander says that the centrifuge cabin and its occupants accelerate away from the freely moving tennis ball, while we say that the tennis ball accelerates away from us under the influence of the centrifugal field.

There is one pitfall that must be avoided: you can't freely mix the two viewpoints. It would be a complete fallacy for the bystander to say ``gee, the folks in the cabin told me the tennis ball accelerated outward; therefore it must move to the south starting from point A''. In fact, the free-flying ball does not accelerate relative to the bystander. It will not wind up even one millimeter south of point A. It will indeed wind up south of our centrifuge cabin, but only because we have peeled off to the north.

Case 2a: Consider from the bystander's point of view what happens to a ball that has *not* been released, but is just sitting on a seat in the centrifuge. The bystander sees the ball subjected to an unbalanced force, causing it to move in a non-straight path relative to the earth.

Case 2b: Consider the seated ball from the centrifuge-riders' point of view. The force on the ball exerted by the seat is just enough to cancel the force due to centrifugal acceleration, so the forces are in balance and the ball does not move.

When analyzing unsteady motion, or when trying to calculate the motion of the centrifuge itself, it is often simpler to analyze everything from the bystander's point of view, in which the centrifugal field will not appear. On the other hand, in a steady turn, is often easy and natural to use the centrifuge-riders' point of view; in which all objects will be subject to centrifugal accelerations.

19.4.3 Centrifuge and Gravity

Now that we understand the basic idea, let's see what happens when our centrifuge operates in the normal gravitational field of the earth. This is shown in [figure 19.2](#). When the tennis ball departs the centrifuge, it once again travels in a purely easterly direction, but this time it also accelerates downward under the influence of gravity.



[Figure 19.2](#): An Object Departing a Centrifuge, with Gravity


Once again, from inside the cabin we observe that the tennis ball initially accelerates away in the direction exactly away from the pivot of the centrifuge. This is no coincidence; it is because the only *difference* between our motion and the free-particle motion comes from the force in the cable that attaches us to the pivot.

(The foregoing applies only to the *initial* acceleration of the dropped ball. As soon as it picks up an appreciable velocity relative to us, we need to account for Coriolis acceleration as well as centrifugal acceleration.)

Remember, the equivalence principle says that at each point in space, a gravitational field is *indistinguishable* from an accelerated reference frame. Therefore we need not know or care whether the tennis ball moves away from us because we are being accelerated, or because there is a gravitating planet in the vicinity, or both.

19.5 Centrifugal Effects in a Turning Airplane

Let's examine the forces felt by the pilot in a turning airplane. We start with a coordinated turn, as shown in [figure 19.3](#).




[Figure 19.3](#): Airplane in a Coordinated Turn

In figures such as this, whenever I am analysing things using the pilot's point of view, the figure will include a rectangular ``frame'' with a little stick figure (the observer) standing in the corner. It is important to carefully specify what frame is being used, because even simple questions like ``which way is down'' have answers that depend on which observer you ask. In particular, I define *N-down* (Newtonian down) to mean the direction straight down toward the center of the earth. In contrast, I define *E-down* (effective down, or Einsteinian down) to be the direction in which a free particle departs if you drop it. In a turning airplane, the two directions are not the same.

Using your inner ear, the seat of your pants, and/or the inclinometer ball, you can tell which way is E-down. Using the natural horizon and/or the artificial horizon, you can tell which way is N-down.

In [figure 19.3](#), assume the airplane's mass is one ton. Real gravity exerts a force of one ton, straight down toward the center of the earth. The airplane is in a 45° bank, so there is one ton of centrifugal force, sideways, parallel to the earth's horizon. All in all, the wings are producing 1.41 tons of lift, angled as shown in the figure.




[Figure 19.4](#): Inclinometer in a Coordinated Turn

[Figure 19.4](#) shows the same situation, and analyzes the forces on the inclinometer ball. Real gravity exerts a downward force on the ball, and centrifugity exerts a sideways force. The tubular race that contains the ball exerts a force perpendicular to the wall of the race (whereas the ball is free to roll in the direction along the race). The race-force balances the other forces when the ball is in the middle, confirming that this is a coordinated turn.

As a contrast, consider [figure 19.5](#), which shows what happens if you make a boat turn. (Boat turns are discussed in [section 8.10](#).)

Because the airplane is turning, it and everything in it will be subjected to a centrifugal acceleration (according to the viewpoint of the centrifuge riders).



[Figure 19.5](#): Airplane in a Boat Turn

[Figure 19.6](#) shows how the inclinometer ball responds to a boat turn. Gravity still exerts a force on the ball, straight down. Centrifugity exerts a force sideways toward the outside of the turn. The ball is subject to a force of constraint, perpendicular to the walls of the race. (It is free to roll in the other direction.) The only place in the race where this constraint is in a direction to balance the other forces is shown in the figure. The ball has been ``centrifuged" toward the outside of the turn. This is a quantitative indication that the E-down direction is not perpendicular to the wings, and some force other than wing-lift is acting on the plane.




Figure 19.6: Inclinator in a Boat Turn




Figure 19.7: Airplane in a Nonturning Slip




Figure 19.8: Inclinator in a Nonturning Slip

Finally, [figure 19.7](#) shows the forces on the airplane in an ordinary nonturning slip. The right rudder pedal is depressed, and the port wing has been lowered just enough that the horizontal component of lift cancels the horizontal force due to the crossflow over the fuselage. The airplane is not turning. The bystander and the riders agree that there is no centrifugal field. [Figure 19.8](#) shows the forces on the inclinometer ball in the same situation.

19.6 Angles and Axes

19.6.1 Axes: Yaw, Pitch, and Roll

The standard definitions of the yaw, pitch, and roll axes are shown in [figure 19.9](#).




Figure 19.9: Yaw, Pitch, and Roll Axes

Many older books refer to the pitch axis as the *lateral axis* and the roll axis as the *longitudinal axis*, which are sensible enough, but then they refer to pitch-axis stability as *longitudinal stability* and roll-axis stability as *lateral stability* — which seems completely reversed and is confusing at best. [Reference 22](#) calls the yaw axis the *normal axis*, since it is normal (i.e. perpendicular) to the other axes — but that isn't very helpful since *all* the axes are normal to all the others. Other references call the yaw axis the *vertical axis*, but that is very confusing since if the bank attitude or pitch attitude is not level, the yaw axis will not be vertical.

Obviously the new terminology is an improvement. The situation is summarized in the following table.

This Book	Older Terminology
pitch axis	lateral axis
pitch-axis stability	longitudinal stability
roll axis	longitudinal axis
roll-axis stability	lateral stability
yaw axis	vertical axis
yaw-axis stability	directional stability

19.6.2 Attitude: Heading, Pitch, Bank

The term *attitude* describes the orientation of the airplane's axes relative to the earth. Attitude is specified in terms of three angles: the heading, the pitch attitude, and the bank attitude.

To place the airplane in a specified attitude, start with the airplane in level flight (i.e. X axis and Y axis horizontal) and with the airplane pointed north. Then

- Rotate the airplane around the yaw axis by the specified *heading* angle. A positive angle specifies a clockwise rotation as seen from above, so that a heading of 090 degrees corresponds to pointing east and a heading of 180 degrees corresponds to pointing south.
- Rotate the airplane around the pitch axis by the specified *pitch attitude* angle. A positive angle specifies a nose-up attitude.
- Rotate the airplane around the roll axis by the specified *bank attitude* angle. A positive angle corresponds to clockwise as seen from the rear.

As discussed below ([section 19.6.4](#)), it is important to perform these rotations in the order specified.

We have just seen how, given a set of angles, we can put the airplane into a specified attitude. We now consider the reverse question: given an airplane in some attitude, how do we determine the angles that describe that attitude?

Answer: just figure out what it would take to return the airplane to level northbound flight. The rotations must be undone in the reverse of the standard order:

- First, rotate the aircraft around the roll axis until the wings are level. This determines the bank attitude.
- Second, rotate the aircraft around the pitch axis until the X axis is level. This determines the pitch attitude. Note that this rotation is not performed around the original pitch axis, but rather around the new pitch axis, which is horizontal as a consequence of the previous step.
- Finally, rotate the aircraft around the (new) yaw axis until the nose is pointing north. This determines the heading.

19.6.3 Angle Terminology

The following table summarizes the various nouns and verbs that apply to angles and motions around the three axes.

	Z axis	Y axis	X axis
Motion	it yaws	it pitches	it rolls
Angle	the heading	the pitch attitude	the bank attitude

There are a few more fine points of *X*-axis terminology:

- Saying that the airplane is "banking" or "in a bank" refers to a definite bank attitude.
- In contrast, saying that the airplane is "rolling" or "in a roll" refers to a definite rate of rotation, i.e. a changing bank angle.
- *Pitch angle* usually means the same thing as pitch attitude.
- *Bank angle* usually means the same thing as bank attitude.
- The term *roll angle* is a rarely-used synonym for bank angle.

* Other Angles

To define the *angle of attack* of the fuselage, take the direction of flight (or its reciprocal, the relative wind) and project it onto the *XZ* plane. The angle of attack is the angle between this projection and the *X* axis or some other convenient reference.

To define the *slip angle*, take the direction of flight (or the relative wind) and project it onto the *XY* plane. The slip angle is the angle between this projection and the *X* axis.

Some aerodynamics texts use the term *sideslip angle*, which is synonymous with slip angle. This may be confusing because some pilot-training books try to draw a distinction between a *forward slip* and a *side slip*, even though the difference is imaginary, as is discussed in conjunction with [figure 11.1](#).

19.6.4 Yaw Does Not Commute with Pitch

It is a fundamental fact of geometry that the result of a sequence of rotations *depends on the order* in which the rotations are performed.

Note that for a sequence of ordinary non-rotational movements, the ordering does not matter. That is, suppose I have two small objects that start out at the same place. I move one object move two feet north, and then three feet west. I move the other object the same distances in the other order: three feet west and then two feet north. Assuming there are no obstructions, both objects will arrive at the same destination. The ordering of the movements does not matter.

However, angles don't play by the same rules as distances. For instance, it is possible to change the yaw angle (i.e., the heading) by 37 degrees (or any other amount) without ever yawing the plane. That is:

- You can pull up into aerobatic vertical flight, use the ailerons to rotate 37 degrees around the now-vertical *roll axis*, and then push back to level flight.
- Another way to do the same thing is to roll into a ninety-degree bank, pull on the yoke to rotate around the now-vertical *pitch axis* by 37 degrees, and then roll back to wings-level attitude.

If the aircraft (and its occupants) can tolerate heavy *G* loads, such maneuvers are perfectly fine ways to make tight turns at high airspeed.

In non-aerobatic flight, a less-extreme statement applies: a rotation in a purely horizontal plane is *not* a pure yaw when the aircraft is not in a level attitude. For instance, suppose you are in level flight, steadily turning to the left. This is, of course, a turn around an axis that runs vertically through the center of the earth. Further suppose that you have a nose-up pitch attitude, while still maintaining a level flight path. This means that the airplane's yaw axis (*Z* axis) is not vertical. You could, in principle, perform the required heading change by pitching down to level pitch attitude, performing a pure yaw, and then pitching back up, but since rotations are not commutative this is not equivalent to maintaining your pitch attitude and performing a pure yaw. Performing the required change of heading without pitching down requires *mostly* pure leftward yaw, but involves some rightward rotation around the airplane's roll axis as well.

The analysis in the previous paragraph is 100% accurate, but completely irrelevant when you are piloting the airplane. [5](#) In particular, the rotation around the roll axis does not raise the angle of attack on one wingtip or lower it on the other. Arguing about whether the heading change is a pure yaw or a yaw plus roll is almost like arguing about whether a glass of water is half full or half empty — the physics is the same. In this case the physics is simple: the inside (left) wing follows a horizontal circular path, while the outside (right) wing follows a slightly longer horizontal path around a larger circle.

It is easy to see why that is so: The turn requires a rotation around a vertical axis. Such a rotation moves the wingtips (and everything else) in purely horizontal directions. As long as the airplane's center-of-mass motion is also horizontal, the rotation can only change the speeds, not the angles, of the airflow.

Now, things get more interesting when the direction of flight is not horizontal. Therefore let us consider a new example in which you are climbing while turning. That means your flight path is inclined above the horizontal. As before, you are turning to the left at a steady rate.

In any halfway-reasonable situation, the yaw axis will be nearly perpendicular to the direction of flight. Having it not exactly perpendicular is just a distraction from the present topic, so I hereby *define* a new ``yaw-like" axis exactly perpendicular the direction of flight. It is also perpendicular to the pitch axis (which remains parallel to the wingtip-to-wingtip direction). The third axis is a new ``roll-like" axis perpendicular to the other two. If we assume zero slip angle, it will be aligned with the direction of flight.

As the airplane flies from point to point along its curving path, its heading must change. This is a rotation around a purely vertical axis. In climbing flight, the yaw-like axis is not vertical, so the turn is not pure yaw. The turn moves the inside wingtip horizontally backwards, relative to where it would be if there were no heading change. In contrast, a pure yaw-like rotation would have moved the wing back *and down*. Therefore we need not just leftward yaw-like rotation but also some rightward roll-like rotation to keep the wingtip moving along the actual flight path.

This roll-like motion means that (other things being equal) the inside wingtip would fly at a lower angle of attack during a climbing turn. Less lift would be produced. You need to deflect the ailerons to the outside to compensate.

Note that I said less lift ``would be" produced, not ``is" produced. Because you have compensated with the ailerons, both wings are in fact producing the same amount of lift, as they should. Remember that this is a *steady* turn, so no force is required to maintain the steady roll rate. (Remember, according to Newton's laws, an unbalanced force would create an *acceleration* around the roll axis, which is not what is happening here.) There are widespread misconceptions about this. Because of the roll-like motion, the air will arrive at the two wings from two different directions. You will have to deflect the ailerons, not in order to create a magnitude of lift, but rather to *keep from* creating a different magnitude of lift.

Although the magnitude of the lift remains the same, the direction of the lift will be twisted, as discussed in [section 8.8.4](#); see in particular [figure 8.7](#). You will need to deflect the rudder to overcome the resulting yawing moment. This will be in the usual direction: right rudder in proportion to right aileron deflection, and left rudder in proportion to left aileron deflection.

In a climbing turn, the differential relative wind combines with the differential wingtip velocity to create a large overbanking tendency. In an ordinary descending turn, the relative wind effect tends to oppose the velocity effect. In a spin, the differential relative wind is a key ingredient, as discussed in [section 18.6.1](#), including [figure 18.6](#). Also, [section 9.7](#) analyzes climbing and descending turns in slightly different words and gives a numerical example.

[19.6.5](#) Yaw Does Not Commute with Bank

As stated above, a rotation in a purely horizontal plane is not a pure yaw when the aircraft is not in a level attitude. In the previous section we considered the consequences of a non-level pitch attitude, but the same logic applies to a non-level *bank* attitude. The latter case is in some sense more significant, since although not all turns involve a non-level pitch attitude, they almost always involve a bank.

You could perform the required rotation by rolling to a level attitude, performing a pure yaw, and then rolling back to the banked attitude. This is not equivalent to performing a pure yaw while maintaining constant bank. For modest bank angles, the constant-bank maneuver is *mostly* pure yaw, but involves some rotation around the pitch axis as well. Because of this rotation around the pitch axis, the relative wind hits the wing and the tail at slightly different angles. You will need to pull back on the yoke slightly to compensate. This pull is in addition to whatever pull you might use for controlling airspeed during the turn. You can see that the two phenomena are definitely distinct, by the following argument: suppose that you maintain constant angle of attack during the turn, so that the required load factor is produced by increased airspeed not increased angle of attack. You would still need to pull back a little bit, to overcome the noncommutativity.

[19.7](#) Torque and Moment

Just as Newton's first law says that to start an object moving you have to apply a force, there is a

corresponding law that says to start an object turning you need to apply a torque.

You may have heard of the word "torque" in conjunction with left-turning tendency on takeoff, and you may have heard of the word "moment" in conjunction with weight & balance problems. In a deep sense, the two words are two names for the same thing. In particular,

- A rolling moment is a torque around the roll axis.
- A pitching moment is a torque around the pitch axis.
- A yawing moment is a torque around the yaw axis.

A familiar example: fuel and cargo cause a pitching moment, depending on how far forward or aft they are loaded. By the same token, they will cause a rolling moment if they are loaded asymmetrically left or right.

Another familiar example: gyroscopic effects can cause torque around the yaw axis. By the same token, they can cause torque around the pitch axis as well.

I will use the terms "torque" and "moment" interchangeably.

Torque is not the same as force. Of the two, force is the more familiar concept. If you attach a string to an object and pull, the object is subjected to a force in the direction of the string. Force is measured in pounds or newtons.

To apply a torque, you need a force and a lever-arm. The amount of torque is given by the following formula:

$$\text{torque} = \text{force} \times \text{distance} \quad (19.2)$$

where the distance is a vector representing the separation between the pivot-point⁶ and the point where the force is applied. In this formula, we are using the vector *cross product*,⁷ indicating that we only count the component of the distance perpendicular to the force. This quantity (the component of the distance perpendicular to the force) is called the *lever arm*, which makes it convenient to say that the torque is the force times the lever arm.

Torque is measured not in pounds but in footpounds (that is, feet times pounds); the corresponding metric unit is newtonmeters. ⁸

Figure 19.10 shows a situation where all the forces and torques are in balance. On the right side of the bar, a group of three springs is exerting a force of 30 pounds. On the left side of the bar, there is a group of two springs (exerting a force of 20 pounds) and a single spring (exerting a force of 10 pounds). Since the total leftward force equals the total rightward force, the forces are in balance.




Figure 19.10: Forces and Torques in Balance

To show that the torques are in balance requires a separate check. Let's choose the point marked "x" as our pivot point. The rightward force produces no torque, because it is attached right at the pivot point — it has a zero-length lever arm. The group of two springs produces a counterclockwise torque, and the single spring produces a clockwise torque of the same magnitude, because even though it has half as much force it has twice the lever arm. The torques cancel. The system is in equilibrium.




Figure 19.11: Forces in Balance but Torques NOT in Balance

Figure 19.11 shows a different situation. The forces are in balance (20 pounds to the right, 20 pounds total to the left) but the torques are not in balance. One of the left-pulling springs has twice the lever arm, producing a net clockwise torque. If you tried to set up a system like this, the bar would immediately start turning clockwise. The system is out of equilibrium.

Torque is a vector; that is, it has direction as well as magnitude. A torque *around* a given axis is represented

by a torque vector *along* that axis. In particular, the clockwise torque in [figure 19.11](#) is represented by a vector perpendicular to the paper, directed away from the reader.

19.8 Angular Momentum

The notion of *angular momentum* is the key to really understanding rotating objects.

Angular momentum is related to ordinary straight-line momentum in the same way that torque is related to ordinary straight-line force. Here is a summary of the correspondences:

Straight-line concept

Force

Momentum

The ordinary momentum of a system won't change unless a force is applied.

The change in straight-line momentum is the force times the time.

Angular concept


Torque (equals force times lever arm)

Angular momentum (equals ordinary momentum times lever arm)

The angular momentum of a system won't change unless a torque is applied.

The change in angular momentum is the torque times the time.

When I give lectures, I illustrate conservation of angular momentum using a demo you can easily set up for yourself. As illustrated in [figure 19.12](#), tie some kite string to a small bean-bag and swing it in a circle. When you pull on the free end of the string (reducing the radius of the circle) the bean-bag speeds up. When you let out the string (increasing the radius of the circle) the bean-bag slows down.⁹



[Figure 19.12](#): Conservation of Angular Momentum

In typical textbooks, conservation of angular momentum is exemplified by spinning ice skaters, but I find it easier to travel with a bean-bag (rather than an ice skater) in my luggage.

In the demonstration, there are some minor torques due to friction that will eventually slow down the bean-bag whether or not you shorten or lengthen the string, but if you perform the experiment quickly enough the torques can be neglected, and the angular momentum of the system is more or less constant. Therefore, if you decrease the lever arm by a factor of N , the straight-line momentum must increase by a factor of N (since their product cannot change).¹⁰

Since the tangential velocity increases by a factor of N , and the radius decreases by a factor of N , the rate of turn (degrees per second) increases by a factor of N squared.

The energy of the system also increases by a factor of N squared. You can feel that you added energy to the system when you pull on the string, pulling against tension.

So far we have analyzed the situation from the point of view of a bystander in a non-rotating reference frame. You can reach the same conclusion by analyzing the situation in the rotating reference frame, as would apply to an ant riding on the bean-bag. The ant would say that as the string is pulled in, the bean-bag accelerates sideways because of the Coriolis effect, as discussed in [section 19.3](#).

Conservation of angular momentum applies to airplanes as well as bean-bags. For instance, consider an airplane in a flat spin, as discussed in [section 18.6.4](#). In order to recover from the spin, you need to push the nose down. This means whatever mass is in the nose and tail will move closer to the axis of rotation. The angular momentum of the airplane doesn't change (in the short run), so the rotation will speed up (in the short run). More rotation may seem like the opposite of what you wanted, but remember you are trying to get rid of angular momentum, not just angular rate. You should persevere and force the nose down. Then the aerodynamic forces (or, rather, torques) will carry angular momentum out of the system and the rotation will decrease.

Angular momentum is a vector; that is, it has direction as well as magnitude. The angular momentum vector points more-or-less¹¹ along the axis of rotation. In particular, the angular momentum of an aircraft propeller

(rotating clockwise as viewed from behind) is directed *forward* along the axis of rotation.

19.9 Gyroscopes

19.9.1 Precession

For any normal object (such as a book) if you apply a force in a given direction, it will respond with motion in that direction. People are so accustomed to this behavior that they lose sight of the fact that force and motion are not exactly the same thing, and they don't always go together.

In particular, for a gyroscope, if you apply a torque around one axis, it will respond with motion around a different axis. When I give my ``See How It Flies'' lectures, I carry around a bicycle wheel with handles, as shown in [figure 19.13](#). The indicated direction of spin corresponds to a normal American engine and propeller, if the nose of the airplane is toward the left side of the diagram.



[Figure 19.13](#): Bicycle Wheel with Handles

To demonstrate the remarkable behavior of a gyroscope, I stand behind the ``propeller'' (on the right side of the diagram) and support its weight by lifting the rear handle only. The force of gravity acts on the center of the system, so there is a pure nose-down / tail-up pitching moment. If this were a normal, non-spinning object, it would respond by pitching in the obvious way, but the gyroscope actually responds with a pure yawing motion. I have to turn around and around to my left to stay behind the wheel.

It is really quite amazing that the wheel does not pitch down. Even though I am applying a pitching force, the wheel doesn't pitch down; it just yaws around and around.



[Figure 19.14](#): Gyroscopic Precession

This phenomenon, where a gyro responds to a torque about one axis with a motion about another axis, is called *gyroscopic precession*.


Precession is often quite noticeable at the point where a taildragger raises the tail, early in the takeoff roll. If the airplane were an ordinary non-spinning object, you could raise the tail using the flippers alone. But note that airflow over the flippers does not actually dictate the *motion* of the airplane; it just produces a *force* and a torque around the pitch axis. For a gyroscope, a torque around the pitch axis produces a motion around the yaw axis. If you try to raise the tail of a real airplane using flippers alone, it will yaw to the left because of precession.

To get a gyroscope to actually *move* around the pitch axis, you need to apply a torque around the yaw axis — using the rudder.

Of course, an airplane has some plain old mass in addition to its gyroscopic properties. In order to lift this ordinary mass you need to use the flippers. Therefore, the tail-raising maneuver requires both flippers and rudder — flippers to change the pitch of the ordinary mass, and rudder to change the pitch of the gyroscope.

19.9.2 The Torque is 90 Degrees Ahead of the Motion

Now that you've got the idea that a torque around the pitch axis causes a motion around the yaw axis, you may be wondering how to figure out the *direction*. Will it be yaw to the right, or yaw to the left? There is a simple rule: the axis of the applied torque will be 90 degrees ahead of the axis of the motion.¹² An application of this rule is depicted in [figure 19.15](#).



[Figure 19.15](#): The Applied Torque is 90 Degrees Ahead of The Motion

If you apply forward pressure on the yoke, it applies a torque around the pitch axis in the direction that your fingers curl in [figure 19.15\(a\)](#).¹³

Now rotate your hand 90 degrees in the direction that the propeller is spinning. Your fingers now specify the direction of motion, as shown in [figure 19.15\(b\)](#). The nose will move to the left.

Conclusion: forward pressure on the yoke makes the nose move to the left.


As another application of this rule, you can convince yourself that right rudder deflection will cause a nose-down / tail-up pitching motion.

The applied torque is 90 degrees ahead of the motion.

Superficially, this gyroscope law is reminiscent of the Coriolis effect (force in one direction, motion in a perpendicular direction) but they do not represent the same physics. The Coriolis law only applies to objects that are moving relative to a rotating observer. In contrast, the gyroscope law applies to a stationary observer; no parts of the gyroscope are moving relative to each other.

[19.9.3](#) Angular Momentum is the Key Idea

Precession, and all the other gyro effects discussed in this section, are direct consequences of the angular momentum laws discussed in the previous section.



[Figure 19.16](#): Angular Momentum Explains Precession

[Figure 19.16](#) is the same as [figure 19.14](#), but with the torque and angular momentum vectors drawn in. The torque around the pitch axis is represented by a vector along the pitch axis — to my left if I am standing behind the wheel. The angular momentum is represented by a vector along the axis of rotation, directed away from me since the wheel is spinning clockwise as seen from behind. The change in angular momentum is (according to the laws of motion) equal to the torque times the time.

The figure shows what happens during one unit of time (one second). The new angular momentum is the vector sum of the old angular momentum, plus the torque times the time. Since we are considering *one* second, the change vector is equal in magnitude to the torque vector.

During subsequent seconds, the torque vector will be a new direction because the whole system has rotated. The successive changes will pull the angular momentum vector around in a horizontal circle.

Gyroscopic effects only occur when the orientation of the gyro axis is changing. You can take a gyro and transport it north/south, east/west, or up/down, without causing any precession, as long as the gyro axis remains pointed in the same direction. You can even roll an airplane without seeing gyroscopic effects, since the roll leaves the gyro axis pointing in the same direction.

If you have a lightweight airframe and a heavy, rapidly spinning propeller, watch out: the flippers will cause

yawing motion and the rudder will cause pitching motion.

If you want to make a gyro change orientation quickly, it will take more torque than doing it slowly.

19.9.4 Inertial Platform

We now consider what happens when a gyro is *not* subjected to any large torques.

Suppose we support a gyroscope on gimbals. The gimbals support its weight but do not transmit any torques to it, even if the airplane to which the gimbals are mounted is turning. We call this a *free gyro* since it is free to not turn when the airplane turns.

Even though the gyro is small, it has a huge amount of angular momentum, because it is spinning so rapidly. Any small torque applied to the gyro (because of inevitable imperfections in the gimbals) will, over time, change the angular momentum — but over reasonably short times the change is negligible compared to the total.

In such a situation, the gyro will tend to maintain fixed orientation in space. We say that the gyro is an *inertial platform* with respect to rotations.¹⁴ Other books say the gyro exhibits *rigidity in space* but that expression seems a bit odd to me.

19.10 Gyroscopic Instruments

We now discuss the principles of operation of the three main gyroscopic instruments: artificial horizon gyro (attitude indicator), the directional gyro (heading indicator), and the rate of turn gyro (turn needle or turn coordinator).

19.10.1 Heading Indicator

The *directional gyro* is a free gyro. It establishes an inertial platform.

The gyro spins in some vertical plane; that is, its angular momentum vector points in some arbitrary horizontal direction. A system of gears measures the angle that the angular momentum vector makes in the *XY* plane¹⁵ and displays it to the pilot. The trick is to measure the angle and support the gyro while minimizing the accidental torques on it. Imperfections in the mechanism cause the gyro to precess; therefore, every so often the heading indication must be corrected, typically by reference to a magnetic compass.

19.10.2 Artificial Horizon

The *artificial horizon* (also known as the *attitude indicator*) is another free gyro. This gyro's plane of rotation is horizontal; that is, its angular momentum vector is vertical. A mechanical linkage measures the angle that this vector makes in the *YZ* (bank) and *XZ* (pitch) planes, and displays it to the pilot.

It is instructive to compare the horizon gyro (which tells you which way is ``down") with the inclinometer ball or a plumb-bob on a string (which has a different notion of which way is ``down"). The distinction is that the plumb-bob tells you which way is E-down, while the gyro is designed to tell you which way is N-down (toward the center of the earth). Whenever the airplane is being accelerated (e.g. during the takeoff roll or during a turn), the two directions are quite different. As seen in [figure 19.17](#), during a turn the E-down vector gets centrifuged to the outside of the turn; the N-down vector always points to the center of the earth.

[Figure 19.17](#): E-Down versus N-down During a Turn

To a first approximation, the horizon gyro works just by *remembering* which way is N-down. However, no gyro can remember anything forever, so the instrument contains an ``erecting mechanism" that makes continual small adjustments. You would like it to align the gyro axis with N-down — but the mechanism doesn't know which way is N-down! It knows which way is E-down (the same way the plumb-bob does), but according to Einstein's principle of equivalence, it cannot possibly know what components of E-down are due

to gravity and what components are due to acceleration. The erecting mechanism does, in fact, continually nudge the gyro axis toward E-down, but the result is a good approximation to N-down, for the following reason: if you *average* the E-down vectors over an entire turn, they average out to N-down. As you can see in the figure:

- sometimes E-down points a little to the north of N-down,
- sometimes E-down points a little to the west of N-down,
- sometimes E-down points a little to the south of N-down,
- sometimes E-down points a little to the east of N-down.

If you average the discrepancies over an entire turn, they cancel.

* Artificial Horizon Errors

Of course, if you only make *half* a turn, the discrepancies don't average to zero, and the attitude indicator will be slightly inaccurate for a while. Analogous errors occur during takeoff, because the gyro's estimate of "down" gets dragged backwards by the acceleration, so the artificial horizon will be a little bit below the true forward horizon for a while thereafter. The averaging time for a typical instrument is about five minutes. When the instrument gets old, its memory gets shorter because the gyro isn't spinning as fast, so the errors become larger.

19.10.3 Rate-of-Turn Gyro

There are two slightly different types of rate-of-turn gyro: (a) the rate-of-turn needle, and (b) the turn coordinator.

In both cases, the gyro is not free; it is a *rate gyro*. That is, its axis of rotation is more-or-less firmly attached to the airplane. It does not have gimbals. It is forced to change orientation when the airplane yaws, i.e. rotates around its *Z* axis.¹⁶ The instrument measures how much torque is required to re-orient the gyro.

Sometimes the *rate-of-turn needle* is built with its spin axis running parallel to the *Y* axis, in which case the airplane's yawing motion requires a torque around the *X* axis. Other models have the spin axis in the *X* direction, in which case yaw requires a torque around the *Y* axis. In principle, you could use any pair of axes that lie in the *XY* plane.¹⁷

The required torque is proportional to (a) the rate of change of orientation, and (b) the angular momentum of the gyro. Therefore an accurate rate-of-turn gyro must spin at exactly the right speed, not too fast or too slow. (This is in contrast to the directional gyro and the artificial horizon gyro, which just have to spin "fast enough".)

Many rate gyros incorporate a sneaky trick. They spin around the *Y* axis, with the top of the gyro spinning toward the rear. That is, the angular momentum vector points in the *+Y* direction. They also use a spring that is weak enough to allow the gyro to precess a little around the *X* axis. In a turn to the left, precession will tilt the gyro a little to the right. That means that during a turn, the gyro's tilt compensates for the airplane's bank, leaving the gyro somewhat more aligned with the earth's vertical axis. The goal, apparently, is to create an instrument that more nearly indicates heading change (relative to the earth's vertical axis) rather than simply rotation around the airplane's *Z* axis. Since the relationship between bank angle and rate of turn depends on airspeed, load factor, et cetera, this trick can't possibly achieve the goal except under special conditions.

The *turn coordinator* is very similar to the rate-of-turn needle. It displays a miniature airplane instead of a needle. The key operational difference is that it is slightly sensitive to rate of *roll* as well as rate of heading change. To create such an instrument, all you have to do is take a rate-of-turn instrument, tilt the mechanism nose-up by 20 or 30 degrees, and change the display.

The advantage of a turn coordinator is that it helps you anticipate what actions you need to take. That is, if the airplane has its wings level but is *rolling* to the right, it will probably be turning to the right pretty soon, so you might want to apply some aileron deflection. The disadvantage has to do with turbulence. Choppy air oftentimes causes the airplane to roll continually left and right. The roll rate can be significant, even if the bank angle never gets very large. The chop has relatively little effect on the yaw axis. In such conditions a plain old rate-of-turn needle gives a more stable indication than a turn coordinator does.

It is rather unfortunate that the display on a turn coordinator is a miniature airplane that banks left and right. This leads some people to assume, incorrectly, that the instrument indicates bank angle, which it most definitely does not.

1 This is traditionally expressed in terms of an "action" and an "equal and opposite reaction", but the meaning of those words has drifted over the centuries. Momentum is the modern term.

2 Troublemakers sometimes point out that lift actually *is* slightly reduced in a steady descent, since part of the weight is being supported by drag. To this I retort: (a) this is an obscure technicality, based on details of the definitions of the four forces (as given in [section 4.1](#)); (b) the magnitude of the reduction is negligible in ordinary flying, (c) the lift is reduced for climbs as well as descents — so this technicality certainly does not explain the motion, and (d) when we consider the total upward force, there is no reduction.

3 It is easy to find hand-waving explanations of the Coriolis effect that overlook one or the other of the two contributions, and are therefore off by a factor of two. Beware.

4 If you consider multiple widely-separated points, you can distinguish gravity versus centrifugity versus straight-line acceleration by checking for nonuniformities in the fields. However, an airplane is so small compared to the planet, and so small compared to its turning radius, that these nonuniformities do not provide a very practical way of telling one field from another.

5 It might be relevant if you are designing an airplane or a flight simulator.

6 In ordinary cases (specifically, when you know the forces are in balance and you are trying to figure out whether the torques are in balance) it doesn't matter what point in the airplane you choose as the pivot-point, provided you measure all lever arms from the same point.

7 Contrast this with the dot product used in [section 4.4](#).

8 Sometimes you see these written as hyphenated words (foot-pounds or newton-meters) in which case the hyphen should not be mistaken for a minus sign. A foot-pound is a foot times a pound, not a foot minus a pound.

9 It is best to feed the string through a small smooth tube, rather than just your bare hand. You might use a poultry baster, or the axial hole in a spool of thread.

10 The bean-bag acquires the necessary straight-line momentum, and energy, via the string. It cannot acquire angular momentum from the string, since that would require a lever arm perpendicular to the force. Since the string can only exert a force parallel to itself, the lever arm is zero, so the torque is zero.

11 For an object rotating around an axis of symmetry, the angular momentum is exactly aligned with the axis; for an oddball axis this might not be true.

12 The word "motion" here refers to moving or re-orienting the airplane as a whole, not to the rotation of the propeller.

13 If you use your right hand, as shown, the torque vector is pointed in the direction of your thumb. Engineers would call this a negative pitching moment, since the torque vector points opposite to the standard (left to right) direction of the pitch axis.

14 An even fancier inertial platform would keep a *position* (not just orientation) independent of straight-line accelerations.

15 See [figure 19.9](#) for the definition of the *X*, *Y*, and *Z* directions.

16 The instrument is not directly sensitive to any change in the direction the airplane is going, just to changes in the direction it is pointing.

17 The *X*, *Y*, and *Z* directions are defined in [figure 19.9](#).

20 The Atmosphere

If you don't like the weather in Ithaca, just wait a few minutes. It'll get worse.
— apologies to Mark Twain

20.1 Circulation Around Fronts and Low Pressure Centers

Because the earth is spinning and the air is moving, there are significant Coriolis effects.¹ You'll never understand how weather systems work unless you pay attention to this.

We are accustomed to seeing the rotation of storm systems depicted on the evening news, but you should remember that even a chunk of air that appears absolutely still on the weather map is rotating, because of the rotation of the earth as a whole. Any chunk of air that appears to rotate *on the map* must be rotating faster or slower than the underlying surface. (In particular, the air in a storm generally rotates faster, not slower.)

Based on their everyday indoor experience, people think they understand how air behaves:

- They know that the stream of air from a fan moves in a straight line, with no particular tendency to curve right or left.
- They know that once the fan is switched off, the airflow won't last very long or travel very far before being overcome by friction.

However, when we consider the outdoor airflow patterns that Mother Nature creates, the story changes completely. In a chunk of air that is many miles across, a mile thick, and a mile away from the surface, there can be airflow patterns that last for hours or days, because there is so much more inertia and so much less friction. During these hours or days, the earth will rotate quite a bit, so Coriolis effects will be very important.

Note: In this chapter, I will use the § symbol to indicate words that are correct in the northern hemisphere but which need to be reversed in the southern hemisphere. Readers in the northern hemisphere can ignore the § symbol.

20.1.1 Flow Around a Low


Suppose we start out in a situation where there is no wind, and where everything is in equilibrium. We choose the rotating Earth as our reference frame, which is a traditional and sensible choice. In this rotating frame we observe a centrifugal field, as well as the usual gravitational field, but the air has long ago distributed itself so that its pressure is in equilibrium with those fields.

Then suppose the pressure is suddenly changed, so there is a region where the pressure is lower than the aforementioned equilibrium pressure.


In some cases the low pressure region is roughly the same size in every direction, in which case it is called a *low pressure center* (or simply a *low*) and is marked with a big ``L" on weather maps. In other cases, the low pressure region is quite long and skinny, in which case it is called a *trough* and is marked ``trof" on the maps. See [figure 20.1](#).

In either case, we have a *pressure gradient*.² Each air parcel is subjected to an unbalanced force due to the pressure gradient.


Initially, each air parcel moves directly inward, in the direction of the pressure gradient, but whenever it moves it is subject to large sideways Coriolis forces, as shown in [figure 20.2](#). Before long, the motion is almost pure counterclockwise§ circulation *around* the low, as shown in [figure 20.3](#), and this pattern persists throughout most of the life of the low-pressure region. If you face downwind at locations such as the one marked *A*, the pressure gradient toward the left§ is just balanced by the Coriolis force to the right§, and the wind blows in a straight line parallel to the trough. At locations such as the one marked *B*, the pressure gradient is stronger than the Coriolis force. The net force deflects the air.



[Figure 20.1](#): Initial Force near a Low Pressure Region



[Figure 20.2](#): Initial Motion near a Low Pressure Region



[Figure 20.3](#): Steady Motion near a Low Pressure Region

Now we must account for friction (in addition to the other forces just mentioned). The direction of the frictional force will be opposite to the direction of motion. This will reduce the circulatory velocity. This allows the air to gradually spiral inward.


The unsophisticated idea that air should flow from a high pressure region toward a low pressure region is only correct in the very lowest layers of the atmosphere, where friction is dominant. If it weren't for friction, the low would never get filled in. At any reasonable altitude, friction is negligible — so the air aloft just spins around and around the low pressure region.

The astute reader may have noticed a similarity between the air in [figure 20.2](#) and the bean-bag in [figure 19.12](#). In one case, something gets pulled inwards and increases its circulatory motion ``because'' of Coriolis force, and in the other case something gets pulled inwards and increases its circulatory motion ``because'' of conservation of angular momentum. For a bean-bag, you can analyze it either way, and get the same answer. Also for a simple low-pressure center, you can analyze it either way, and get the same answer. For a trough, however, there is no convenient way to apply the conservation argument.

In any case, please do not get the idea that the air spins around a low partly because of conservation of angular momentum and partly because of the Coriolis force. Those are just two ways of looking at the same thing; they are not cumulative.

[20.1.2](#) Fronts and Troughs

As mentioned above, whenever the wind is blowing in a more-or-less straight line, there must low pressure on the left§ to balance the Coriolis force to the right§ (assuming you are facing downwind). In particular, the classic cold front wind pattern (shown in [figure 20.4](#)) is associated with a trough, (as shown in [figure 20.5](#)). The force generated by the low pressure is the only thing that could set up the characteristic frontal flow pattern.



[Figure 20.4](#): Wind Near a Front



[Figure 20.5](#): Pressure Near a Front



Figure 20.6: Temperature Near a Front

The wind shift is what defines the existence of the front. Air flows one way on one side of the front, and the other way on the other side (as shown in [figure 20.4](#)).

Usually the front is oriented approximately north/south, and the whole system is being carried west-to-east by the prevailing westerlies. In this case, we have the classic cold front scenario, as shown in [figure 20.4](#), [figure 20.5](#), and [figure 20.6](#). Ahead of the front, warm moist air flows in from the south. Behind the front, the cold dry air flows in from the north. Therefore the temperature drops when the front passes. In between cold fronts, there is typically a non-frontal gradual warming trend, with light winds.

You can use wind patterns to your advantage when you fly cross-country. If there is a front or a pressure center near your route, explore the winds aloft forecasts. Start by choosing a route that keeps the low pressure to your left. By adjusting your altitude and/or route you can often find a substantial tailwind (or at least a substantially decreased headwind).

Note: by ancient tradition, meteorologists name winds by the direction *from whence* they come. A south wind (or southerly wind) blows from south to north. Almost everything else is named the other way. An aircraft on a southerly heading is flying *toward* the south. Physicists and mathematicians name all vectors by the direction *toward* which they point. To avoid confusion, it is better to say "wind from the south" rather than "south wind".

A warm front is in many ways the *same* as a cold front. It is certainly not the opposite of a cold front. In particular, it is also a trough, and has the same cyclonic flow pattern.

A warm front typically results when a piece of normal cold front gets caught and spun backwards by the east-to-west flow just north of a strong low pressure center, as shown in [figure 20.7](#). That is, near the low pressure center, the wind circulating around the center is stronger than the overall west-to-east drift of the whole system.




Figure 20.7: Warm Front

If a warm front passes a given point, a cold front must have passed through a day or so earlier. The converse does not hold — cold front passage does *not* mean you should expect a warm front a day or so later. More commonly, the pressure is more-or-less equally low along most of the trough. There will be no warm front, and the cold front will be followed by fair weather until the next cold front.

Low pressure — including cold fronts and warm fronts — is associated with bad weather for a simple reason. The low pressure was created by an updraft that removed some of the air, carrying it up to the stratosphere. The air cools adiabatically as it rises. When it cools to its dew point, clouds and precipitation result. The latent heat of condensation makes the air warmer than its surroundings, strengthening the updraft.

Ascending air $\Rightarrow$ low pressure at the surface
 Ascending air $\Rightarrow$ clouds

The return flow down from the stratosphere (high pressure, very dry descending air, and no clouds) generally occurs over a wide area, not concentrated into any sort of front. There is no sudden wind shift, and no

sudden change in temperature. This is not considered "significant weather" and is not marked on the charts at all.

20.2 Pressure and Winds Aloft


Air shrinks when it gets cold. This simple idea has some important consequences. It affects your altimeter, as will be discussed in [section 20.2.3](#). It also explains some basic facts about the winds aloft, which we will discuss now.

20.2.1 Thermal Gradient Wind

Most non-pilots are not very aware of the winds aloft. Any pilot who has ever flown westbound in the winter is keenly aware of some basic facts:

- The winds aloft tend to come from the west.
- They are much stronger in the winter.
- They get stronger and stronger as altitude increases.

A typical situation is shown in [figure 20.8](#). In January, the average temperature in Vero Beach, Florida, is about 15 Centigrade (59 Fahrenheit), while the average temperature in Oshkosh, Wisconsin is about -10 Centigrade (14 Fahrenheit). Imagine a day where surface winds are very weak, and the sea-level barometric pressure is the same everywhere, namely 1013 millibars (29.92 inches of mercury).



[Figure 20.8](#): Thermal Gradient Wind

The pressure above Vero Beach will decrease with altitude. According to the International Standard Atmosphere (ISA), we expect the pressure to be 697 millibars at 10,000 feet.

Of course the pressure above Oshkosh will decrease with altitude, too, but it will not exactly follow the ISA, because the air is 25 centigrade colder than standard. Air shrinks when it gets cold. In the figure, I have drawn a stack of ten boxes at each site. Each box at VRB contains the same number of air molecules as the corresponding box at OSH.³ The pile of boxes is shorter at OSH than it is at VRB.

The fact that the OSH air column has shrunk (while the VRB air column has not) produces a big effect on the winds aloft. As we mentioned above, the pressure at VRB is 697 millibars at 10,000 feet. In contrast, the pressure at OSH is 672 millibars at the same altitude — a difference of 25 millibars.

This puts a huge force on the air. This force produces a motion, namely a wind of 28 knots out of the west. (Once again, during most of the life of this pressure pattern, the pressure gradient toward the left is just balanced by the Coriolis force to the right, assuming you are facing downwind.) This is the average wind at 10,000 feet, everywhere between VRB and OSH.

More generally, suppose surface pressures are reasonably uniform (which usually is the case) and temperatures are not uniform (which is usually the case, especially in winter). If you have low temperature on your left and high temperature on your right, you will have a tailwind aloft. The higher you go, the stronger the wind. This is called *thermal gradient wind*.

The wind speed will be proportional to the temperature gradient. Above a large airmass with uniform temperature, there will be no thermal gradient wind. But if there is a front between a warm airmass and a cold airmass, there will be a large temperature change over a short distance, and this can lead to truly enormous winds aloft.

In July, OSH warms up considerably, to about 20 centigrade, while VRB only warms up slightly, to about 25 centigrade. This is why the thermal gradient winds are typically much weaker in summer than in winter — only about 5 knots on the average at 10,000 feet.

In reality, the temperature change from Florida to Wisconsin does not occur perfectly smoothly; there may be large regions of relatively uniform temperature separated by rather abrupt temperature gradients — cold fronts or warm fronts. Above the uniform regions the thermal gradient winds will be weak, while above the fronts they will be much stronger.

For simplicity, the foregoing discussion assumed the sea-level pressure was the same everywhere. It also assumed that the temperature profile above any given point was determined by the surface temperature and the "standard atmosphere" lapse rate. You don't need to worry about such details; as a pilot you don't need to calculate your own winds-aloft forecasts. The purpose here is to make the official forecasts less surprising, less confusing, and easier to remember.

20.2.2 Altimetry

An aircraft altimeter does not really measure altitude. It really measures *pressure*, which is related to altitude, but it's not quite the same thing.

In order to get an estimate of the altitude, the altimeter depends on a two-step process. First, the altimeter has a knob whereby you can adjust things to account for how the local weather raises or lowers the atmospheric pressure. You should make this adjustment on the runway before takeoff, and for extended flights you should get updated settings via radio. If you neglect this, you could find yourself at a too-low altitude, if you fly to a region where the barometric pressure is lower. The mnemonic is: "High to low, look out below".

Altimeters are not perfect. Even if the altimeter and airplane were inspected yesterday, and found to be within tolerances,

±

The altimeter could be off by 30 feet when it reads 2500 (according to the tolerances in FAR 43 Appendix E).

-

If the airplane is moving at 100 knots, the indicated altitude could be off by another 30 feet, due to nonidealities in the arrangement of the static port (FAR 23.1325).

-

If the airplane is descending at 750 FPM, the altimeter could be off by an additional 70 feet, due to friction in the mechanism (FAR 43 Appendix E).

-

There could be 30 feet of hysteresis, if you have recently descended from a very high altitude (FAR 43 Appendix E).

-

Wind flowing over a nice airfoil-shaped hill can produce low pressure there. A 35-knot wind could produce a 50-foot altimetry error. See [section 3.4.1](#).

The first item could be off in either direction, but the other items will almost certainly be off in the bad direction when you are descending. Also, if the airplane has been in service for a few months since the last inspection, the calibration could have drifted a bit. All in all, it would be perfectly plausible to find that your altimeter was off by 50 feet when parked on the ground, and off by 200 feet in descending flight over hilly terrain.

20.2.3 High Altimeter due to Low Temperature

The altimeter measures a pressure and converts it to a so-called altitude. The conversion is based on the assumption that the actual atmospheric pressure varies with altitude the same way the standard atmosphere would. The pressure decreases by roughly 3.5% per thousand feet, more or less, depending on temperature.

The problem is that the instrument does not account for nonstandard temperature. Therefore if you set the altimeter to indicate correctly on the runway at a cold place, it will be inaccurate in flight. It will indicate that you are higher than you really are. This could get you into trouble if you are relying on the altimeter for terrain clearance. The mnemonic is HALT — High Altimeter because of Low Temperature.

As an example: Suppose you are flying an instrument approach into Saranac Lake, NY, according to the FAA-approved "Localizer Runway 23" procedure. The airport elevation is 1663 feet. You obtain an altimeter setting from the airport by radio, since you want your altimeter to be as accurate as possible when you reach the runway.

You also learn that the surface temperature is -32 Centigrade, which is rather cold but not unheard-of at this

location. That means the atmosphere is about 45 C colder than the standard atmosphere. That in turn means the air has shrunk by about 16%. Throughout the approach, you will be too low by an amount that is 16% of your height above the airport.

The procedure calls for crossing the outer marker at 3600 MSL and then descending to 2820 MSL, which is the Minimum Descent Altitude. That means that on final approach, you are *supposed* to be 1157 feet above the airport. If you blindly trust your altimeter, you will be 1157 "shrunk feet" above the airport, which is only about 980 real feet. You will be 180 real feet (210 shrunk feet) lower than you think. To put that number in perspective, remember that localizer approaches are designed to provide only 250 feet of obstacle clearance.⁴

You must combine this HALT error with the ordinary altimetry errors discussed in [section 20.2.2](#). The combination means you could be 400 feet lower than what the altimeter indicates — well below the protected airspace. You could hit the trees on Blue Hill, 3.9 nm northeast of the airport.

Indeed, you may be wondering why there haven't been lots of crashes already -- especially since the Minimum Descent Altitude used to be lower (1117 feet, until mid-year 2001). Possible explanations include:

- Most people use the ILS approach instead of the localizer approach. That provides electronic vertical guidance that isn't affected by temperature.
- In winter, the real atmosphere usually has a smaller lapse rate than the standard atmosphere, especially at the lower altitudes, so the errors are usually less than what simple theory would suggest.
- The FAA has overestimated the height of the trees. They routinely assume there will be small structures and trees rising 200 feet above the land surface, but the trees on Blue Hill are probably closer to 100 feet. This is helpful, but we shouldn't rely on it. The trees are still growing, and other trees in the vicinity are over 150 feet tall. Furthermore, if somebody built a 190-foot tower atop Blue Hill, the FAA would not change the Minimum Descent Altitude for this procedure, and there would be problems for sure.
- The new 1157-foot Minimum Descent Altitude is about 40 feet higher than what you would expect just based on the height of the hill. I'm told this represents an allowance for the effect of wind blowing over hilly terrain as mentioned above.

Even if people don't "usually" crash, we still need to do something to increase the margin for error.

There is an obvious way to improve the situation: In cold weather, you need to apply temperature compensation to all critical obstacle-clearance altitudes.

You can do an approximate calculation in your head: If it's cold, add 10%. If it's really, really cold, add 20%. Approximate compensation is a whole lot better than no compensation.

The percentages here are applied to the height *above the field*, or, more precisely, to the height above the station that is giving you your altimeter setting. In the present example, 20% of 1157 is about 230. Add that to 2820 to get 3050, which is the number you want to see on your altimeter during final approach. Note that this number, 3050, represents a peculiar mixture: 1663 real feet plus 1387 shrunk feet.

For better accuracy, you can use the following equation. The indicated altitude you want to see is:

$$288.15 - I F$$

$$A_i = F + \quad (20.1)$$

$$273.15 + T_f$$

where F is the field elevation, A_i is the true altitude you want to fly (so $A_i - F$ is the height above the field, in real feet), I is the standard lapse rate (2 °C per thousand feet), T_f is the temperature at the field, 273.15 is the conversion from Centigrade to absolute temperature (Kelvin), and 15 C = 288.15 K is the sea-level temperature of the standard atmosphere.

You might want to do this for a range of temperatures, and tabulate the results. An example is shown in [table](#)

[20.1](#). Make a row for each of the critical altitudes, not just the Minimum Descent Altitude. Then, for each flight, find the column that applies to the current conditions and pencil-in each number where it belongs on the approach plate.

Field Temperature, °C	12	0	-10	-20	-30	-40
South Sector	6600	6820	7000	7220	7440	7700
Northeast Sector	5100	5240	5380	5540	5680	5860
Northwest Sector	4100	4200	4300	4400	4520	4640
Procedure Turn	4800	4940	5060	5200	5340	5500
Crossing Outer Marker	3600	3680	3760	3840	3940	4020
Minimum Descent Alt	2820	2860	2920	2960	3020	3080

[Table 20.1](#): Saranac Lake Critical Altimeter Readings

It is dangerously easy to get complacent about the temperature compensation. You could live in New Jersey for years without needing to think about it -- but then you could fly to Saranac Lake in a couple of hours, and get a nasty surprise.

The HALT correction is important whenever temperatures are below standard and your height above the terrain is a small fraction of your height above the station that gave you your altimeter setting. This can happen in a couple of ways:

- When flying an instrument approach, you might need to apply 200 feet of compensation; this might make the difference between crashing and not crashing.
- When flying over tall mountains, you might need to apply a great deal of compensation; 12,000 shrunk feet might not be enough to get you over a 10,000-foot mountain.

[20.3](#) Prevailing Winds and Seasonal Winds

A parcel of air will have less density if it has


- a higher temperature,
- a higher dewpoint, and/or
- a lower pressure.

If a parcel of air is less dense than the surrounding air, it will be subject to an upward force.[5](#)

[20.3.1](#) Primary Circulation Patterns

As everyone knows, the tropics are hotter and more humid than the polar regions. Therefore there tends to be permanently rising air at the equator, and permanently sinking air at both poles.[6](#) This explains why equatorial regions are known for having a great deal of cloudy, rainy weather, and why the polar regions have remarkably clear skies.

You might think that the air would rise at the equator, travel to the poles at high altitude, descend at the poles, and travel back to the equator at low altitude. The actual situation is a bit more complicated, more like what is shown in [figure 20.9](#). In each hemisphere, there are actually *three* giant cells of circulation. Roughly speaking, there is rising air at the equator, descending air at 25 degrees latitude, rising air at 55 degrees latitude, and descending air at the poles. This helps explain why there are great deserts near latitude 25 degrees in several parts of the world.



[Figure 20.9](#): Primary Circulation Cells

The three cells are named as follows: the *Hadley cell* (after the person who first surmised that such things existed, 250 years ago), the *Ferrel cell*, and the *polar cell*. The whole picture is called the *tricellular theory*. It correctly describes some interesting features of the real-world situation, but there are other features that it does not describe correctly, so it shouldn't be taken overly-seriously.

You may be wondering why there are three cells in each hemisphere, as opposed to one, or five, or some other number. The answer has to do with the size of the earth (24,000 miles in circumference), its speed of rotation, the thickness of the atmosphere (a few miles), the viscosity of the air, the brightness of the sun, and so forth. I don't know how to prove that three is the right answer — so let's just take it as an observed fact.

Low pressure near 55 degrees coupled with high pressure near 25 degrees creates a force pushing the air towards the north§ in the temperate regions. The air responds to this force with motion in the perpendicular direction, namely from west to east. As shown in [figure 20.10](#), these are the *prevailing westerlies* that are familiar to people who live in these areas.

According to the same logic, low pressure near the equator coupled with high pressure near 25 degrees creates a force toward the equator. The air responds to this force with motion in the perpendicular direction, namely eastward. These are the famous *trade winds*, which are typically found at low latitudes in each hemisphere, as shown in [figure 20.10](#).



[Figure 20.10](#): Primary Prevailing Winds

In days of old, sailing-ship captains would use the trade winds to travel in one direction and use the prevailing westerlies to travel in the other direction. The regions in between, where there was sunny weather but no prevailing wind, were named the *horse latitudes*. The region near the equator where there was cloudy weather and no prevailing wind was called the *doldrums*.

The boundaries of these great circulatory cells move with the sun. That is, they are found in more northerly positions in July and in more southerly positions in January. In certain locales, this can produce a tremendous seasonal shift in the prevailing wind, which is called a *monsoon*.⁷

20.3.2 Continental / Oceanic Patterns

Now let us add a couple more facts:

1. The sun is not very effective at heating the air, especially dry air. Normally, the sun heats the surface of the planet, then the air gains heat from the surface — partly by simple contact, and partly by absorbing energy-rich water vapor that evaporates from the surface.
2. When we change from winter to summer, solar heating warms the dry land much more quickly than the ocean.⁸ This is because the ocean is constantly being stirred. To heat up the land, you need only heat up the top few inches of soil. To heat up the ocean, you need to heat up several feet of water.

As a consequence, in temperate latitudes, we find that in summer, the land is *hotter* than the ocean (other things, such as latitude, being constant), whereas in winter the land is *colder* than the ocean.

This dissimilar heating of land and water creates huge areas of low pressure, rising air, and cyclonic flow over the oceans in winter, along with a huge area of high pressure and descending air over Siberia. Conversely there are huge areas of high pressure, descending air, and anticyclonic flow over the oceans in summer.

These continental / oceanic patterns are superimposed on the primary circulation patterns. In some parts of the world, one or the other is dominant. In other parts of the world, there is a day-by-day struggle between them.

20.4 Summary

Very near the surface (where friction dominates), air flows from high pressure to low pressure, just as water flows downhill. Meanwhile, in the other 99% of the atmosphere (where Coriolis effects dominate) the motion tends to be *perpendicular* to the applied force. The air flows clockwise§ around a high pressure center and counterclockwise§ around a low pressure center, cold front, or warm front.

Although trying to figure out all the details of the atmosphere from first principles is definitely not worth the trouble, it is comforting to know that the main features of the wind patterns make sense. They do not arise by magic; they arise as consequences of ordinary physical processes like thermal expansion and the Coriolis effect.

If you really want to know what the winds are doing at 10,000 feet, get the latest 700 millibar constant pressure analysis chart and have a look. These charts used to be nearly impossible for general-aviation pilots to obtain, but the situation is improving. Now you can get them by computer network or fax. On a trip of any length, this is well worth the trouble when you think of the time and fuel you can save by finding a good tailwind.

A few rules of thumb: eastbound in the winter, fly high. Westbound in the winter, fly lower. In the summer, it doesn't matter nearly as much. In general, try to keep low pressure to your left§ and high pressure to your right§.

1

The origin of the Coriolis effect is discussed in [section 19.3](#).

2

In general, a *gradient* has to do with how steeply something changes from place to place.

3

The bottom box starts at sea level at both sites. We ignore the fact that OSH is actually 808 feet above sea level. The fact that the ground ``sticks up" into the bottom box doesn't change the essence of the argument. This is consistent with the notion that you adjust your altimeter to read 808 (not zero) on the ground at OSH.

4

The idea to get you down as low as possible, to maximize your ability to get below the cloud ceiling so you can find the airport in bad weather.

5

It would be simpler, but less accurate, to say ``hot air rises". For one thing, if all the air is hot, none of it will rise. Secondly, it is important to keep in mind that an upward force is not necessarily the same as upward motion.

6

Although there is, as expected, somewhat low pressure at the equator (and very low density, when you take humidity into account), there is not any noticeable high pressure at the poles. In fact, there is phenomenally low pressure at the south pole. I have no idea why this is. Sorry.

7

Many people take the word ``monsoon" to mean ``lots of rain", but that's not the only (or even the primary) meaning. It comes from an Arabic word meaning ``season", hence ``seasonal wind". Now in parts of India and various other places, one of the seasonal winds comes from the ocean, bringing lots of rain.

8

A similar thing happens, on a smaller scale, when we change from night to day.

21 Pilot-In-Command Decisionmaking

The really good pilots use their superior judgment to keep them out of situations where they might be required to demonstrate their superior skill.

21.1 Decisionmaking In General

Piloting requires a range of skills:

- At one extreme are “reflexes” that involve lots of eye-hand coordination and tactics rather than strategy. There are also some basic thinking skills, such as giving answers to clearly-posed questions. These skills are relatively easy to teach, and relatively easy to evaluate.
- At the other extreme are advanced decisionmaking skills. As pilot in command, you will need to make decisions in situations where it’s not even obvious that a decision needs to be made. This requires being able to notice small things as well as being able to see the big picture. It also requires not assuming that whatever procedure you used last time is going to work next time. These skills are extremely important. They are, alas, not so easy to teach or easy to evaluate.

As an example: In a lesson, you might be asked to demonstrate short-field landing procedure. It isn’t your decision, since you have been *told* to use short-field procedure. In contrast, suppose that years from now, you want to fly yourself to XYZ airport. Your former instructor knows the XYZ runway is very short, but do you? The instructor won’t be there to tell you to use short-field procedure. What’s worse, there won’t even be anybody to *ask* you what procedure is called for. Will you be wise enough to ask the question on your own? Perhaps questions such as “should we be using short-field procedure?” and “should we be using this runway at all?” ought to be on the takeoff and landing checklists, but in light aircraft they typically aren’t.

- Let’s assume you know how to calculate the runway requirements using the POH. That’s easy, if and when somebody asks the question.
- Let’s assume you know how to use short-field technique. That’s not the issue.
- The question is, will you know *whether* to use short-field technique, when there’s nobody there to ask the question.

This sort of decisionmaking is a high-level skill. It is not well understood. Common sense is good, but this goes beyond common sense. Planning ahead is good, but this goes beyond planning. Logic is good, but this goes beyond mere syllogisms. Strategy is good, but this goes beyond the usual definition of strategy. Let’s just call it decisionmaking.

There are innumerable situations that require good decisionmaking.

As just discussed, one aspect of decisionmaking is to recognize that a decision is needed, even though it’s not in the form of a clearly-posed question. A somewhat-related aspect involves starting from a seemingly-small, seemingly-isolated fact, then seeing how it connects to the other facts you know, and working out the implications and ramifications.

As an example of ramifications: Suppose you were expecting a 20-knot tailwind, but you’ve actually got a 20-knot headwind. For starters, you have to notice this. It won’t be obvious, unless you’re checking arrival-time at enroute waypoints, and/or checking DME or GPS groundspeed. Then you have to work out the implications. The obvious implication is that you’re going to be late. A less-obvious implication is that you might not have enough fuel to reach your originally-intended destination. If you have to divert, don’t wait until you’re low on fuel to do it. An even-less-obvious implication is that if the forecast is wrong about the winds, it’s probably wrong about everything else: ceiling, visibility, temperature, precipitation, icing, turbulence, et cetera. You’d be well advised to get an updated forecast.

21.2 Decisionmaking during Lessons

I have adopted an instructional style designed to exercise your decisionmaking muscle. (Other instructors may have different styles. Discuss it in advance with your instructor, to avoid misunderstandings.)

21.2.1 Please Act As PIC During Lessons

During introductory lessons, you start out with minimal responsibilities. Then you gradually acquire more and more responsibilities. In non-introductory situations, I want you to act like pilot-in-command as much as possible.

With one exception ([section 21.2.3](#)), everything I say is merely a suggestion. My suggestions aren't meant to relieve you of your responsibility as pilot-in-command. For instance, if I ask you to turn right, you remain responsible for clearing the area. Please clear the area just as you would if you were solo. Also think about whether the new heading will take us into restricted airspace or some such. If you don't want to turn right, we can discuss it.

If we are in a situation that you ought to be able to handle on your own, I will generally let you handle it. If I need to contribute, I like to use a multi-stage "escalation" process:

- Ideally, I don't need to say anything. If we are facing an energy-management challenge, you can notice it (the sooner the better) and deal with it.
- If you don't deal with it on your own, I'll start asking questions, such as: "How's your energy? Are we high and fast, or low and slow?"
- Then come more-explicit statements: "It looks like the angle from the horizon to the aim point is growing. If you don't do something we're going to overshoot the runway."
- Then it escalates to an instruction: "Go around."
- Then the instructions become more detailed: "Add power. Raise the nose. Start retracting the flaps."
- Finally: **"I've got it."**

Remember, being a pilot means making decisions, even during lessons (except introductory lessons). During the escalation process, I'm gradually shifting more of the decisionmaking onto my shoulders. Your goal should be to take the hints at the earliest possible stage, so that further stages of escalation are not needed.

As another illustration of the same idea, I try to avoid giving an instruction such as "go around". If I see a deer on the runway, I'll say there's a deer on the runway, and you can come to your own conclusion about going around. If we need to do some go-arounds just for practice, I might say there's a hippopotamus on the runway. You know it's not real, but I want you to *pretend* there's an obstruction, and come to your own conclusions about how to deal with it. Most likely you will decide to go around.

At the other extreme, if you are struggling with an unfamiliar situation, I'll just tell you how to deal with it. No hints, no escalation. There are lots of good but non-obvious techniques, and I don't expect you to re-invent them on your own.

Also note that one element of good PIC decisionmaking is knowing when to ask for help. This includes asking for clarification of an overly-vague hint.

[21.2.2](#) How's Your Workload?

From time to time I will ask you "How's your workload?" There are many possible answers, including:

- Swamped. Could use some help.
- Pretty busy right now.
- Workload's not too bad right now.

If you're swamped, I'll help. If you're busy, I'll leave you alone. If you're not busy, I might strike up a conversation about strategy or tactics, or suggest an exercise.

[21.2.3](#) I've Got It

If I say **"I've got it"**, that means I am taking command of the airplane and I don't want any delay or any question about it. (We will discuss it afterward.)

Notice the important distinction:

- "How about I fly for a bit?" or "Would you like me to demonstrate that maneuver?" Those are simply questions, perhaps verging on suggestions. Those are negotiable.

- **"I've got it."** This is not a suggestion. This is absolutely not negotiable. This is necessary to preserve safety.

21.2.4 Hood Work

When you are under the hood, practicing flight by reference to instruments, you should start by telling your safety pilot (whether it's me or somebody else) "I'm delegating the traffic-spotting to you". The safety pilot should give you a readback on this, saying something like "I accept the delegation". You should insist on this.

You remain pilot-in-command and you even retain a share of the responsibility for traffic separation. Before turning (except small shallow turns), ask "Clear right?" or "Clear left?"; don't assume your safety-pilot has pre-cleared all turns.

When practicing an instrument approach, I will tell you if/when we break out of the simulated clouds. If in doubt, you can ask whether we have broken out. This is another decisionmaking exercise. If we reach the missed-approach point or decision height and haven't broken out, *do not expect me to say anything* at this point. I want you to decide on your own when it's time to begin the missed-approach procedure.

If I say "you've got the approach lights" it means you haven't entirely broken out of the simulated clouds, but you have the option of continuing the approach in accordance with FAR 91.175(c)3(i).

21.3 Layers of Safety

One of the standard ways to achieve a high level of safety is to use a layered approach: layers and layers of backups and crosschecks.

For example, before takeoff, always check the fuel level by looking in the tanks. If you can't accurately judge the fuel level by eye, use a pipette to measure it. Then see what the cockpit fuel gauges are reading. Crosscheck the two types of measurement. If they disagree, you've got a problem.

Similarly, during the course of the flight, you have two ways of estimating how much fuel remains: (a) Start with what you had at takeoff, and decrement it according to the expected fuel-burn rate, and (b) look at the gauges. Method (a) will fool you if there is a leak or other problem in the fuel system, and method (b) will fool you if the gauge is stuck, but the chance of both problems happening at the same time is remote. (The third layer of safety is to make a forced landing, but you hope it doesn't come to that.)

Don't lightly give away layers of safety. For example, if you don't do a magneto check before each flight, it's just a matter of time before one mag fails. You won't notice this, especially if it is the right mag (since only the left mag is used for starting). Then it's just a matter of time before the other mag fails in flight. You will notice it then, because the engine will suddenly quit.

The notion of layers of safety applies to many aspects of flying:

- Having two magnetos doesn't make the engine twice as reliable; it makes it *thousands* of times more reliable.
- Pipetting the tanks *and* keeping an eye on the fuel gauges is thousands of times more reliable than either one separately.
- As discussed in [section 12.7.4](#) there is a long list of cues you can use to make sure you aren't landing with a tailwind and/or excessive airspeed. Sometimes some cues will be useless or worse: perhaps the windsock is not visible from pattern altitude, and perhaps the previous airplane used the wrong runway. And perhaps you will occasionally overlook one or two cues. But that still leaves many, many cues that will keep you out of trouble.
- You have multiple sources of navigation information (pilotage, dead reckoning, VOR, GPS, ATC radar) which can be cross-checked against each other.

You should pay attention to anything that peels away one or more layers of safety. Keep track of how many layers remain.

If one magneto fails, park the airplane until it is fixed!

If the fuel gauges cannot be trusted, park the airplane until it is fixed! Do not rely on clock-and-dipstick alone, or on the gauges alone.

21.4 Example: Obstacle Clearance

In other publications, obstacle clearance is commonly discussed under the heading “controlled flight into terrain” (CFIT). The term obstacle clearance is preferable, partly because it puts a more positive spin on things: it is better to talk about your obstacle clearance successes than your CFIT failures. Also, the CFIT statistics are misnamed because they include collisions with trees, man-made structures, bodies of water, etc. that you might not have thought of as “terrain”. They also include taxiing into potholes and other things that you might not have thought of as “flight”.

We ought to pay serious attention to the obstacle clearance issue, because statistics show a surprisingly large number of accidents where a perfectly good aircraft collides with an obstacle. You would think such accidents would be entirely preventable, so even one occurrence is far too many.

Obstacles can be a factor during any phase of flight, including departure, enroute, or approach. A typical accident scenario goes something like this: At night (or in hazy weather), at an unfamiliar airport, the pilots crash into power lines or into a hillside.

Let’s analyze this scenario using our notion of layers of safety. Let’s ask what “caused” this accident.

- Did they crash “because” of the obstacles? If they had been flying somewhere with more benign topography, they wouldn’t have gotten into trouble.
- Did they crash “because” they made a wrong turn? Presumably every airport has *some* path that airplanes can safely follow on takeoff. Perhaps they weren’t familiar with the correct procedure, or perhaps they just neglected to follow the correct procedure.
- Did they crash “because” of darkness or hazy weather? If it had been daytime in clear weather, they would have seen the approaching obstacle in time to turn away.

We say this accident had at least three causative factors. Each of the factors was “a” contributory cause of the accident, but none was “the” sole cause of the accident.

Multi-factor situations like this can be a challenge to your decisionmaking skills. [Section 21.3](#) says you should not lightly give away layers of safety. But what does “lightly” mean? Sometimes there are good reasons for accepting some risk. Sometimes it’s OK to fly at night, or in hazy weather. Sometimes it’s OK to fly to an unfamiliar field. But don’t get too complacent. If you get complacent about each risk factor separately you can get into big trouble if/when multiple risk factors gang up on you.

Anything that involves operating at low altitudes peels off one or two layers of safety. In addition to ordinary approaches and departures, there are many examples including patrol, photo work, crop dusting, scud running (i.e. flying at low altitudes to avoid clouds), buzzing (i.e. flying at low altitudes to show off), and mountain flying.

The departure phase and approach phase account for a huge proportion of the obstacle clearance problems. You need to worry about this even in regions that are not considered mountainous. A modest hill or a modest structure can be a serious threat if it’s near a runway.

VFR at an unfamiliar field at night (or in hazy weather) is particularly risky, as discussed in [section 12.1.3](#) and [section 13.7.5](#).

In the enroute phase, the primary obstacle-clearance technique is to choose a suitable route and a suitable altitude, as discussed in [section 14.8](#). A good secondary technique, to reduce the chance of mistakes, is to get radar advisories. The ATC computers know the minimum safe enroute altitude in each sector, and will sound an alarm if you get too low. Similarly, some fancy RNAV units now contain obstacle-clearance data and will give you warning of approaching threats. Another thing that may be of some help is an altitude alerter. This is a simple, cheap instrument. You tell it what you have chosen as your intended altitude, and it will beep if you inadvertently drift above or below that. Alas, this won’t help much if you punch an unsuitable number into the instrument (due to bad planning or whatever) and it won’t help if you are trying to fly through a mountain pass and get off-course horizontally.

21.5 Flow Pattern

During the preflight check, you should walk around the airplane and check everything that you come to, in order. This is an example of using the *flow pattern*.

After examining things according to the flow pattern, you should run the checklist to see what you overlooked.

Checklists are good. Flow patterns are good. Neither one is a substitute for the other; instead, each is a *backup* for the other. Using them both is much, much better than using either one alone.

There are many situations where you can use a flow pattern, including:

- Preflight walk-around.
- Pre-takeoff instrument check.
- Enroute instrument check.
- Engine failure / restart ([section 15.1.1](#)).
- Et cetera.

21.6 Checklists

If you don't use a written checklist, it is just a matter of time until you forget something.

There are some checklists that you should commit to memory, such as rejected takeoff, spins, fire in flight, go-around, and possibly others, depending on how complex your airplane is. But even these should not be entrusted to long-term memory. Short-term memory is better than long-term memory, so refresh your memory at frequent intervals. An excellent method is to recite the checklist out loud, while somebody else checks your version against the written version.

I know a lot of pilots who fastidiously use a written checklist for preflight, but rely on memory for the approach and landing checklists. It is fairly easy to see how people fall into this trap: During preflight you are not strapped into your seat, and you are not busy flying the airplane. You can always take a minute to find the POH and read through it. In contrast, when you are setting up for a landing, the book is likely to be somewhere hard to reach and you're too busy to do much reading anyway.

Therefore, here are some constructive suggestions. Pick the one(s) you like best:

- Make a *pocket checklist*. Print up the checklists you are going to need during flight and fold them in such a way that they fit in a pocket. During preflight, put the list in a convenient pocket.
- If you habitually use a lap desk, tape a copy of the in-flight checklists to the lap desk itself.
- If space can be found, glue a copy of the in-flight checklists to the instrument panel.

The pocket checklist should include the approach, landing, and go-around checklists, as discussed in [section 12.1](#).

If you fly more than one airplane, make sure you have an appropriate pocket checklist or lap-desk checklist for each of them. As you progress in your pilot career, you will be flying progressively more complex aircraft, and if you persist in using the same old checklist you will get into trouble some day. Some aircraft have retractable landing gear; some don't. Some aircraft have cowl flaps; some don't. Some aircraft require using carburetor heat; some don't. Some aircraft require switching on the electric fuel pump for landing; others forbid it.

Discipline yourself to pay attention to the checklist. Don't just keep it in your pocket as a good-luck charm.

My pocket checklist also includes an enroute checklist, which is only three words long: "indications, configuration, location". That is shorthand for the following:

Indications: Go left-to-right across the panel and check everything, including the gauges that aren't part of your ordinary moment-to-moment scan: Fuel level, engine instruments, et cetera. Check that the directional gyro is aligned with the compass.

Configuration: See if the fuel/air mixture is appropriate for this phase of flight. Make sure you're not flying cross-country with the flaps extended. On more complex airplanes, there are more things to check: landing gear, cowl flaps, speed brakes, et cetera.

Location: Where are we? Where's the nearest airport? Mark this location on the chart, along with the time, as discussed in [section 14.2](#).

21.7 Personal Minimums

All too often, pilots get into risky situations without even realizing it. They don't consciously decide to run a

risk. They just take off on a supposedly routine flight, and by the time they notice a problem it is already too late to solve the problem.

In theory, you can avoid such problems by paying meticulous attention to "everything". But in reality, it is unreasonable to expect people to be super-meticulous all the time. The trick is to be aware of what's routine and what's not. It helps to have a clear-cut set of *personal minimums*.

Personal minimums are distinct from regulatory minimums. For instance, the regulations might permit a pilot with little experience to fly an unfamiliar aircraft on a maximum-range mission over water at night in moderate turbulence, then land in a 25-knot crosswind on an unfamiliar narrow obstructed unlighted runway, having had little sleep the night before ... but I don't recommend it.

Write down your personal minimums in the form of a checklist. Review the list before each proposed flight. If you are within the limitations, fine. If the proposed mission is slightly outside the envelope in one or two aspects, you might want to go ahead with it anyway -- *provided* you are super-meticulous. The checklist is warning you that this flight is not routine.

Issues to consider include:

1. At least _____ runway length for _____ airplane for density altitude below 2000 feet.
2. At least _____ runway length for _____ airplane for density altitude between 2000 and 5000 feet.
3. At least _____ runway width.
4. At most _____ knots gust component along the runway.
5. At most _____ knots across the runway, including gusts.
6. At least ____:____ fuel reserves.
7. Ceiling _____ and visibility _____ for VFR.
8. Ceiling _____ and visibility _____ at destination for IFR.
9. Survival equipment for flight over wilderness.
10. Survival equipment for flight beyond gliding range from dry land.
11. At least _____ hours experience in this make & model.
12. At most _____ pressure altitude without oxygen mask.
13. Physiological issues: "I-M-SAFE"
 - Illness?
 - Medication?
 - Stress?
 - Alcohol?
 - Fatigue?
 - Eating?
14. Turbulence?
15. Dark VFR?
16. Mountain flying? Bush flying? Obstructions enroute? Obstructions near the airports to be used?
Low-altitude operations?
17. In-flight "how-goes-it" checkpoints and "proceed/divert" checkpoints.

There are lots of other things you can legally do with just a private pilot certificate that are, alas, not covered

in the basic private-pilot training. Examples include

- If you're flying on a moonless night over unlighted areas, it can be a challenge to keep the airplane right-side-up. This requires instrument-flying skills far beyond what is required for the private-pilot checkride. (Night VFR is prohibited in many countries, but not all.)
- Specialized skills are required for mountain flying, bush flying, low-level patrol, crop dusting, aerobatics, formation flying, et cetera.
- Beware of the obstacle-clearance issues listed in [section 21.4](#).
- The first time you fly into a big, busy place like O'Hare, you might want to take along an instructor or at least a copilot who's been there before.
- Hand-propping (i.e. starting the engine by pulling on the propeller) is potentially very dangerous. You may have seen it done in some old Hollywood movie, but that doesn't mean you're qualified to do it. Don't try it unless you've got a good reason *and* have been thoroughly trained on the procedure.
- If you routinely fly solo out of a short field, it doesn't prove you can depart from there carrying three large passengers and their luggage. Maybe you can, maybe you can't, but it's not *routine*, so you need to get out the book and redo the weight & balance calculations, the performance calculations, et cetera.

These are just a few examples out of many. It is your responsibility to recognize when a situation is outside the envelope of your training and experience. It is your responsibility to acquire whatever skills are required.

Don't try to impose your personal minimums on the pilot next door. Yours will be too strict in some aspects and too lax in others. Your personal minimums are designed by you, for you. That's why they're called "personal" minimums.

The idea of personal minimums applies during the whole flight, not just during preflight. In some cases you should establish specific personal checkpoints for a specific flight. For instance, if you're flying into questionable weather, pick a specific checkpoint where you will get an updated weather report and decide whether to divert or not. Similarly, if it's a maximum-range mission, establish checkpoints along the way, where you will re-evaluate the headwinds, fuel quantity, et cetera. The idea is to notice *early* if Plan A isn't working, so you can execute Plan B while there's still time.

21.8 Skepticism, and Crisp Execution of Plan B

You need to notice things. That means more than just seeing things; you need to appreciate the significance of what you are seeing.

Skepticism means, among other things, not assuming that the way things are is the way things should be. For example:

- Suppose that during preflight you see a red light on one wingtip and a green light on the other wingtip. Are you sure you would notice if they were interchanged, so that both wingtips had the wrong color?
- Suppose you see that the tow-bar is lying loose in the back of the airplane. You could just leave it that way, on the theory that people have flown the airplane hundreds of times in that condition without getting into trouble. But suppose you get into an unusual attitude or a minor crash; you don't want the tow-bar to come whizzing toward you like a spear. Leaving it loose is needlessly throwing away one of your layers of safety. Anything that can't be stowed securely should be left on the ground until you get back.
- Suppose you find a puddle of oil underneath the engine. You must be skeptical about where it came from. I've seen puddles on about 20 occasions. Usually it's just because some klutz spilled oil while trying to add it to the crankcase. But on one occasion it was due to a moderately serious leak in the engine, and on another occasion it was due to a very serious crack in a cylinder, the sort of crack that will get rapidly worse during flight. The only way to know for sure is to take off the cowlings, mop up the mess, run the engine for a few minutes, and check for leakage. That's a lot of bother, but it's preferable to risking in-flight engine failure.

You want to be properly skeptical without being paranoid. Nothing in this world is perfect. If you cancel the flight whenever the airplane is not perfect, you'll never go flying. Judgement is required.

Also: Piloting requires flexible thinking. Do not think you can plan a flight and then fly it exactly as planned.

- First of all, you have to have a Plan B. (See [section 21.3.](#))
- Secondly, you need to promptly recognize when Plan A has gone to pot. (See also [section 21.7.](#))
- If Plan A isn't working, don't persevere with Plan A! We need crisp execution of Plan B.

The very first takeoff on my private pilot checkride was supposed to be a simulated soft-field takeoff. That requires, among other things, not stopping after leaving the run-up area, lest we sink into the simulated mud. I was nervous. I wanted to make a good first impression.

We were cleared for takeoff on runway 15 Left. Another aircraft was cleared to land on runway 15 Right. As I turned onto the runway, I heard Tower yelling at the other guy, pointedly reminding him to land on the *right*. I didn't hear any response. I had no idea whether the other aircraft was near or far; all I could do was try to imagine what could cause Tower to say such things. I imagined that the other pilot was planning on landing on the left and was persevering with his Plan A even though he'd been cleared to do something else.

I didn't like what I was imagining, so I pulled the throttle to idle and stomped on the brakes. I also keyed the transmitter and said "Tower, Two-Four-Kilo is gonna hold our position for a moment". Now I was *really* nervous. I had planned to comply with all ATC clearances, such as takeoff clearance. I had also planned to comply with the examiner's request for soft-field procedure. But there I was, stopped on the runway. My Plan A was in shambles.

I had no idea whether holding my position would make things better or worse. I figured it was a 50/50 chance. The deciding factor was not the odds but the relative payoff: if there's going to be a collision, I'd rather have a collision on the ground than in the air.

A moment later, the other aircraft flew right over top of us, about 10 feet up. My Plan B was starting to look pretty good.

Tower said "Two-Four-Kilo, advise when you're ready for departure". I said "as soon as that guy is out of our hair, we're ready". Tower acted like it was no big deal, and just said "Two-Four-Kilo, runway 15 Left, cleared for takeoff".

The examiner sat there with his poker-face on. I couldn't tell whether or not he approved of what I'd done. Looking back with the benefit of years of PIC experience, I'm quite sure I did the right thing, and I'm quite sure the examiner and the tower controller were glad that I deviated from their instructions.

Still, it is worth remembering that at the time, I was uncertain about this. I found the decision difficult and stressful. People aren't born with advanced decisionmaking skills. Training is needed.

21.9 Leadership and Resource Management

Be smart about using all the resources available to you.

If you have a co-pilot, that's an important resource. You can delegate certain tasks to your co-pilot.

- A very useful technique for expediting a flight is to let one pilot do the pre-flight walkaround while the other deals with weather, flight plans, et cetera.
- In busy airspace, it works well to let one pilot work the radios while the other does everything else.
- Co-pilots and even passengers can help spot traffic.
- There are many other examples of effective collaboration. Entire books have been written on the subject.

On the other hand, make sure you don't get into a situation where two pilots are worse than one. This can happen more easily than you might think:

- I've seen situations where each pilot assumed the other would take care of something, but it never got taken care of.
- I've even seen the following: It was time to toggle to a new frequency. One pilot pushed the "toggle" button. The other pilot pushed it also. As a result, the frequency was not what either pilot was expecting.

To prevent such problems, make sure there is agreement about who is pilot-in-command. If you are PIC, you retain final authority and responsibility for everything. If you delegate something, make sure the delegation is understood and carried out. If you are second-in-command, make sure the PIC knows what you've done and not done.

Learn from the story of Eastern Airlines flight 401. They crashed an airliner into the swamp partly because all three pilots were preoccupied with debugging a landing-gear indicator light. (Of course this is only part of the story.) Whether you have one pilot on board or three, don't let a small problem interfere with your basic responsibility to *fly the airplane*. That includes maintaining a safe altitude, seeing and avoiding other traffic, et cetera. If you have more than one pilot, let one fly the airplane while another debugs the problem.

Fly the airplane.

Resource management is often called "cockpit" resource management (CRM), but it should include resources outside the cockpit, notably Flight Service and ATC.

A related issue is *cockpit leadership*. If you are PIC, don't act like Captain Bligh. Encourage your crew members to speak up if they see anything questionable. Keep them informed as to your intentions, so that they will be better able to notice if something unintended is happening. If you are a crewmember, it is your duty to speak up if you see something amiss, no matter how surly the PIC is.

21.10 Learn from the Experience of Others

When pilots get together, they often trade stories about accidents or near-accidents. Non-pilots are sometimes shocked; they think it's ghoulish. But that's not the point at all. The point is to learn what led up to the problem, and what could have been done to prevent it.

You should study the accident statistics, so you know what are the big worries and what are the relatively minor worries. There are many sources of such information, including:

- Various pilot-oriented magazines feature every month a new anecdote from some pilot who learned something the hard way.
- The NTSB keeps records on all accidents and incidents. These are available on their web site.
- The ASRS (Aviation Safety Reporting System) collects reports on untoward events, even those that don't result in accidents. Every month they publish a discussion called *Callback* which is available from the NASA website and otherwise.
- The AOPA Air Safety Foundation puts out the annual *Nall Report* summarizing statistics from multiple sources, taking a general-aviation point of view. This is available from their web site and otherwise.
- There are many published books containing aviation biographies or collected stories.

21.11 Try to Outdo Yourself

Part of the romance of aviation is to do everything better than necessary. If the runway is 50 feet wide and the airplane's wheelbase is 10 feet wide, it is technically possible to land on the left half or the right half of the runway. But everybody tries to land exactly on the centerline. If you were off by one foot last time, try to be off by half a foot next time.

Safety is not directly affected by your best performance, or even your average performance. What matters, directly, is your worst-ever performance. This is called the *minimax* principle: make sure your worst-case performance is good enough. This partly involves skill, but largely involves using judgement to stay out of situations (distractions, fatigue, bad weather, etc.) that might cause your performance to be significantly worse than usual.

High standards contribute indirectly to safety in the following way: If your usual tolerances are tight enough, then even on the super-rare occasions when your performance is ten times worse than usual, you will still have a wide margin of safety.