NONLINEAR DYNAMICS AND CHAOS WITH STUDENT SOLUTIONS MANUAL

WITH APPLICATIONS TO PHYSICS, BIOLOGY, CHEMISTRY, AND ENGINEERING, SECOND EDITION

SECOND EDITION

Steven H. Strogatz

ROUTLEDGE

CONTENTS

PREFACE TO THE SECOND EDITION

Welcome to this second edition of *Nonlinear Dynamics and Chaos*, now available in e-book format as well as traditional print.

In the twenty years since this book first appeared, the ideas and techniques of nonlinear dynamics and chaos have found application in such exciting new fields as systems biology, evolutionary game theory, and sociophysics. To give you a taste of these recent developments, I've added about twenty substantial new exercises that I hope will entice you to learn more. The fields and applications include (with the associated exercises listed in parentheses):

Animal behavior: calling rhythms of Japanese tree frogs (8.6.9)
Classical mechanics: driven pendulum with quadratic damping (8.5.5)
Ecology: predator-prey model; periodic harvesting (7.2.18, 8.5.4)
Evolutionary biology: survival of the fittest (2.3.5, 6.4.8)
Evolutionary game theory: rock-paper-scissors (6.5.20, 7.3.12)
Linguistics: language death (2.3.6)
Prebiotic chemistry: hypercycles (6.4.10)
Psychology and literature: love dynamics in *Gone with the Wind* (7.2.19)
Macroeconomics: Keynesian cross model of a national economy (6.4.9)
Mathematics: repeated exponentiation (10.4.11)
Neuroscience: binocular rivalry in visual perception (8.1.14, 8.2.17)
Sociophysics: opinion dynamics (6.4.11, 8.1.15)
Systems biology: protein dynamics (3.7.7, 3.7.8)

Thanks to my colleagues Danny Abrams, Bob Behringer, Dirk Brockmann, Michael Elowitz, Roy Goodman, Jeff Hasty, Chad Higdon-Topaz, Mogens Jensen, Nancy Kopell, Tanya Leise, Govind Menon, Richard Murray, Mary Silber, Jim Sochacki, Jean-Luc Thiffeault, John Tyson, Chris Wiggins, and Mary Lou Zeeman for their suggestions about possible new exercises. I am especially grateful to Bard Ermentrout for devising the exercises about Japanese tree frogs (8.6.9) and binocular rivalry (8.1.14, 8.2.17), and to Jordi Garcia-Ojalvo for sharing his exercises about systems biology (3.7.7, 3.7.8).

In all other respects, the aims, organization, and text of the first edition have been left intact, except for a few corrections and updates here and there. Thanks to all the teachers and students who wrote in with suggestions.

It has been a pleasure to work with Sue Caulfield, Priscilla McGeehon, and Cathleen Tetro at Westview Press. Many thanks for your guidance and attention to detail.

Finally, all my love goes out to my wife Carole, daughters Leah and Jo, and dog Murray, for putting up with my distracted air and making me laugh.

Steven H. Strogatz
Ithaca, New York
2014

PREFACE TO THE FIRST EDITION

This textbook is aimed at newcomers to nonlinear dynamics and chaos, especially students taking a first course in the subject. It is based on a one-semester course I've taught for the past several years at MIT. My goal is to explain the mathematics as clearly as possible, and to show how it can be used to understand some of the wonders of the nonlinear world.

The mathematical treatment is friendly and informal, but still careful. Analytical methods, concrete examples, and geometric intuition are stressed. The theory is developed systematically, starting with first-order differential equations and their bifurcations, followed by phase plane analysis, limit cycles and their bifurcations, and culminating with the Lorenz equations, chaos, iterated maps, period doubling, renormalization, fractals, and strange attractors.

A unique feature of the book is its emphasis on applications. These include mechanical vibrations, lasers, biological rhythms, superconducting circuits, insect outbreaks, chemical oscillators, genetic control systems, chaotic water-wheels, and even a technique for using chaos to send secret messages. In each case, the scientific background is explained at an elementary level and closely integrated with the mathematical theory.

Prerequisites

The essential prerequisite is single-variable calculus, including curve-sketching, Taylor series, and separable differential equations. In a few places, multivariable calculus (partial derivatives, Jacobian matrix, divergence theorem) and linear algebra (eigenvalues and eigenvectors) are used. Fourier analysis is not assumed, and is developed where needed. Introductory physics is used throughout. Other scientific prerequisites would depend on the applications considered, but in all cases, a first course should be adequate preparation.

Possible Courses

The book could be used for several types of courses:

- A broad introduction to nonlinear dynamics, for students with no prior exposure to the subject. (This is the kind of course I have taught.) Here one goes straight through the whole book, covering the core material at the beginning of each chapter, selecting a few applications to discuss in depth and giving light treatment to the more advanced theoretical topics or skipping them altogether. A reasonable schedule is seven weeks on Chapters 1–8, and five or six weeks on Chapters 9–12. Make sure there's enough time left in the semester to get to chaos, maps, and fractals.
- A traditional course on nonlinear ordinary differential equations, but with more emphasis on applications and less on perturbation theory than usual. Such a course would focus on Chapters 1–8.
- A modern course on bifurcations, chaos, fractals, and their applications, for students who have already been exposed to phase plane analysis. Topics would be selected mainly from Chapters 3–4, and 8–12.

For any of these courses, the students should be assigned homework from the exercises at the end of each chapter. They could also do computer projects; build chaotic circuits and mechanical systems; or look up some of the references to get a taste of current research. This can be an exciting course to teach, as well as to take. I hope you enjoy it.

Conventions

Equations are numbered consecutively within each section. For instance, when we're working in Section 5.4, the third equation is called (3) or Equation (3), but elsewhere it is called (5.4.3) or Equation (5.4.3). Figures, examples, and exercises are always called by their full names, e.g., Exercise 1.2.3. Examples and proofs end with a loud thump, denoted by the symbol ■.

Acknowledgments

Thanks to the National Science Foundation for financial support. For help with the book, thanks to Diana Dabby, Partha Saha, and Shinya Watanabe (students); Jihad Touma and Rodney Worthing (teaching assistants); Andy Christian, Jim Crutchfield, Kevin Cuomo, Frank DeSimone, Roger Eckhardt, Dana Hobson, and Thanos Siapas (for providing figures); Bob Devaney, Irv Epstein, Danny Kaplan, Willem Malkus, Charlie Marcus, Paul Matthews, Arthur Mattuck, Rennie Mirollo, Peter Renz, Dan Rockmore, Gil Strang, Howard Stone, John Tyson, Kurt Wiesenfeld, Art Winfree, and Mary Lou Zeeman (friends and colleagues who gave advice); and to my editor Jack Repcheck, Lynne Reed, Production Supervisor, and all the other helpful people at Addison-Wesley. Finally, thanks to my family and Elisabeth for their love and encouragement.

Steven H. Strogatz
Cambridge, Massachusetts
1994

1

OVERVIEW

1.0 Chaos, Fractals, and Dynamics

There is a tremendous fascination today with chaos and fractals. James Gleick's book *Chaos* (Gleick 1987) was a bestseller for months—an amazing accomplishment for a book about mathematics and science. Picture books like *The Beauty of Fractals* by Peitgen and Richter (1986) can be found on coffee tables in living rooms everywhere. It seems that even nonmathematical people are captivated by the infinite patterns found in fractals (Figure 1.0.1). Perhaps most important of all, chaos and fractals represent hands-on mathematics that is alive and changing. You can turn on a home computer and create stunning mathematical images that no one has ever seen before.

The aesthetic appeal of chaos and fractals may explain why so many people have become intrigued by these ideas. But maybe you feel the urge to go deeper— to learn the mathematics behind the pictures, and to see how the ideas can be applied to problems in science and engineering. If so, this is a textbook for you.

The style of the book is informal (as you can see), with an emphasis on concrete examples and geometric thinking, rather than proofs and abstract arguments. It is also an extremely "applied" book—virtually every idea is illustrated by some application to science or engineering. In many cases, the applications are drawn from the recent research literature. Of course, one problem with such an applied approach is that not everyone is an expert in physics *and* biology *and* fluid mechanics … so the science as well as the mathematics will need to be explained from scratch. But that should be fun, and it can be instructive to see the connections among different fields.

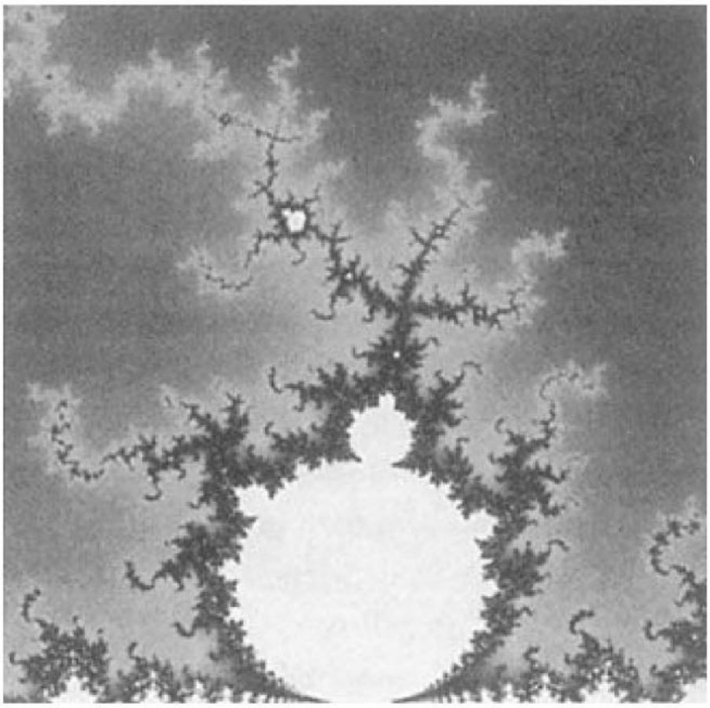

Figure 1.0.1

Before we start, we should agree about something: chaos and fractals are part of an even grander subject known as ***dynamics.*** This is the subject that deals with change, with systems that evolve in time. Whether the system in question settles down to equilibrium, keeps repeating in cycles, or does something more complicated, it is dynamics that we use to analyze the behavior. You have probably been exposed to dynamical ideas in various places—in courses in differential equations, classical mechanics, chemical kinetics, population biology, and so on. Viewed from the perspective of dynamics, all of these subjects can be placed in a common framework, as we discuss at the end of this chapter.

Our study of dynamics begins in earnest in Chapter 2. But before digging in, we present two overviews of the subject, one historical and one logical. Our treatment is intuitive; careful definitions will come later. This chapter concludes with a "dynamical view of the world," a framework that will guide our studies for the rest of the book.

1.1 Capsule History of Dynamics

Although dynamics is an interdisciplinary subject today, it was originally a branch of physics. The subject began in the mid-1600s, when Newton invented differential equations, discovered his laws of motion and universal gravitation, and combined them to explain Kepler's laws of planetary motion. Specifically, Newton solved the two-body problem—the problem of calculating the motion of the earth around the sun, given the inverse-square law of gravitational attraction between them. Subsequent generations of mathematicians and physicists tried to extend Newton's analytical methods to the three-body problem (e.g., sun, earth, and moon) but curiously this problem turned out to be much more difficult to solve. After decades of effort, it was eventually realized that the three-body problem was essentially *impossible* to solve, in the sense of obtaining explicit formulas for the motions of the three bodies. At this point the situation seemed hopeless.

The breakthrough came with the work of Poincaré in the late 1800s. He introduced a new point of view that emphasized qualitative rather than quantitative questions. For example, instead of asking for the exact positions of the planets at all times, he asked "Is the solar system stable forever, or will some planets eventually fly off to infinity?" Poincaré developed a powerful *geometric* approach to analyzing such questions. That approach has flowered into the modern subject of dynamics, with applications reaching far beyond celestial mechanics. Poincaré was also the first person to glimpse the possibility of ***chaos***, in which a deterministic system exhibits aperiodic behavior that depends sensitively on the initial conditions, thereby rendering long-term prediction impossible.

But chaos remained in the background in the first half of the twentieth century; instead dynamics was largely concerned with nonlinear oscillators and their applications in physics and engineering. Nonlinear oscillators played a vital role in the development of such technologies as radio, radar, phase-locked loops, and lasers. On the theoretical side, nonlinear oscillators also stimulated the invention of new mathematical techniques—pioneers in this area include van der Pol, Andronov, Littlewood, Cartwright, Levinson, and Smale. Meanwhile, in a separate development, Poincaré's geometric methods were being extended to yield a much deeper understanding of classical mechanics, thanks to the work of Birkhoff and later Kolmogorov, Arnol'd, and Moser.

The invention of the high-speed computer in the 1950s was a watershed in the history of dynamics. The computer allowed one to experiment with equations in a way that was impossible before, and thereby to develop some intuition about nonlinear systems. Such experiments led to Lorenz's discovery in 1963 of chaotic motion on a strange attractor. He studied a simplified model of convection rolls in the atmosphere to gain insight into the notorious unpredictability of the weather. Lorenz found that the solutions to his equations never settled down to equilibrium or to a periodic state—instead they continued to oscillate in an irregular, aperiodic fashion. Moreover, if he started his simulations from two slightly different initial conditions, the resulting behaviors would soon become totally different. The implication was that the system was *inherently* unpredictable—tiny errors in measuring the current state of the atmosphere (or any other chaotic system) would be amplified rapidly, eventually leading to embarrassing forecasts. But Lorenz also showed that there was structure in the chaos—when plotted in three dimensions, the solutions to his equations fell onto a butterfly-shaped set of points (Figure 1.1.1). He argued that this set had to be "an infinite complex of surfaces"— today we would regard it as an example of a fractal.

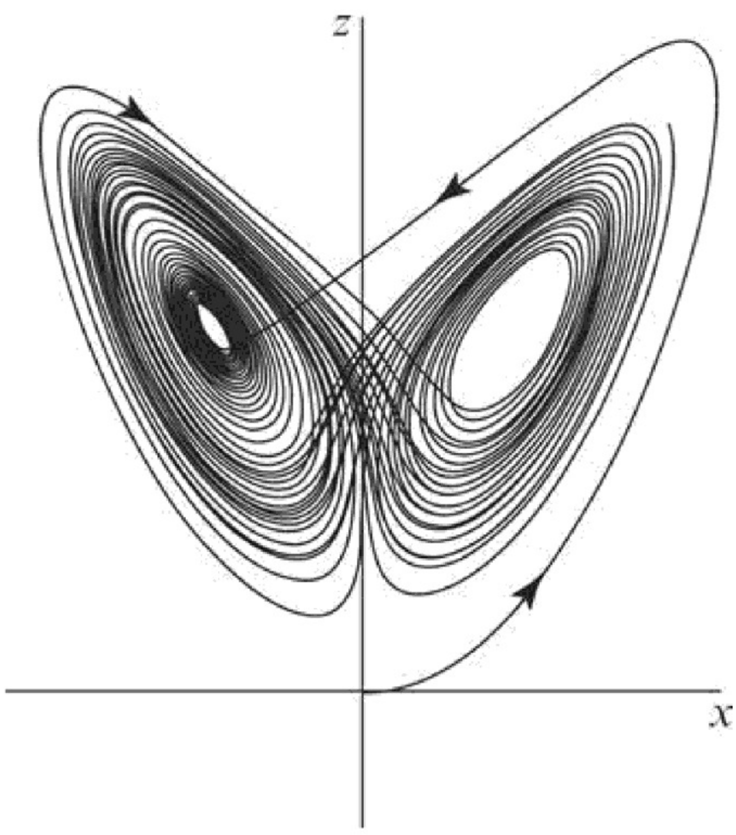

Figure 1.1.1

Lorenz's work had little impact until the 1970s, the boom years for chaos. Here are some of the main developments of that glorious decade. In 1971, Ruelle and Takens proposed a new theory for the onset of turbulence in fluids, based on abstract considerations about strange attractors. A few years later, May found examples of chaos in iterated mappings arising in population biology, and wrote an influential review article that stressed the pedagogical importance of studying simple nonlinear systems, to counterbalance the often misleading linear intuition fostered by traditional education. Next came the most surprising discovery of all, due to the physicist Feigenbaum. He discovered that there are certain universal laws governing the transition from regular to chaotic behavior; roughly speaking, completely different systems can go chaotic in the same way. His work established a link between chaos and phase transitions, and enticed a generation of physicists to the study of dynamics. Finally, experimentalists such as Gollub, Libchaber, Swinney, Linsay, Moon, and Westervelt tested the new ideas about chaos in experiments on fluids, chemical reactions, electronic circuits, mechanical oscillators, and semiconductors.

Although chaos stole the spotlight, there were two other major developments in dynamics in the 1970s. Mandelbrot codified and popularized fractals, produced magnificent computer graphics of them, and showed how they could be applied in a variety of subjects. And in the emerging area of mathematical biology, Winfree applied the geometric methods of dynamics to biological oscillations, especially circadian (roughly 24-hour) rhythms and heart rhythms.

By the 1980s many people were working on dynamics, with contributions too numerous to list. Table 1.1.1 summarizes this history.

1.2 The Importance of Being Nonlinear

Now we turn from history to the logical structure of dynamics. First we need to introduce some terminology and make some distinctions.

There are two main types of dynamical systems: ***differential equations*** and ***iterated maps*** (also known as difference equations). Differential equations describe the evolution of systems in continuous time, whereas iterated maps arise in problems where time is discrete. Differential equations are used much more widely in science and engineering, and we shall therefore concentrate on them. Later in the book we will see that iterated maps can also be very useful, both for providing simple examples of chaos, and also as tools for analyzing periodic or chaotic solutions of differential equations.

Now confining our attention to differential equations, the main distinction is between ordinary and partial differential equations. For instance, the equation for a damped harmonic oscillator

Dynamics — A Capsule History

1666	Newton	Invention of calculus, explanation of planetary motion
1700s		Flowering of calculus and classical mechanics
1800s		Analytical studies of planetary motion
1890s	Poincaré	Geometric approach, nightmares of chaos
1920–1950		Nonlinear oscillators in physics and engineering, invention of radio, radar, laser
1920–1960	Birkhoff	Complex behavior in Hamiltonian mechanics
	Kolmogorov	
	Arnol'd	
	Moser	
1963	Lorenz	Strange attractor in simple model of convection
1970s	Ruelle & Takens	Turbulence and chaos

	May	Chaos in logistic map
	Feigenbaum	Universality and renormalization, connection between chaos and phase transitions
		Experimental studies of chaos
	Winfree	Nonlinear oscillators in biology
	Mandelbrot	Fractals
1980s		Widespread interest in chaos, fractals, oscillators, and their applications

Table 1.1.1

$$m\frac{d^2x}{dt^2}+b\frac{dx}{dt}+kx=0 \tag{1}$$

is an ordinary differential equation, because it involves only ordinary derivatives dx/dt and d^2x/dt^2. That is, there is only one independent variable, the time t. In contrast, the heat equation

$$\frac{\partial u}{\partial t}=\frac{\partial^2 u}{\partial x^2}$$

is a partial differential equation—it has both time t and space x as independent variables. Our concern in this book is with purely temporal behavior, and so we deal with ordinary differential equations almost exclusively.

A very general framework for ordinary differential equations is provided by the system

$$\dot{x}_1=f_1(x_1,\ldots,x_n) \quad \vdots \quad \dot{x}_n=f_n(x_1,\ldots,x_n). \tag{2}$$

Here the overdots denote differentiation with respect to t. Thus $_i=dx_i/dt$. The variables $x_1, \ldots, x_n$ might represent concentrations of chemicals in a reactor, populations of different species in an ecosystem, or the positions and velocities of the planets in the solar system. The functions $f_1, \ldots, f_n$ are determined by the problem at hand.

For example, the damped oscillator (1) can be rewritten in the form of (2), thanks to the following trick: we introduce new variables $x_1 = x$ and $x_2=$. Then $_1 = x_2$ from the definitions, and

$$\ddot{x}_2=\ddot{x}=-\frac{b}{m}\dot{x}-\frac{k}{m}x=-\frac{b}{m}x_2-\frac{k}{m}x_1$$

from the definitions and the governing equation (1). Hence the equivalent system (2) is

$$\dot{x}_1=x_2 \quad \dot{x}_2=-\frac{b}{m}x_2-\frac{k}{m}x_1.$$

This system is said to be ***linear,*** because all the x_i on the right-hand side appear to the first power only. Otherwise the system would be ***nonlinear.*** Typical nonlinear terms are products, powers, and functions of the x_i, such as $x_1 x_2$, $(x_1)^3$, or $\cos x_2$.

For example, the swinging of a pendulum is governed by the equation

$$\ddot{x}+\frac{g}{L}\sin x=0,$$

where x is the angle of the pendulum from vertical, g is the acceleration due to gravity, and L is the length of the pendulum. The equivalent system is nonlinear:

$$\dot{x}_1=x_2 \quad \dot{x}_2=-\frac{g}{L}\sin x_1.$$

Nonlinearity makes the pendulum equation very difficult to solve analytically. The usual way around this is to fudge, by invoking the small angle approximation $\sin x \approx x$ for $x \ll 1$. This converts the problem to a linear one, which can then be solved easily. But by restricting to small x, we're throwing out some of the physics, like motions where the pendulum whirls over the top. Is it really necessary to make such drastic approximations?

It turns out that the pendulum equation *can* be solved analytically, in terms of elliptic functions. But there ought to be an easier way. After all, the motion of the pendulum is simple: at low energy, it swings back and forth, and at high energy it whirls over the top. There should be some way of extracting this information from the system directly. This is the sort of problem we'll learn how to solve, using geometric methods.

Here's the rough idea. Suppose we happen to know a solution to the pendulum system, for a particular initial condition. This solution would be a pair of functions $x_1(t)$ and $x_2(t)$, representing the position and velocity of the pendulum. If we construct an abstract space with coordinates (x_1, x_2), then the solution $(x_1(t), x_2(t))$ corresponds to a point moving along a curve in this space (Figure 1.2.1).

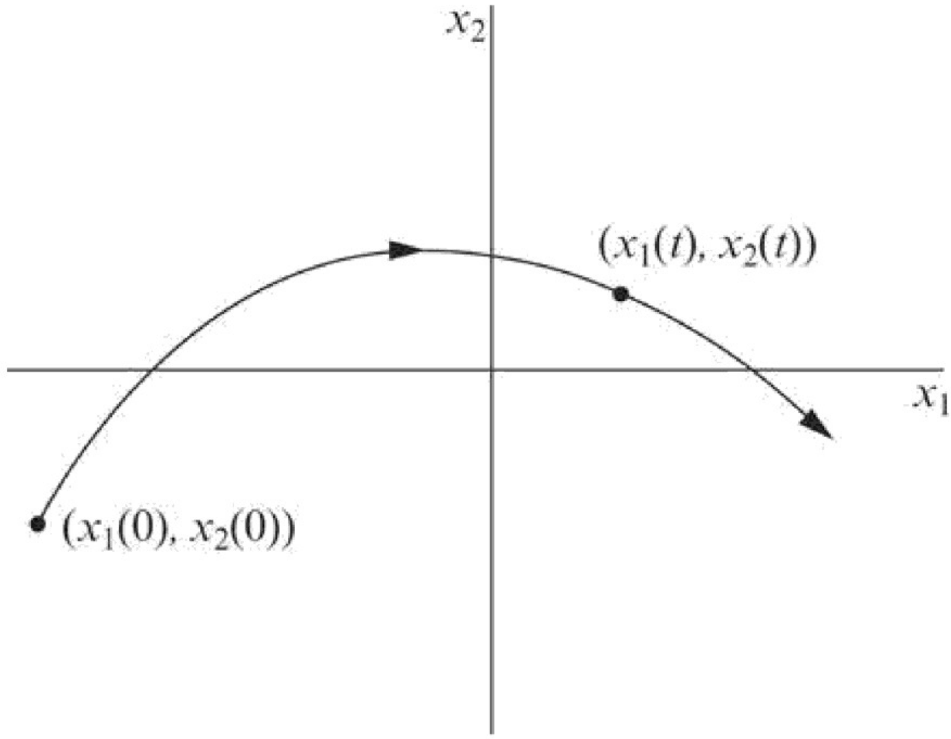

Figure 1.2.1

This curve is called a ***trajectory,*** and the space is called the ***phase space*** for the system. The phase space is completely filled with trajectories, since each point can serve as an initial condition.

Our goal is to run this construction *in reverse:* given the system, we want to draw the trajectories, and thereby extract information about the solutions. In many cases, geometric reasoning will allow us to draw the trajectories *without actually solving the system!*

Some terminology: the phase space for the general system (2) is the space with coordinates $x_1, \ldots, x_n$. Because this space is n-dimensional, we will refer to (2) as an ***n-dimensional system*** or an ***nth-order*** system. Thus n represents the dimension of the phase space.

Nonautonomous Systems

You might worry that (2) is not general enough because it doesn't include any explicit *time dependence.* How do we deal with time-dependent or ***nonautonomous*** equations like the forced harmonic oscillator m $+ b$ $+ kx + F \cos t$? In this case too there's an easy trick that allows us to rewrite the system in the form (2). We let $x_1 = x$ and $x_2 =$ as before but now we introduce $x_3 = t$. Then $_3 = 1$

and so the equivalent system is

$$\dot{x}_1 = x_2 \quad \dot{x}_2 = \frac{1}{m}(-kx_1 - bx_2 + F\cos x_3) \quad \dot{x}_3 = 1 \tag{3}$$

which is an example of a *three*-dimensional system. Similarly, an *n*th-order time-dependent equation is a special case of an $(n + 1)$-dimensional system. By this trick, we can always remove any time dependence by adding an extra dimension to the system.

The virtue of this change of variables is that it allows us to visualize a phase space with trajectories *frozen* in it. Otherwise, if we allowed explicit time dependence, the vectors and the trajectories would always be wiggling—this would ruin the geometric picture we're trying to build. A more physical motivation is that the ***state*** of the forced harmonic oscillator is truly three-dimensional: we need to know three numbers, *x*, , and *t*, to predict the future, given the present. So a three-dimensional phase space is natural.

The cost, however, is that some of our terminology is nontraditional. For example, the forced harmonic oscillator would traditionally be regarded as a second-order linear equation, whereas we will regard it as a third-order nonlinear system, since (3) is nonlinear, thanks to the cosine term. As we'll see later in the book, forced oscillators have many of the properties associated with nonlinear systems, and so there are genuine conceptual advantages to our choice of language.

Why Are Nonlinear Problems So Hard?

As we've mentioned earlier, most nonlinear systems are impossible to solve analytically. Why are nonlinear systems so much harder to analyze than linear ones? The essential difference is that *linear systems can be broken down into parts.* Then each part can be solved separately and finally recombined to get the answer. This idea allows a fantastic simplification of complex problems, and underlies such methods as normal modes, Laplace transforms, superposition arguments, and Fourier analysis. In this sense, a linear system is precisely equal to the sum of its parts.

But many things in nature don't act this way. Whenever parts of a system interfere, or cooperate, or compete, there are nonlinear interactions going on. Most of everyday life is nonlinear, and the principle of superposition fails spectacularly. If you listen to your two favorite songs at the same time, you won't get double the pleasure! Within the realm of physics, nonlinearity is vital to the operation of a laser, the formation of turbulence in a fluid, and the superconductivity of Josephson junctions.

1.3 A Dynamical View of the World

Now that we have established the ideas of nonlinearity and phase space, we can present a framework for dynamics and its applications. Our goal is to show the logical structure of the entire subject. The framework presented in Figure 1.3.1 will guide our studies thoughout this book.

The framework has two axes. One axis tells us the number of variables needed to characterize the state of the system. Equivalently, this number is the *dimension of the phase space.* The other axis tells us whether the system is linear or *nonlinear.*

For example, consider the exponential growth of a population of organisms. This system is described by the first-order differential equation $= rx$ where x is the population at time t and r is the growth rate. We place this system in the column labeled "$n = 1$" because *one* piece of information—the current value of the population x—is sufficient to predict the population at any later time. The system is also classified as linear because the differential equation $= rx$ is linear in x.

As a second example, consider the swinging of a pendulum, governed by

$$\ddot{x} + \frac{g}{L}\sin x = 0,$$

In contrast to the previous example, the state of this system is given by *two* variables: its current angle x and angular velocity . (Think of it this way: we need the initial values of both x and to determine the solution uniquely. For example, if we knew only x, we wouldn't know which way the pendulum was swinging.) Because two variables are needed to specify the state, the pendulum belongs in the $n = 2$ column of Figure 1.3.1. Moreover, the system is nonlinear, as discussed in the previous section. Hence the pendulum is in the lower, nonlinear half of the $n = 2$ column.

One can continue to classify systems in this way, and the result will be something like the framework shown here. Admittedly, some aspects of the picture are debatable. You might think that some topics should be added, or placed differently, or even that more axes are needed—the point is to think about classifying systems on the basis of their dynamics.

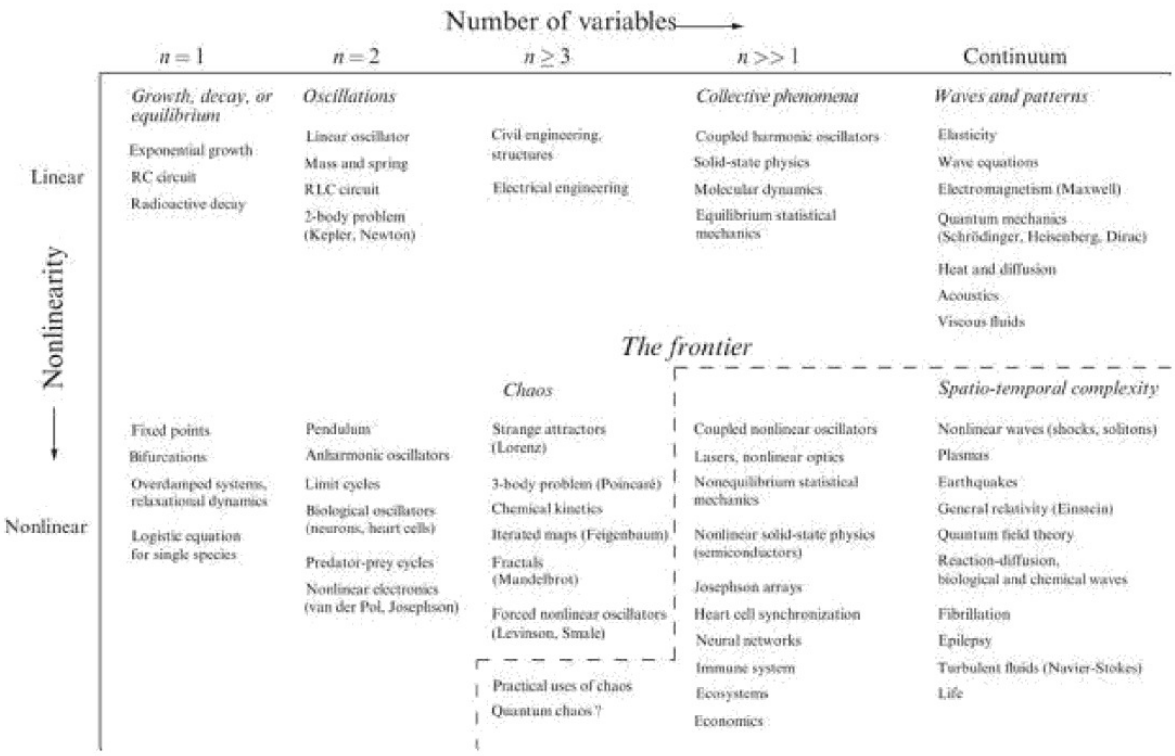

Figure 1.3.1

There are some striking patterns in Figure 1.3.1. All the simplest systems occur in the upper left-hand corner. These are the small linear systems that we learn about in the first few years of college. Roughly speaking, these linear systems exhibit growth, decay, or equilibrium when $n = 1$, or oscillations when $n = 2$. The italicized phrases in Figure 1.3.1 indicate that these broad classes of phenomena first arise in this part of the diagram. For example, an *RC* circuit has $n = 1$ and cannot oscillate, whereas an *RLC* circuit has $n = 2$ and can oscillate.

The next most familiar part of the picture is the upper right-hand corner. This is the domain of classical applied mathematics and mathematical physics where the linear partial differential equations live. Here we find Maxwell's equations of electricity and magnetism, the heat equation, Schrödinger's wave equation in quantum mechanics, and so on. These partial differential equations involve an infinite "continuum" of variables because each point in space contributes additional degrees of freedom. Even though these systems are large, they are tractable, thanks to such linear techniques as Fourier analysis and transform methods.

In contrast, the lower half of Figure 1.3.1—the nonlinear half—is often ignored or deferred to later courses. But no more! In this book we start in the lower left corner and systematically head to the right.

As we increase the phase space dimension from $n = 1$ to $n = 3$, we encounter new phenomena at every step, from fixed points and bifurcations when $n = 1$, to nonlinear oscillations when $n = 2$, and finally chaos and fractals when $n = 3$. In all cases, a geometric approach proves to be very powerful, and gives us most of the information we want, even though we usually can't solve the equations in the traditional sense of finding a formula for the answer. Our journey will also take us to some of the most exciting parts of modern science, such as mathematical biology and condensed-matter physics.

You'll notice that the framework also contains a region forbiddingly marked "The frontier." It's like in those old maps of the world, where the mapmakers wrote, "Here be dragons" on the unexplored parts of the globe. These topics are not completely unexplored, of course, but it is fair to say that they lie at the limits of current understanding. The problems are very hard, because they are both large and nonlinear. The resulting behavior is typically complicated in *both space and time*, as in the motion of a turbulent fluid or the patterns of electrical activity in a fibrillating heart. Toward the end of the book we will touch on some of these problems—they will certainly pose challenges for years to come.

PART I

ONE-DIMENSIONAL FLOWS

2

FLOWS ON THE LINE

2.0 Introduction

In Chapter 1, we introduced the general system

$$\dot{x}_1 = f_1(x_1,\ldots,x_n) \quad \vdots \quad \dot{x}_n = f_n(x_1,\ldots,x_n)$$

and mentioned that its solutions could be visualized as trajectories flowing through an n-dimensional phase space with coordinates $(x_1, \ldots, x_n)$. At the moment, this idea probably strikes you as a mind-bending abstraction. So let's start slowly, beginning here on earth with the simple case $n = 1$. Then we get a single equation of the form

$$\dot{x} = f(x).$$

Here $x(t)$ is a real-valued function of time t, and $f(x)$ is a smooth real-valued function of x. We'll call such equations ***one-dimensional*** or ***first-order systems.***

Before there's any chance of confusion, let's dispense with two fussy points of terminology:

1. The word *system* is being used here in the sense of a dynamical system, not in the classical sense of a collection of two or more equations. Thus a single equation can be a "system."
2. We do not allow f to depend explicitly on time. Time-dependent or "nonautonomous" equations of the form $\dot{x} = f(x, t)$ are more complicated, because one needs *two* pieces of information, x and t, to predict the future state of the system. Thus $\dot{x} = f(x, t)$ should really be regarded as a *two-dimensional* or *second-order* system, and will therefore be discussed later in the book.

2.1 A Geometric Way of Thinking

Pictures are often more helpful than formulas for analyzing nonlinear systems. Here we illustrate this point by a simple example. Along the way we will introduce one of the most basic techniques of dynamics: *interpreting a differential equation as a vector field.*

Consider the following nonlinear differential equation:

$$\dot{x} = \sin x. \tag{1}$$

To emphasize our point about formulas versus pictures, we have chosen one of the few nonlinear equations that can be solved in closed form. We separate the variables and then integrate:

dt=dxsinx,

which implies

t=∫cscx dx=−lnl cscx+cotx l+C.

To evaluate the constant C, suppose that $x = x_0$ at $t = 0$. Then C=lnl cscx0+cotx0 l. Hence the solution is

t=lnl cscx0+cotx0cscx+cotx l. *(2)*

This result is exact, but a headache to interpret. For example, can you answer the following questions?

1. Suppose $x_0 = \pi/4$; describe the qualitative features of the solution $x(t)$ for all $t > 0$. In particular, what happens as $t \to \infty$?
2. For an *arbitrary* initial condition x_0, what is the behavior of $x(t)$ as $t \to \infty$?

Think about these questions for a while, to see that formula (2) is not transparent.

In contrast, a graphical analysis of (1) is clear and simple, as shown in Figure 2.1.1. We think of t as time, x as the position of an imaginary particle moving along the real line, and as the velocity of that particle. Then the differential equation = sin x represents a *vector field* on the line: it dictates the velocity vector at each x. To sketch the vector field, it is convenient to plot versus x, and then draw arrows on the x-axis to indicate the corresponding velocity vector at each x. The arrows point to the right when > 0 and to the left when < 0.

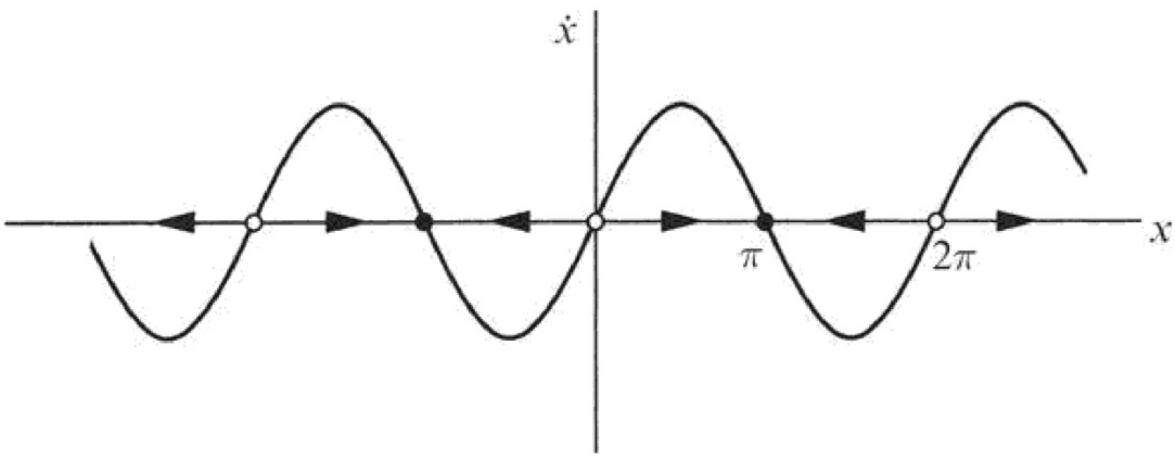

Figure 2.1.1

Here's a more physical way to think about the vector field: imagine that fluid is flowing steadily along the x-axis with a velocity that varies from place to place, according to the rule = sin x. As shown
$\dot{x}$ in Figure 2.1.1, the ***flow*** is to the right when > 0 and to the left when < 0. At points where = 0, there is no flow; such points are therefore called ***fixed points.*** You can see that there are two kinds of fixed points in Figure 2.1.1: solid black dots represent ***stable*** fixed points (often called *attractors* or *sinks*, because the flow is toward them) and open circles represent ***unstable*** fixed points (also known as *repellers* or *sources*).

Armed with this picture, we can now easily understand the solutions to the differential equation = sin x. We just start our imaginary particle at x_0 and watch how it is carried along by the flow.

This approach allows us to answer the questions above as follows:

1. Figure 2.1.1 shows that a particle starting at $x_0 = \pi/4$ moves to the right faster and faster until it crosses $x = \pi/2$ (where $\sin x$ reaches its maximum). Then the particle starts slowing down and eventually approaches the stable fixed point $x = \pi$ from the left. Thus, the qualitative form of the solution is as shown in Figure 2.1.2.

   Note that the curve is concave up at first, and then concave down; this corresponds to the initial acceleration for $x < \pi/2$, followed by the deceleration toward $x = \pi$.
2. The same reasoning applies to any initial condition x_0. Figure 2.1.1 shows that if $x > 0$ initially, the particle heads to the right and asymptotically approaches the nearest stable fixed point. Similarly, if < 0 initially, the particle approaches the nearest stable fixed point to its left. If $= 0$, then x remains constant. The qualitative form of the solution for any initial condition is sketched in Figure 2.1.3.

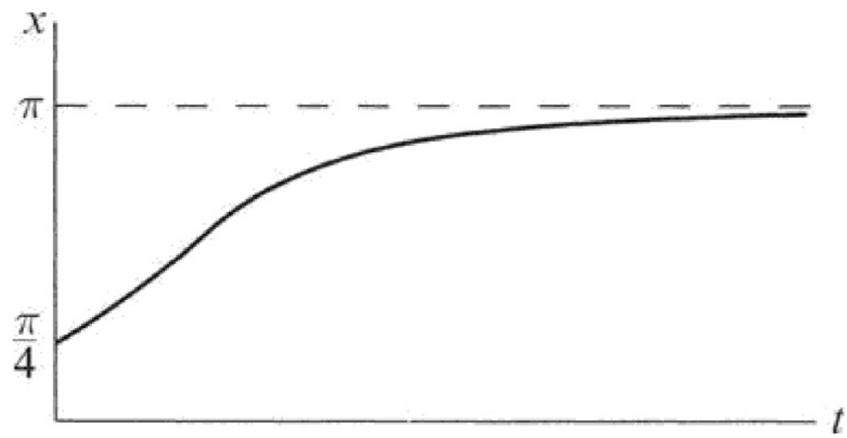

Figure 2.1.2

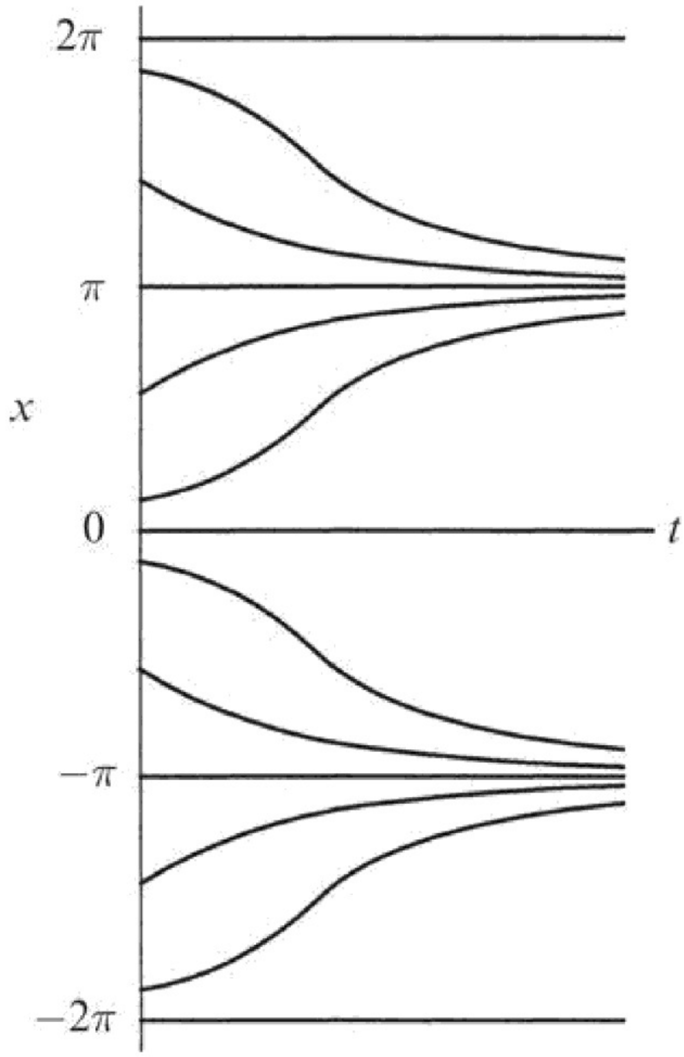

Figure 2.1.3

In all honesty, we should admit that a picture can't tell us certain *quantitative* things: for instance, we don't know the time at which the speed | | is greatest. But in many cases *qualitative* information is what we care about, and then pictures are fine.

2.2 Fixed Points and Stability

The ideas developed in the last section can be extended to any one-dimensional system $x = f(x)$. We just need to draw the graph of $f(x)$ and then use it to sketch the vector field on the real line (the x-axis in Figure 2.2.1).

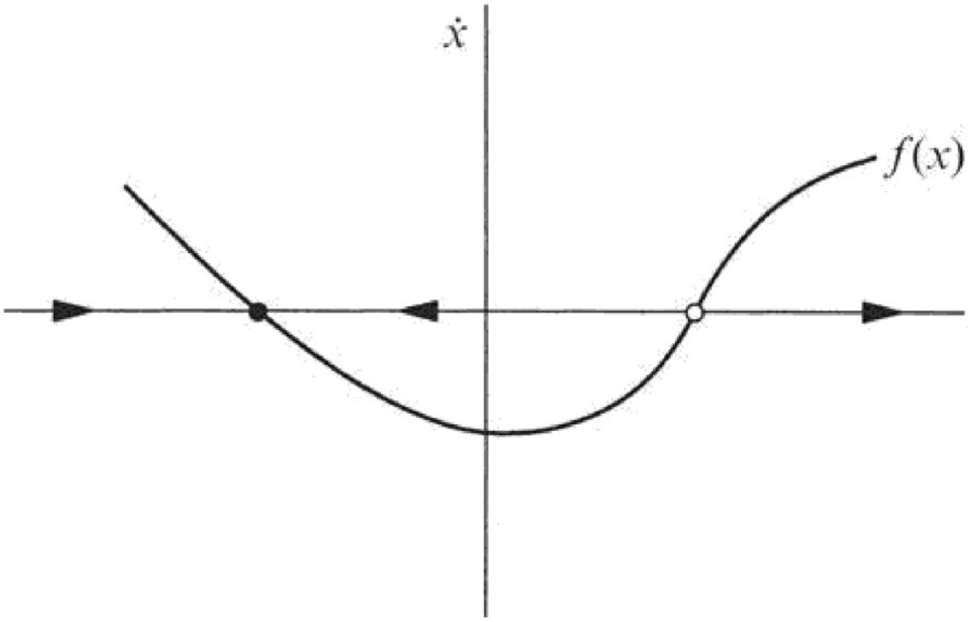

Figure 2.2.1

As before, we imagine that a fluid is flowing along the real line with a local velocity $f(x)$. This imaginary fluid is called the phase fluid, and the real line is the phase space. The flow is to the right where $f(x) > 0$ and to the left where $f(x) < 0$. To find the solution to $x = f(x)$ starting from an arbitrary initial condition x_0, we place an imaginary particle (known as a ***phase point***) at x_0 and watch how it is carried along by the flow. As time goes on, the phase point moves along the x-axis according to some function $x(t)$. This function is called the ***trajectory*** based at x_0, and it represents the solution of the differential equation starting from the initial condition x_0. A picture like Figure 2.2.1, which shows all the qualitatively different trajectories of the system, is called a ***phase portrait.***

The appearance of the phase portrait is controlled by the fixed points x^*, defined by $f(x^*) = 0$; they correspond to stagnation points of the flow. In Figure 2.2.1, the solid black dot is a stable fixed point (the local flow is toward it) and the open dot is an unstable fixed point (the flow is away from it).

In terms of the original differential equation, fixed points represent ***equilibrium*** solutions (sometimes called steady, constant, or rest solutions, since if $x = x^*$ initially, then $x(t) = x^*$ for all time). An equilibrium is defined to be stable if all sufficiently small disturbances away from it damp out in time. Thus stable equilibria are represented geometrically by stable fixed points. Conversely, unstable equilibria, in which disturbances grow in time, are represented by unstable fixed points.

EXAMPLE 2.2.1:

Find all fixed points for $= x^2 - 1$, and classify their stability.

Solution: Here $f(x) = x^2 - 1$. To find the fixed points, we set $f(x^*) = 0$ and solve for x^*. Thus $x^* = \pm 1$. To determine stability, we plot $x^2 - 1$ and then sketch the vector field (Figure 2.2.2). The flow is to the right where $x^2 - 1 > 0$ and to the left where $x^2 - 1 < 0$. Thus $x^* = -1$ is stable, and $x^* = 1$ is unstable.

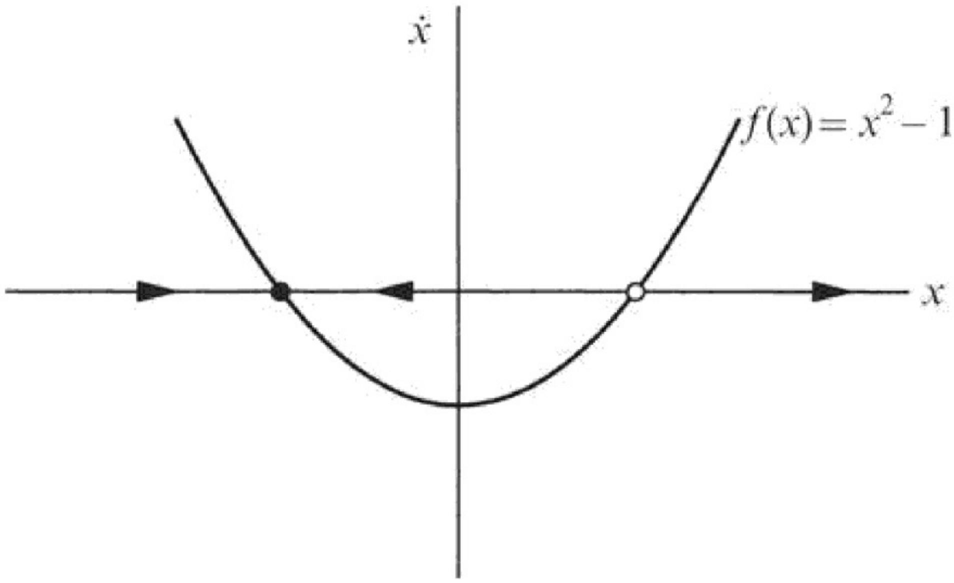

Figure 2.2.2

Note that the definition of stable equilibrium is based on *small* disturbances; certain large disturbances may fail to decay. In Example 2.2.1, all small disturbances to $x^* = -1$ will decay, but a large disturbance that sends x to the right of $x = 1$ will *not* decay—in fact, the phase point will be repelled out to $+\infty$. To emphasize this aspect of stability, we sometimes say that $x^* = -1$ is *locally stable*, but not globally stable.

EXAMPLE 2.2.2:

Consider the electrical circuit shown in Figure 2.2.3. A resistor R and a capacitor C are in series with a battery of constant dc voltage V_0 Suppose that the switch is closed at $t = 0$, and that there is no charge on the capacitor initially. Let $Q\,(t)$ denote the charge on the capacitor at time $t \geq 0$. Sketch the graph of $Q\,(t)$.

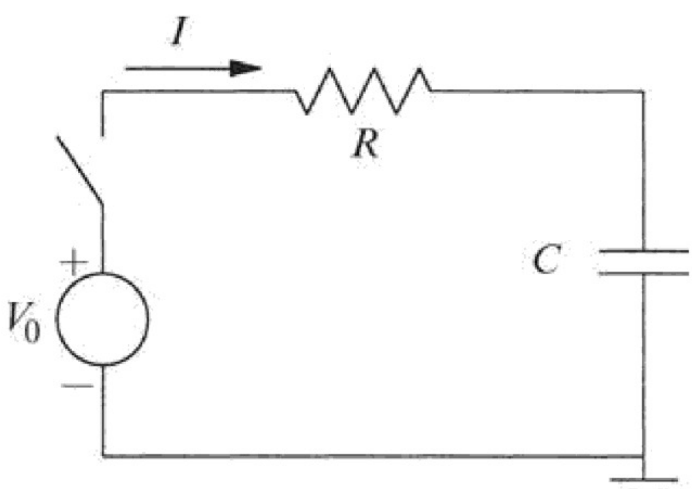

Figure 2.2.3

Solution: This type of circuit problem is probably familiar to you. It is governed by linear equations and can be solved analytically, but we prefer to illustrate the geometric approach.

First we write the circuit equations. As we go around the circuit, the total voltage drop must equal zero; hence $-V_0 + RI + Q/C = 0$, where I is the current flowing through the resistor. This current causes charge to accumulate on the capacitor at a rate $\dot{Q}=I$. Hence

$$-V0+R\dot{Q}+Q/C=0 \quad \text{or} \dot{Q}=f(Q)=V0R-QRC.$$

The graph of $f(Q)$ is a straight line with a negative slope (Figure 2.2.4). The corresponding vector field has a fixed point where $f(Q) = 0$, which occurs at $Q^* = CV_0$. The flow is to the right where $f(Q) > 0$ and to the left where $f(Q) < 0$. Thus the flow is always toward Q^*—it is a *stable* fixed point. In fact, it is ***globally stable***, in the sense that it is approached from *all* initial conditions.

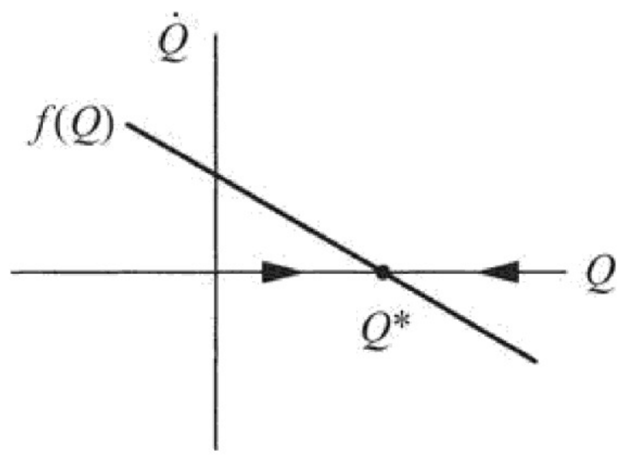

Figure 2.2.4

To sketch $Q(t)$, we start a phase point at the origin of Figure 2.2.4 and imagine how it would move. The flow carries the phase point monotonically toward Q^*. Its speed $\dot{Q}$ decreases linearly as it approaches the fixed point; therefore $Q(t)$ is increasing and concave down, as shown in Figure 2.2.5. ■

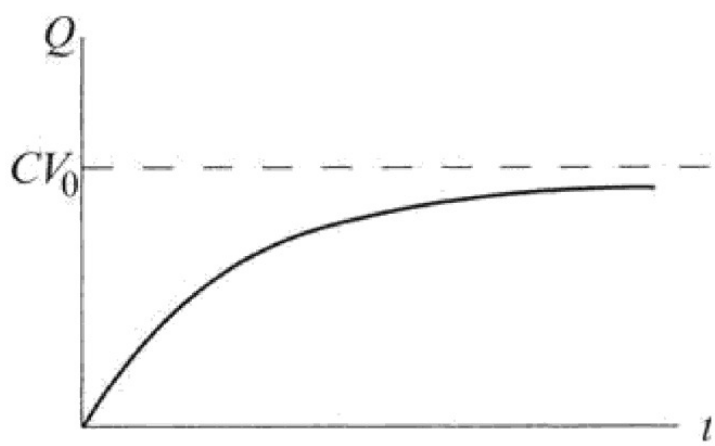

Figure 2.2.5

EXAMPLE 2.2.3:

Sketch the phase portrait corresponding to $= x - \cos x$, and determine the stability of all the fixed points.

Solution: One approach would be to plot the function $f(x) = x - \cos x$ and then sketch the associated vector field. This method is valid, but it requires you to figure out what the graph of $x - \cos x$ looks like.

There's an easier solution, which exploits the fact that we know how to graph $y = x$ and $y = \cos x$ *separately*. We plot both graphs on the same axes and then observe that they intersect in exactly one point (Figure 2.2.6).

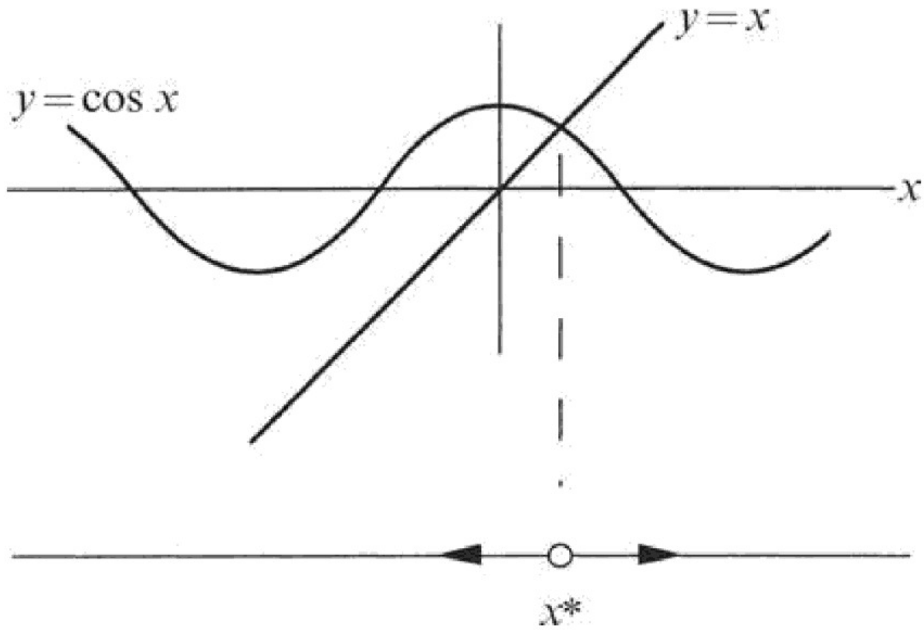

Figure 2.2.6

This intersection corresponds to a fixed point, since $x^* = \cos x^*$ and therefore $f(x^*) = 0$. Moreover, when the line lies above the cosine curve, we have $x > \cos x$ and so > 0: the flow is to the right. Similarly, the flow is to the left where the line is below the cosine curve. Hence x^* is the only fixed point, and it is unstable. Note that we can classify the stability of x^*, even though we don't have a formula for x^* itself!

2.3 Population Growth

The simplest model for the growth of a population of organisms is = rN, where $N(t)$ is the population at time t, and $r > 0$ is the growth rate. This model predicts exponential growth: $N(t) = N_0 e^{rt}$, where N_0 is the population at $t = 0$.

Of course such exponential growth cannot go on forever. To model the effects of overcrowding and limited resources, population biologists and demographers often assume that the per capita growth rate /N decreases when N becomes sufficiently large, as shown in Figure 2.3.1. For small N, the growth rate equals r, just as before. However, for populations larger than a certain ***carrying capacity*** K, the growth rate actually becomes negative; the death rate is higher than the birth rate.

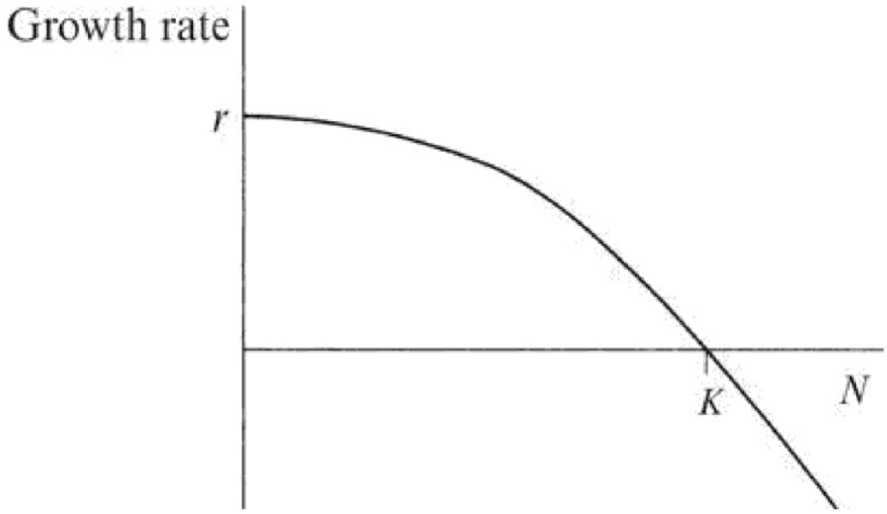

Figure 2.3.1

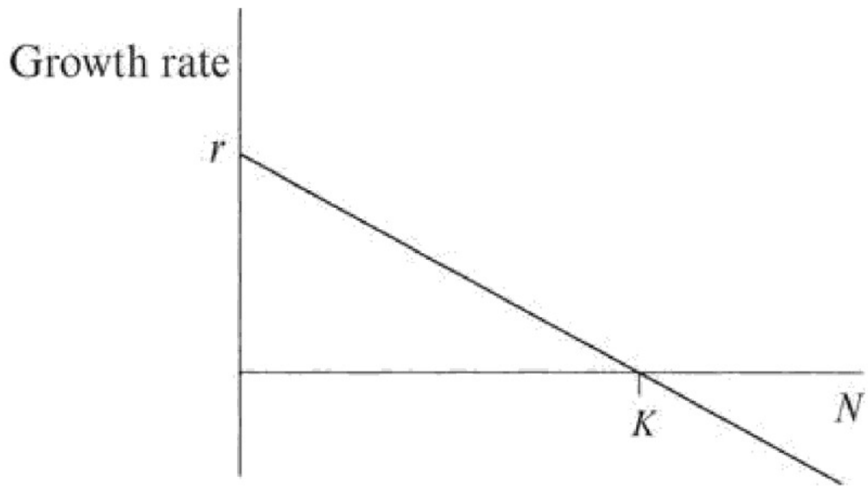

Figure 2.3.2

A mathematically convenient way to incorporate these ideas is to assume that the per capita growth rate /N decreases *linearly* with N (Figure 2.3.2).

This leads to the ***logistic equation***

$$N^{\cdot}=rN(1-NK)$$

first suggested to describe the growth of human populations by Verhulst in 1838. This equation can be solved analytically (Exercise 2.3.1) but once again we prefer a graphical approach. We plot N versus N to see what the vector field looks like. Note that we plot only $N \geq 0$, since it makes no sense to think about a negative population (Figure 2.3.3). Fixed points occur at $N^* = 0$ and $N^* = K$, as found by setting = 0 and solving for N. By looking at the flow in Figure 2.3.3, we see that $N^* = 0$ is an unstable fixed point and $N^* = K$ is a stable fixed point. In biological terms, $N = 0$ is an unstable equilibrium: a small population will grow exponentially fast and run away from $N = 0$. On the other hand, if N is disturbed slightly from K, the disturbance will decay monotonically and $N(t) \rightarrow K$ as $t \rightarrow \infty$.

In fact, Figure 2.3.3 shows that if we start a phase point at *any* $N_0 > 0$, it will always flow toward $N = K$. Hence *the population always approaches the carrying capacity.*

The only exception is if $N_0 = 0$; then there's nobody around to start reproducing, and so $N = 0$ for all time. (The model does not allow for spontaneous generation!)

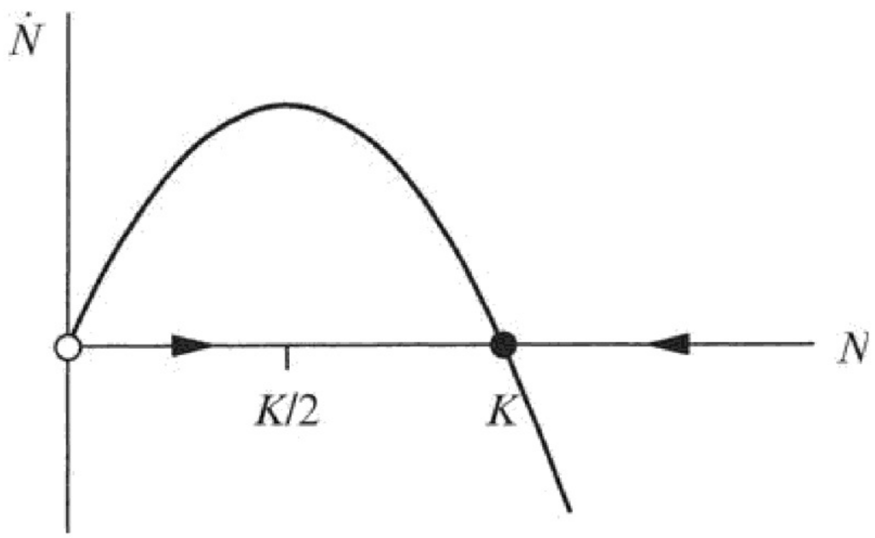

Figure 2.3.3

Figure 2.3.3 also allows us to deduce the qualitative shape of the solutions. For example, if $N_0 < K/2$, the phase point moves faster and faster until it crosses = K/ 2, where the parabola in Figure 2.3.3 reaches its maximum. Then the phase point slows down and eventually creeps toward $N = K$. In biological terms, this means that the population initially grows in an accelerating fashion, and the graph of $N(t)$ is concave up. But after $N = K/2$, the derivative begins to decrease, and so $N(t)$ is concave down as it asymptotes to the horizontal line $N = K$ (Figure 2.3.4). Thus the graph of $N(t)$ is S-shaped or *sigmoid* for $N_0 < K/2$.

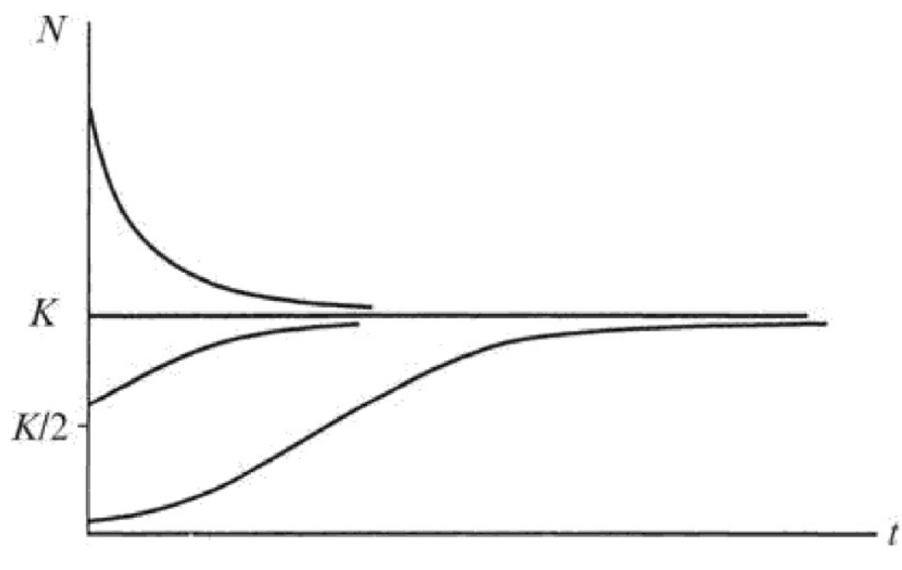

Figure 2.3.4

Something qualitatively different occurs if the initial condition N_0 lies between $K/2$ and K; now the solutions are decelerating from the start. Hence these solutions are concave down for all t. If the population initially exceeds the carrying capacity ($N_0 > K$), then $N(t)$ decreases toward $N = K$ and is concave up. Finally, if $N_0 = 0$ or $N_0 = K$, then the population stays constant.

Critique of the Logistic Model

Before leaving this example, we should make a few comments about the biological validity of the logistic equation. The algebraic form of the model is not to be taken literally. The model should really be regarded as a metaphor for populations that have a tendency to grow from zero population up to some carrying capacity K.

Originally a much stricter interpretation was proposed, and the model was argued to be a universal law of growth (Pearl 1927). The logistic equation was tested in laboratory experiments in which colonies of bacteria, yeast, or other simple organisms were grown in conditions of constant climate, food supply, and absence of predators. For a good review of this literature, see Krebs (1972, pp. 190–200). These experiments often yielded sigmoid growth curves, in some cases with an impressive match to the logistic predictions.

On the other hand, the agreement was much worse for fruit flies, flour beetles, and other organisms that have complex life cycles involving eggs, larvae, pupae, and adults. In these organisms, the predicted asymptotic approach to a steady carrying capacity was never observed—instead the populations exhibited large, persistent fluctuations after an initial period of logistic growth. See Krebs (1972) for a discussion of the possible causes of these fluctuations, including age structure and time-delayed effects of overcrowding in the population.

For further reading on population biology, see Pielou (1969) or May (1981). Edelstein-Keshet (1988) and Murray (2002, 2003) are excellent textbooks on mathematical biology in general.

2.4 Linear Stability Analysis

So far we have relied on graphical methods to determine the stability of fixed points. Frequently one would like to have a more quantitative measure of stability, such as the rate of decay to a stable fixed point. This sort of information may be obtained by *linearizing* about a fixed point, as we now explain.

Let x^* be a fixed point, and let $\eta(t) = x(t) - x^*$ be a small perturbation away from x^*. To see whether the perturbation grows or decays, we derive a differential equation for η. Differentiation yields

$$\dot{\eta} = \frac{d}{dt}(x - x^*) = \dot{x},$$

since x^* is constant. Thus $\dot{\eta} = \dot{x} = f(x) = f(x^* + \eta)$. Now using Taylor's expansion we obtain

$$f(x^* + \eta) = f(x^*) + \eta f'(x^*) + O(\eta^2),$$

where $O(\eta^2)$ denotes quadratically small terms in η. Finally, note that $f(x^*) = 0$ since x^* is a fixed point. Hence

$$\dot{\eta} = \eta f'(x^*) + O(\eta^2).$$

Now if $f'(x^*) \neq 0$, the $O(\eta^2)$ terms are negligible and we may write the approximation

$$\dot{\eta} \approx \eta f'(x^*).$$

This is a linear equation in η, and is called the ***linearization about*** x^*. It shows that *the perturbation* $\eta(t)$ *grows exponentially if* $f'(x^*) > 0$ *and decays if* $f'(x^*) < 0$. If $f'(x^*) = 0$, the $O(\eta^2)$ terms are not negligible and a nonlinear analysis is needed to determine stability, as discussed in Example 2.4.3 below.

The upshot is that the slope $f'(x^*)$ at the fixed point determines its stability. If you look back at the earlier examples, you'll see that the slope was always negative at a stable fixed point. The importance of the *sign* of $f'(x^*)$ was clear from our graphical approach; the new feature is that now we have a measure of *how* stable a fixed point is—that's determined by the *magnitude* of $f'(x^*)$. This magnitude plays the role of an exponential growth or decay rate. Its reciprocal $1/|f'(x^*)|$ is a ***characteristic time scale;*** it determines the time required for $x(t)$ to vary significantly in the neighborhood of x^*.

EXAMPLE 2.4.1:

Using linear stability analysis, determine the stability of the fixed points for $= \sin x$.

Solution: The fixed points occur where $f(x) = \sin x = 0$. Thus $x^* = k\pi$, where k is an integer. Then

$$f'(x^*) = \cos k\pi = \begin{cases} 1, & k \text{ even} \\ -1, & k \text{ odd.} \end{cases}$$

Hence x^* is unstable if k is even and stable if k is odd. This agrees with the results shown in Figure 2.1.1.

EXAMPLE 2.4.2:

Classify the fixed points of the logistic equation, using linear stability analysis, and find the characteristic time scale in each case.

Solution: Here $f(N) = rN(1 - \frac{N}{K})$, with fixed points $N^* = 0$ and $N^* = K$. Then $f'(N) = r - \frac{2rN}{K}$ and so $f'(0) = r$ and $f'(K) = -r$. Hence $N^* = 0$ is unstable and $N^* = K$ is stable, as found earlier by graphical arguments. In either case, the characteristic time scale is $1/|f'(N^*)| = 1/r$.

EXAMPLE 2.4.3:

What can be said about the stability of a fixed point when $f'(x^*) = 0$?

Solution: Nothing can be said in general. The stability is best determined on a case-by-case basis, using graphical methods. Consider the following examples:

$$\text{(a) } \dot{x} = -x^3 \quad \text{(b) } \dot{x} = x^3 \quad \text{(c) } \dot{x} = x^2 \quad \text{(d) } \dot{x} = 0$$

Each of these systems has a fixed point $x^* = 0$ with $f'(x^*) = 0$. However the stability is different in each case. Figure 2.4.1 shows that (a) is stable and (b) is unstable. Case (c) is a hybrid case we'll call ***half-stable***, since the fixed point is attracting from the left and repelling from the right. We therefore indicate this type of fixed point by a half-filled circle. Case (d) is a whole line of fixed points; perturbations neither grow nor decay.

These examples may seem artificial, but we will see that they arise naturally in the context of *bifurcations*—more about that later. ■

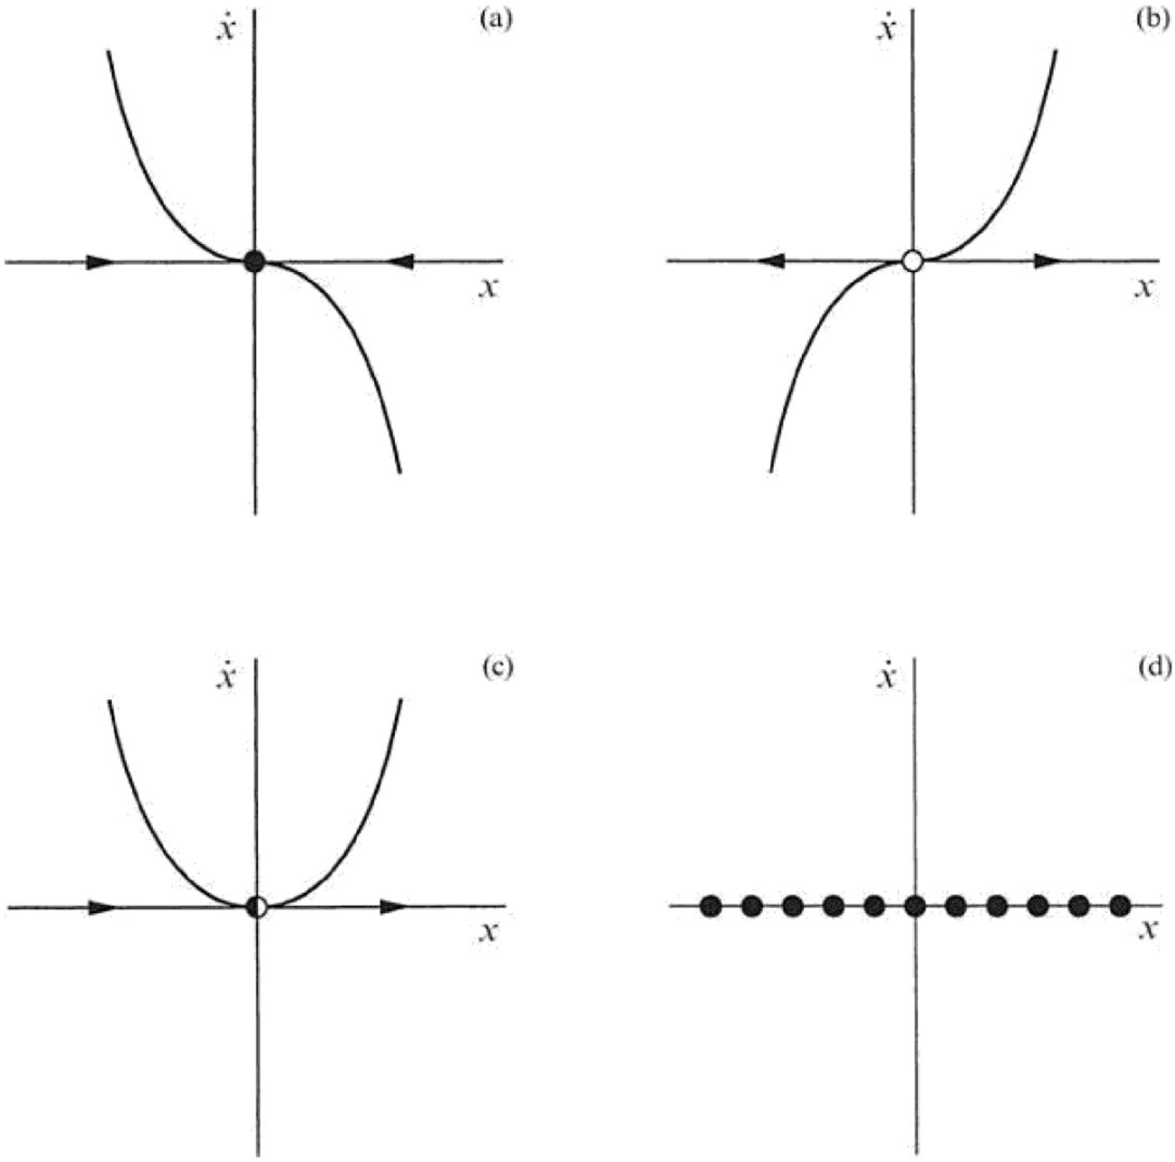

Figure 2.4

2.5 Existence and Uniqueness

Our treatment of vector fields has been very informal. In particular, we have taken a cavalier attitude toward questions of existence and uniqueness of solutions to the system $= f(x)$. That's in keeping with the "applied" spirit of this book. Nevertheless, we should be aware of what can go wrong in pathological cases.

EXAMPLE 2.5.1:

Show that the solution to $= x^{1/3}$ starting from $x_0 = 0$ is not unique.

Solution: The point $x = 0$ is a fixed point, so one obvious solution is $x(t) = 0$ for all t. The surprising fact is that there is *another* solution. To find it we separate variables and integrate:

∫x−1/3dx=∫dt

so 32x2/3=t+C. Imposing the initial condition $x(0) = 0$ yields $C = 0$. Hence x(t)=(23t)3/2 is also a solution! ■

When uniqueness fails, our geometric approach collapses because the phase point doesn't know how to move; if a phase point were started at the origin, would it stay there or would it move according to x(t)=(23t)3/2? (Or as my friends in elementary school used to say when discussing the problem of the irresistible force and the immovable object, perhaps the phase point would explode!)

Actually, the situation in Example 2.5.1 is even worse than we've let on—there are *infinitely* many solutions starting from the same initial condition (Exercise 2.5.4).

What's the source of the non-uniqueness? A hint comes from looking at the vector field (Figure 2.5.1). We see that the fixed point $x^* = 0$ is *very* unstable—the slope f' (0) is infinite.

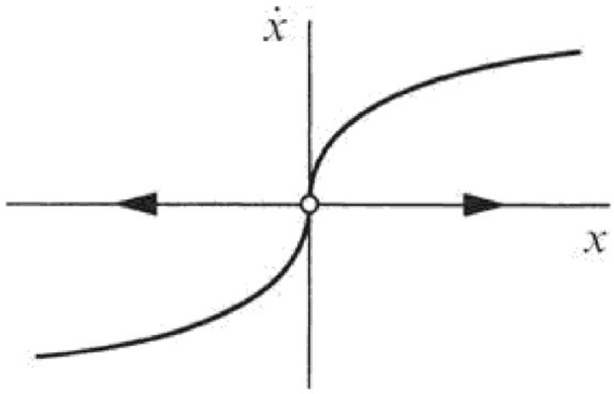

Figure 2.5.1

Chastened by this example, we state a theorem that provides sufficient conditions for existence and uniqueness of solutions to $= f(x)$.

Existence and Uniqueness Theorem: Consider the initial value problem

x˙=f(x),x(0)=x0.

Suppose that $f(x)$ and f' (x) are continuous on an open interval R of the x-axis, and suppose that x_0 is a point in R. Then the initial value problem has a solution $x(t)$ on some time interval (− ,) about $t = 0$, and the solution is unique.

For proofs of the existence and uniqueness theorem, see Borrelli and Coleman (1987), Lin and Segel (1988), or virtually any text on ordinary differential equations.

This theorem says that *if $f(x)$ is smooth enough*, then solutions exist and are unique. Even so, there's no guarantee that solutions exist forever, as shown by the next example.

EXAMPLE 2.5.2:

Discuss the existence and uniqueness of solutions to the initial value problem $= 1 + x^2$, $x(0) = x_0$. Do solutions exist for all time?

Solution: Here $f(x) = 1 + x^2$ This function is continuous and has a continuous derivative for all x. Hence the theorem tells us that solutions exist and are unique for any initial condition x_0. But *the theorem does not say that the solutions exist for all time;* they are only guaranteed to exist in a (possibly very short) time interval around $t = 0$.

For example, consider the case where $x(0) = 0$. Then the problem can be solved analytically by separation of variables:

$$\int \frac{dx}{1+x^2}=\int dt,$$

which yields

$$\tan^{-1} x = t + C.$$

The initial condition $x(0) = 0$ implies $C = 0$. Hence $x(t) = \tan t$ is the solution. But notice that this solution exists only for $-\pi/2 < t < \pi/2$, because $x(t) \to \pm\infty$ as $t \to \pm\pi/2$. Outside of that time interval, there is no solution to the initial value problem for $x_0 = 0$. ■

The amazing thing about Example 2.5.2 is that the system has solutions that reach infinity *in finite time*. This phenomenon is called ***blow-up.*** As the name suggests, it is of physical relevance in models of combustion and other runaway processes.

There are various ways to extend the existence and uniqueness theorem. One can allow f to depend on time t, or on several variables $x_1, \ldots, x_n$. One of the most useful generalizations will be discussed later in Section 6.2.

From now on, we will not worry about issues of existence and uniqueness—our vector fields will typically be smooth enough to avoid trouble. If we happen to come across a more dangerous example, we'll deal with it then.

2.6 Impossibility of Oscillations

Fixed points dominate the dynamics of first-order systems. In all our examples so far, all trajectories either approached a fixed point, or diverged to $\pm\infty$. In fact, those are the *only* things that can happen for a vector field on the real line. The reason is that trajectories are forced to increase or decrease monotonically, or remain constant (Figure 2.6.1). To put it more geometrically, the phase point never reverses direction.

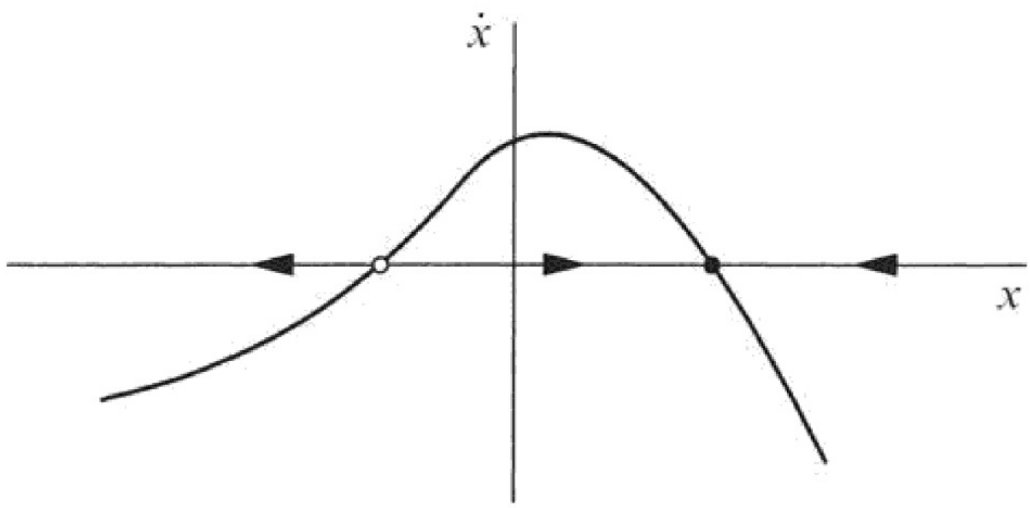

Figure 2.6.1

Thus, if a fixed point is regarded as an equilibrium solution, the approach to equilibrium is always *monotonic*—overshoot and damped oscillations can never occur in a first-order system. For the same reason, undamped oscillations are impossible. Hence *there are no periodic solutions to* $= f(x)$.

These general results are fundamentally topological in origin. They reflect the fact that $= f(x)$ corresponds to flow on a *line.* If you flow monotonically on a line, you'll never come back to your starting place—that's why periodic solutions are impossible. (Of course, if we were dealing with a *circle* rather than a line, we *could* eventually return to our starting place. Thus vector fields on the circle can exhibit periodic solutions, as we discuss in Chapter 4.)

Mechanical Analog: Overdamped Systems

It may seem surprising that solutions to $= f(x)$ can't oscillate. But this result becomes obvious if we think in terms of a mechanical analog. We regard $= f(x)$ as a limiting case of Newton's law, in the limit where the "inertia term" m is negligible.

For example, suppose a mass m is attached to a nonlinear spring whose restoring force is $F(x)$, where x is the displacement from the origin. Furthermore, suppose that the mass is immersed in a vat of very viscous fluid, like honey or motor oil (Figure 2.6.2), so that it is subject to a damping force b . Then Newton's law is $m + b = F(x)$.

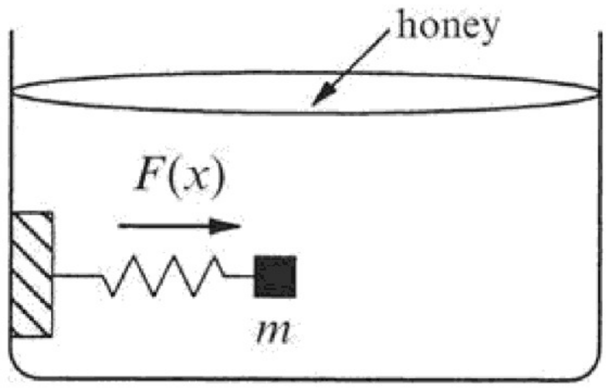

Figure 2.6.2

If the viscous damping is strong compared to the inertia term (b m), the system should behave like $b = F(x)$, or equivalently $= f(x)$, where $f(x) = b^{-1}F(x)$. In this ***over damped*** limit, the behavior of the mechanical system is clear. The mass prefers to sit at a stable equilibrium, where $f(x) = 0$ and $f'(x) < 0$. If displaced a bit, the mass is slowly dragged back to equilibrium by the restoring force. No overshoot can occur, because the damping is enormous. And undamped oscillations are out of the question! These conclusions agree with those obtained earlier by geometric reasoning.

Actually, we should confess that this argument contains a slight swindle. The neglect of the inertia term m is valid, but only after a rapid initial transient during which the inertia and damping terms are of comparable size. An honest discussion of this point requires more machinery than we have available. We'll return to this matter in Section 3.5.

2.7 Potentials

There's another way to visualize the dynamics of the first-order system $= f(x)$, based on the physical idea of potential energy. We picture a particle sliding down the walls of a potential well, where the ***potential*** $V(x)$ is defined by

$$f(x)=-dVdx.$$

As before, you should imagine that the particle is heavily damped—its inertia is completely negligible compared to the damping force and the force due to the potential. For example, suppose that the particle has to slog through a thick layer of goo that covers the walls of the potential (Figure 2.7.1).

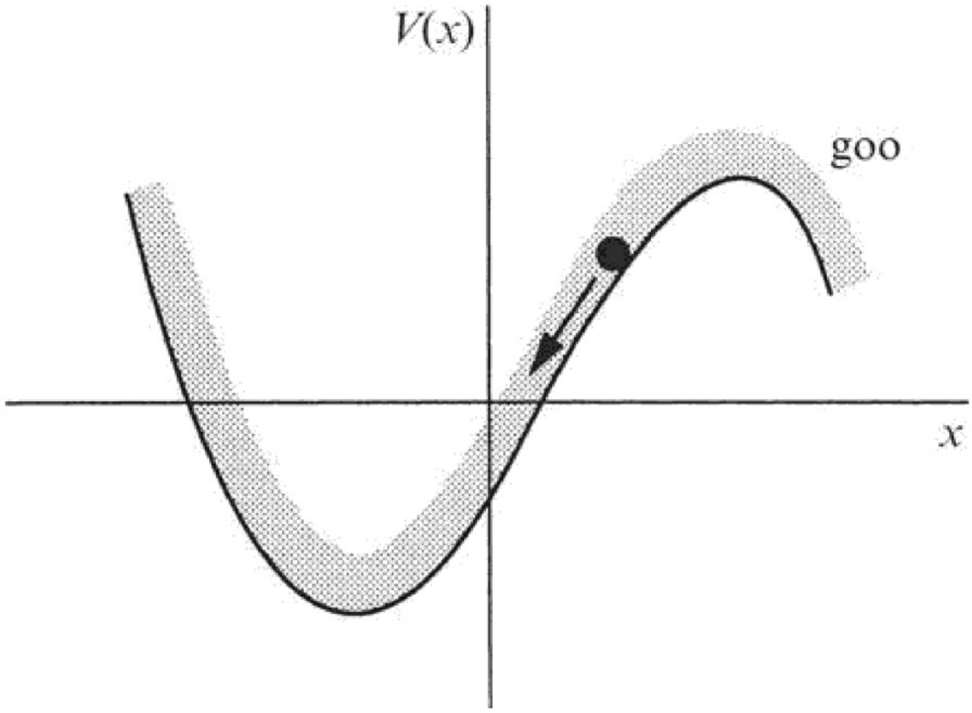

Figure 2.7.1

The negative sign in the definition of V follows the standard convention in physics; it implies that the particle always moves “downhill” as the motion proceeds. To see this, we think of x as a function of t, and then calculate the time-derivative of $V(x(t))$. Using the chain rule, we obtain

$$dVdx=dVdxdxdt.$$

Now for a first-order system,

$$dxdt=-dVdx,$$

since $= f(x) = -dV/dx$, by the definition of the potential. Hence,

$$dVdt=-(dVdx)2\leq 0.$$

Thus $V(t)$ *decreases along trajectories*, and so the particle always moves toward lower potential. Of course, if the particle happens to be at an ***equilibrium*** point where $dV/dx = 0$, then V remains constant. This is to be expected, since $dV/dx = 0$ implies $= 0$; equilibria occur at the fixed points of the vector field. Note that local minima of $V(x)$ correspond to *stable* fixed points, as we’d expect intuitively, and local maxima correspond to unstable fixed points.

EXAMPLE 2.7.1:

Graph the potential for the system $= -x$, and identify all the equilibrium points.

Solution: We need to find $V(x)$ such that $-dV/dx = -x$. The general solution is V(x)=12x2+C, where C is an arbitrary constant. (It always happens that the potential is only defined up to an additive constant. For convenience, we usually choose $C = 0$.) The graph of $V(x)$ is shown in Figure 2.7.2. The only equilibrium point occurs at $x = 0$, and it’s stable. ■

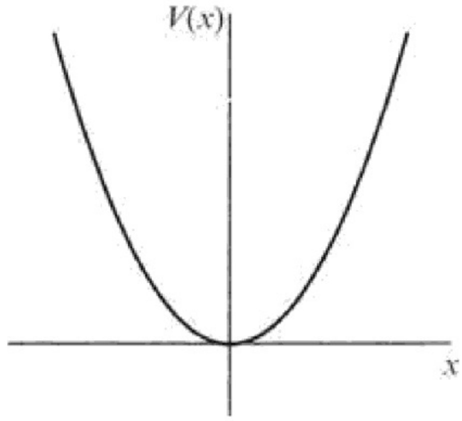

Figure 2.7.2

EXAMPLE 2.7.2:

Graph the potential for the system $= x - x^3$, and identify all equilibrium points.

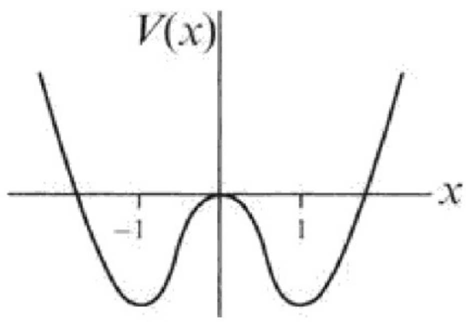

Figure 2.7.3

Solution: Solving $-dV/dx = x - x^3$ yields V=−12x2+14x4+C. Once again we set $C = 0$. Figure 2.7.3 shows the graph of V. The local minima at $x = \pm 1$ correspond to stable equilibria, and the local maximum at $x = 0$ corresponds to an unstable equilibrium. The potential shown in Figure 2.7.3 is often called a ***double-well potential,*** and the system is said to be ***bistable,*** since it has two stable equilibria. ■

2.8 Solving Equations on the Computer

Throughout this chapter we have used graphical and analytical methods to analyze first-order systems. Every budding dynamicist should master a third tool: numerical methods. In the old days, numerical methods were impractical because they required enormous amounts of tedious hand-calculation. But all that has changed, thanks to the computer. Computers enable us to approximate the solutions to analytically intractable problems, and also to visualize those solutions. In this section we take our first look at dynamics on the computer, in the context of ***numerical integration** of* $= f(x)$.

Numerical integration is a vast subject. We will barely scratch the surface. See Chapter 17 of Press et al. (2007) for an excellent treatment.

Euler's Method

The problem can be posed this way: given the differential equation $= f(x)$, subject to the condition $x = x_0$ at $t = t_0$, find a systematic way to approximate the solution $x(t)$.

Suppose we use the vector field interpretation of $= f(x)$. That is, we think of a fluid flowing steadily on the x-axis, with velocity $f(x)$ at the location x. Imagine we're riding along with a phase point being carried downstream by the fluid. Initially we're at x_0, and the local velocity is $f(x_0)$. If we flow for a short time Δt, we'll have moved a distance $f(x_0)\Delta t$, because distance = rate × time. Of course, that's not quite right, because our velocity was changing a little bit throughout the step. But over a sufficiently *small* step, the velocity will be nearly constant and our approximation should be reasonably good. Hence our new position $x(t_0 + \Delta t)$ is approximately $x_0 + f(x_0)$ At. Let's call this approximation x_1. Thus

$$x(t_0+\Delta t)\approx x_1=x_0+f(x_0)\Delta t.$$

Now we iterate. Our approximation has taken us to a new location x_1; our new velocity is $f(x_1)$; we step forward to $x_2 = x_1 + f(x_1)\Delta t$; and so on. In general, the update rule is

$$x_{n+1}=x_n+f(x_n)\Delta t.$$

This is the simplest possible numerical integration scheme. It is known as ***Euler's method.***

Euler's method can be visualized by plotting x versus t (Figure 2.8.1). The curve shows the exact solution $x(t)$, and the open dots show its values $x(t_n)$ at the discrete times $t_n = t_0 + n\Delta t$. The black dots show the approximate values given by the Euler method. As you can see, the approximation gets bad in a hurry unless At is extremely small. Hence Euler's method is not recommended in practice, but it contains the conceptual essence of the more accurate methods to be discussed next.

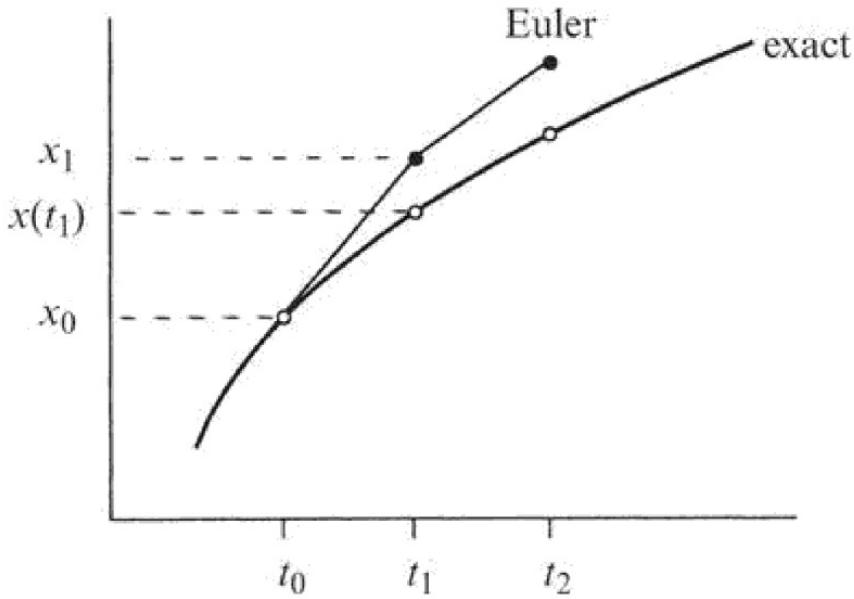

Figure 2.8.1

Refinements

One problem with the Euler method is that it estimates the derivative only at the left end of the time interval between t_n and t_{n+1}. A more sensible approach would be to use the *average* derivative across this interval. This is the idea behind the *improved Euler method.* We first take a trial step across the interval, using the Euler method. This produces a trial value $\tilde{x}_{n+1}=x_n+f(x_n)\Delta t$; the tilde above the x indicates that this is a tentative step, used only as a probe. Now that we've estimated the derivative on both ends of the interval, we average $f(x_n)$ and $f(\tilde{x}_{n+1})$, and use that to take the *real* step across the interval. Thus the improved Euler method is

$$\tilde{x}_{n+1} = x_n + f(x_n)\,\Delta t. \quad x_{n+1} = x_n + \tfrac{1}{2}\left[\, f(x_n) + f(\tilde{x}_{n+1}) \,\right] \Delta t.$$

This method is more accurate than the Euler method, in the sense that it tends to make a smaller ***error*** $E = |x(t_n) - x_n|$ for a given ***stepsize*** Δt. In both cases, the error $E \to 0$ as $\Delta t \to 0$, but the error decreases *faster* for the improved Euler method. One can show that $E \propto \Delta t$ for the Euler method, but $E \propto (\Delta t)^2$ for the improved Euler method (Exercises 2.8.7 and 2.8.8). In the jargon of numerical analysis, the Euler method is first order, whereas the improved Euler method is second order.

Methods of third, fourth, and even higher orders have been concocted, but you should realize that higher order methods are not necessarily superior. Higher order methods require more calculations and function evaluations, so there's a computational cost associated with them. In practice, a good balance is achieved by the ***fourth-order Runge–Kutta method.*** To find x_{n+1} in terms of x_n, this method first requires us to calculate the following four numbers (cunningly chosen, as you'll see in Exercise 2.8.9):

$$k_1 = f(x_n)\,\Delta t \quad k_2 = f(x_n + \tfrac{1}{2}k_1)\,\Delta t \quad k_3 = f(x_n + \tfrac{1}{2}k_2)\,\Delta t \quad k_4 = f(x_n + k_3)\,\Delta t.$$

Then x_{n+1} is given by

$$x_{n+1} = x_n + \tfrac{1}{6}(k_1 + 2k_2 + 2k_3 + k_4).$$

This method generally gives accurate results without requiring an excessively small stepsize Δt. Of course, some problems are nastier, and may require small steps in certain time intervals, while permitting very large steps elsewhere. In such cases, you may want to use a Runge–Kutta routine with an automatic stepsize control; see Press et al. (2007) for details.

Now that computers are so fast, you may wonder why we don't just pick a tiny Δt once and for all. The trouble is that excessively many computations will occur, and each one carries a penalty in the form of ***round-off error.*** Computers don't have infinite accuracy—they don't distinguish between numbers that differ by some small amount . For numbers of order 1, typically $\approx 10^{-7}$ for single precision and $\approx 10^{-16}$ for double precision. Round-off error occurs during every calculation, and will begin to accumulate in a serious way if Δt is too small. See Hubbard and West (1991) for a good discussion.

Practical Matters

You have several options if you want to solve differential equations on the computer. If you like to do things yourself, you can write your own numerical integration routines in your favorite programming language, and plot the results using whatever graphics programs are available. The information given above should be enough to get you started. For further guidance, consult Press et al. (2007).

A second option is to use existing packages for numerical methods. *Matlab, Mathematica*, and *Maple* all have programs for solving ordinary differential equations and graphing their solutions.

The final option is for people who want to explore dynamics, not computing. Dynamical systems software is available for personal computers. All you have to do is type in the equations and the parameters; the program solves the equations numerically and plots the results. Some recommended programs are *PPlane* (written by John Polking and available online as a Java applet; this is a pleasant choice for beginners) and *XPP* (by Bard Ermentrout, available on many platforms including iPhone and iPad; this is a more powerful tool for researchers and serious users).

EXAMPLE 2.8.1:

Solve the system $= x(1 - x)$ numerically.

Solution: This is a logistic equation (Section 2.3) with parameters $r = 1$, $K = 1$. Previously we gave a rough sketch of the solutions, based on geometric arguments; now we can draw a more quantitative picture.

As a first step, we plot the ***slope field*** for the system in the (t, x) plane (Figure 2.8.2). Here the equation $= x(1 - x)$ is being interpreted in a new way: for each point (t, x), the equation gives the slope dx/dt of the solution passing through that point. The slopes are indicated by little line segments in Figure 2.8.2.

Finding a solution now becomes a problem of drawing a curve that is always tangent to the local slope. Figure 2.8.3 shows four solutions starting from various points in the (t, x) plane.

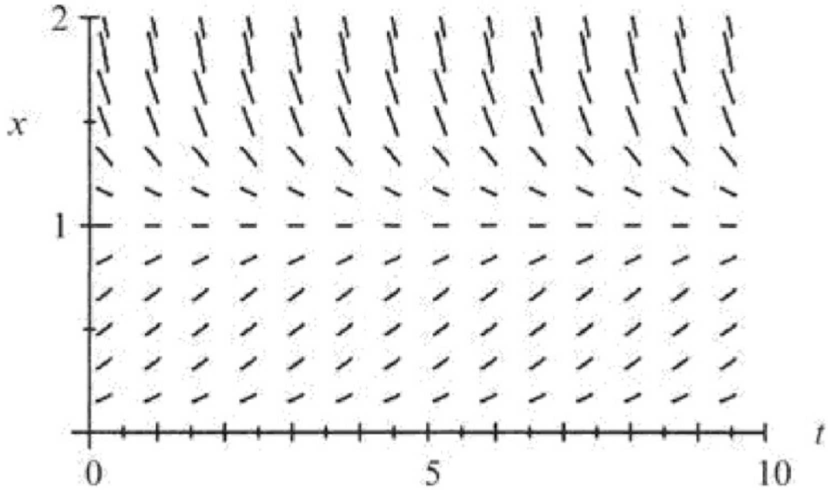

Figure 2.8.2

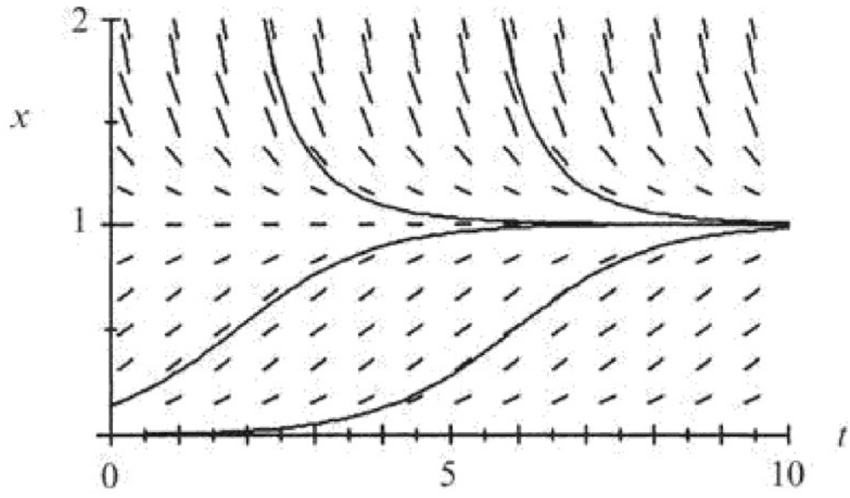

Figure 2.8.3

These numerical solutions were computed using the Runge–Kutta method with a stepsize $\Delta t = 0.1$. The solutions have the shape expected from Section 2.3. ■

Computers are indispensable for studying dynamical systems. We will use them liberally throughout this book, and you should do likewise.

EXERCISES FOR CHAPTER 2

2.1 A Geometric Way of Thinking

In the next three exercises, interpret $= \sin x$ as a flow on the line.

2.1.1 Find all the fixed points of the flow.

2.1.2 At which points does the flow have greatest velocity to the right?

2.1.3

a) Find the flow's acceleration x as a function of x.
b) Find the points where the flow has maximum positive acceleration.

2.1.4 (Exact solution of $= \sin x$) As shown in the text, $= \sin x$ has the solution $t = \ln |$ (csc x_0 + cot x_0)/(csc x + cot x)|, where $x_0 = x(0)$ is the initial value of x.

a) Given the specific initial condition $x_0 = \pi/4$, show that the solution above can be inverted to obtain

x(t)=2tan−1(et1+2).

Conclude that $x(t) \rightarrow \pi$ as $t \rightarrow \infty$, as claimed in Section 2.1. (You need to be good with trigonometric identities to solve this problem.)

b) Try to find the analytical solution for $x(t)$, given an *arbitrary* initial condition x_0.

2.1.5 (A mechanical analog)

a) Find a mechanical system that is approximately governed by $= \sin x$.
b) Using your physical intuition, explain why it now becomes obvious that $x^* = 0$ is an unstable fixed point and $x^* = \pi$ is stable.

2.2 Fixed Points and Stability

Analyze the following equations graphically. In each case, sketch the vector field on the real line, find all the fixed points, classify their stability, and sketch the graph of $x(t)$ for different initial conditions. Then try for a few minutes to obtain the analytical solution for $x(t)$; if you get stuck, don't try for too long since in several cases it's impossible to solve the equation in closed form!

2.2.1 x˙=4x2−16

2.2.2 x˙=1−x14

2.2.3 x˙=x−x3

2.2.4 x˙=e−xsinx

2.2.5 x˙=1+12cosx

2.2.6 x˙=1−2cosx

2.2.7 $= e^x - \cos x$ (Hint: Sketch the graphs of e^x and $\cos x$ on the same axes, and look for intersections. You won't be able to find the fixed points explicitly, but you can still find the qualitative behavior.)

2.2.8 (Working backwards, from flows to equations) Given an equation $= f(x)$,` we know how to sketch the corresponding flow on the real line. Here you are asked to solve the opposite problem: For the phase portrait shown in Figure 1, find an equation that is consistent with it. (There are an infinite number of correct answers—and wrong ones too.)

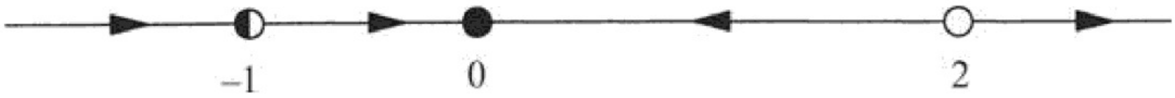

Figure 1

2.2.9 (Backwards again, now from solutions to equations) Find an equation = $f(x)$ whose solutions $x(t)$ are consistent with those shown in Figure 2.

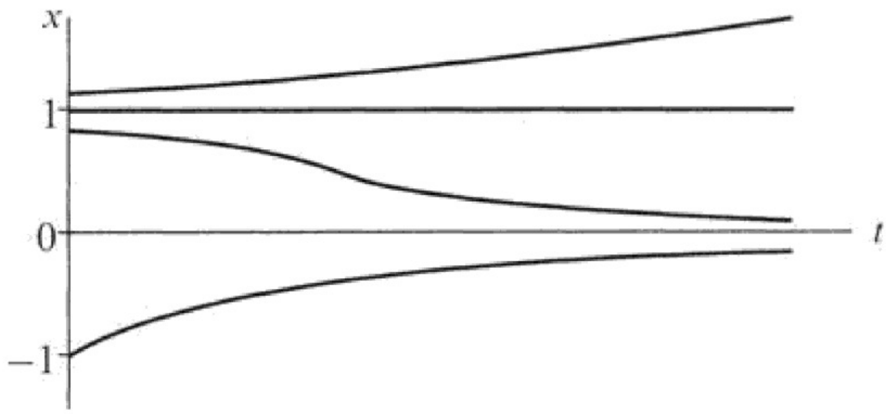

Figure 2

2.2.10 (Fixed points) For each of (a)–(e), find an equation = $f(x)$ with the stated properties, or if there are no examples, explain why not. (In all cases, assume that $f(x)$ is a smooth function.)

a) Every real number is a fixed point.
b) Every integer is a fixed point, and there are no others.
c) There are precisely three fixed points, and all of them are stable.
d) There are no fixed points.
e) There are precisely 100 fixed points.

2.2.11 (Analytical solution for charging capacitor) Obtain the analytical solution of the initial value problem Q˙=V0R−QRC, with $Q(0) = 0$, which arose in Example 2.2.2.

2.2.12 (A nonlinear resistor) Suppose the resistor in Example 2.2.2 is replaced by a nonlinear resistor. In other words, this resistor does not have a linear relation between voltage and current. Such nonlinearity arises in certain solid-state devices. Instead of $I_R = V/R$, suppose we have $I_R = g(V)$, where $g(V)$ has the shape shown in Figure 3.

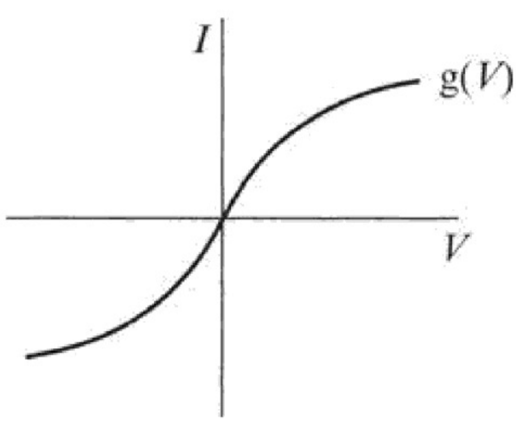

Figure 3

Redo Example 2.2.2 in this case. Derive the circuit equations, find all the fixed points, and analyze their stability. What qualitative effects does the nonlinearity introduce (if any)?

2.2.13 (Terminal velocity) The velocity $v(t)$ of a skydiver falling to the ground is governed by $m\dot{v} = mg - kv^2$, where m is the mass of the skydiver, g is the acceleration due to gravity, and $k > 0$ is a constant related to the amount of air resistance.

a) Obtain the analytical solution for $v(t)$, assuming that $v(0) = 0$.
b) Find the limit of $v(t)$ as $t \to \infty$. This limiting velocity is called the *terminal velocity*. (Beware of bad jokes about the word *terminal* and parachutes that fail to open.)
c) Give a graphical analysis of this problem, and thereby re-derive a formula for the terminal velocity.
d) An experimental study (Carlson et al. 1942) confirmed that the equation $m\dot{v} = mg - kv^2$ gives a good quantitative fit to data on human skydivers. Six men were dropped from altitudes varying from 10,600 feet to 31,400 feet to a terminal altitude of 2,100 feet, at which they opened their parachutes. The long free fall from 31,400 to 2,100 feet took 116 seconds. The average weight of the men and their equipment was 261.2 pounds. In these units, $g = 32.2$ ft/sec^2. Compute the average velocity V_{avg}.
e) Using the data given here, estimate the terminal velocity, and the value of the drag constant k. (Hints: First you need to find an exact formula for $s(t)$, the distance fallen, where $s(0) = 0$, $= v$, and $v(t)$ is known from part (a). You should get $s(t) = \frac{V^2}{g}\ln\left(\cosh\frac{gt}{V}\right)$, where V is the terminal velocity. Then solve for V graphically or numerically, using $s = 29{,}300$, $t = 116$, and $g = 32.2$.)

A slicker way to estimate V is to suppose $V \approx V_{avg}$, as a rough first approximation. Then show that $gt/V \approx 15$. Since $gt/V \gg 1$, we may use the approximation $\ln(\cosh x) \approx x - \ln 2$ for $x \gg 1$. Derive this approximation and then use it to obtain an analytical estimate of V. Then k follows from part (b). This analysis is from Davis (1962).

2.3 Population Growth

2.3.1 (Exact solution of logistic equation) There are two ways to solve the logistic equation $= rN(1 - N/K)$ analytically for an arbitrary initial condition N_0.

a) Separate variables and integrate, using partial fractions.
b) Make the change of variables $x = 1/N$. Then derive and solve the resulting differential equation for x.

2.3.2 (Autocatalysis) Consider the model chemical reaction

$$A + X \underset{k_{-1}}{\overset{k_1}{\rightleftharpoons}} 2X$$

in which one molecule of X combines with one molecule of A to form two molecules of X. This means that the chemical X stimulates its own production, a process called *autocatalysis*. This positive feedback process leads to a chain reaction, which eventually is limited by a "back reaction" in which $2X$ returns to $A + X$.

According to the *law of mass action* of chemical kinetics, the rate of an elementary reaction is proportional to the product of the concentrations of the reactants. We denote the concentrations by lowercase letters $x = [X]$ and $a = [A]$ Assume that there's an enormous surplus of chemical A, so that its concentration a can be regarded as constant. Then the equation for the kinetics of x is

$$\dot{x} = k_1 a x - k_{-1} x^2$$

where k_1 and k_{-1} are positive parameters called rate constants.

a) Find all the fixed points of this equation and classify their stability.

b) Sketch the graph of $x(t)$ for various initial values x_0.

2.3.3 (Tumor growth) The growth of cancerous tumors can be modeled by the Gompertz law $= -aN \ln(bN)$, where $N(t)$ is proportional to the number of cells in the tumor, and a, $b > 0$ are parameters.

a) Interpret a and b biologically.

b) Sketch the vector field and then graph $N(t)$ for various initial values.

The predictions of this simple model agree surprisingly well with data on tumor growth, as long as N is not too small; see Aroesty et al. (1973) and Newton (1980) for examples.

2.3.4 (The Allee effect) For certain species of organisms, the effective growth rate $/N$ is highest at intermediate N. This is called the Allee effect (Edelstein-Keshet 1988). For example, imagine that it is too hard to find mates when N is very small, and there is too much competition for food and other resources when N is large.

a) Show that $/N = r - a(N - b)^2$ provides an example of the Allee effect, if r, a, and b satisfy certain constraints, to be determined.

b) Find all the fixed points of the system and classify their stability.

c) Sketch the solutions $N(t)$ for different initial conditions.

d) Compare the solutions $N(t)$ to those found for the logistic equation. What are the qualitative differences, if any?

2.3.5 (Dominance of the fittest) Suppose X and Y are two species that reproduce exponentially fast: $= aX$ and $= bY$, respectively, with initial conditions X_0, $Y_0 > 0$ and growth rates $a > b > 0$. Here X is "fitter" than Y in the sense that it reproduces faster, as reflected by the inequality $a > b$. So we'd expect X to keep increasing its share of the total population $X + Y$ as $t \to \infty$. The goal of this exercise is to demonstrate this intuitive result, first analytically and then geometrically.

a) Let $x(t) = X(t)/[X(t) + Y(t)]$ denote X's share of the total population. By solving for $X(t)$ and $Y(t)$, show that $x(t)$ increases monotonically and approaches 1 as $t \to \infty$.

b) Alternatively, we can arrive at the same conclusions by deriving a differential equation for $x(t)$. To do so, take the time derivative of $x(t) = X(t)/[X(t) + Y(t)]$ using the quotient and chain rules. Then substitute for and and thereby show that $x(t)$ obeys the logistic equation $= (a - b)x(1 - x)$. Explain why this implies that $x(t)$ increases monotonically and approaches 1 as $t \to \infty$.

2.3.6 (Language death) Thousands of the world's languages are vanishing at an alarming rate, with 90 percent of them being expected to disappear by the end of this century. Abrams and Strogatz (2003) proposed the following model of language competition, and compared it to historical data on the decline of Welsh, Scottish Gaelic, Quechua (the most common surviving indigenous language in the Americas), and other endangered languages.

Let X and Y denote two languages competing for speakers in a given society. The proportion of the population speaking X evolves according to

$$\dot{x} = (1 - x)P_{YX} - xP_{XY}$$

where $0 \le x \le 1$ is the current fraction of the population speaking X, $1 - x$ is the complementary fraction speaking Y, and P_{YX} is the rate at which individuals switch from Y to X. This deliberately idealized model assumes that the population is well mixed (meaning that it lacks all spatial and social structure) and that all speakers are monolingual.

Next, the model posits that the attractiveness of a language increases with both its number of speakers and its perceived status, as quantified by a parameter $0 \le s \le 1$ that reflects the social or economic opportunities afforded to its speakers. Specifically, assume that $P_{YX} = sx^a$ and, by symmetry, $P_{XY} = (1 - s)(1-x)^a$, where the exponent $a > 1$ is an adjustable parameter. Then the model becomes

$$\dot{x}=s(1-x)xa-(1-s)x(1-x)a.$$

a) Show that this equation for has three fixed points.
b) Show that for all $a > 1$, the fixed points at $x = 0$ and $x = 1$ are both stable.
c) Show that the third fixed point, $0 < x^* < 1$, is unstable.

This model therefore predicts that two languages cannot coexist stably—one will eventually drive the other to extinction. For a review of generalizations of the model that allow for bilingualism, social structure, etc., see Castellano et al. (2009).

2.4 Linear Stability Analysis

Use linear stability analysis to classify the fixed points of the following systems. If linear stability analysis fails because $f'(x^*) = 0$, use a graphical argument to decide the stability.

2.4.1 $\dot{x}=x(1-x)$

2.4.2 $\dot{x}=x(1-x)(2-x)$

2.4.3 $\dot{x}=\tan x$

2.4.4 $\dot{x}=x2(6-x)$

2.4.5 $\dot{x}=1-e-x2$

2.4.6 $\dot{x}=\ln x$

2.4.7 $\dot{x} = ax - x^3$, where a can be positive, negative, or zero. Discuss all three cases.

2.4.8 Using linear stability analysis, classify the fixed points of the Gompertz model of tumor growth $= -aN \ln(bN)$ (As in Exercise 2.3.3, $N(t)$ is proportional to the number of cells in the tumor and $a,b > 0$ are parameters.)

2.4.9 (Critical slowing down) In statistical mechanics, the phenomenon of "critical slowing down" is a signature of a second-order phase transition. At the transition, the system relaxes to equilibrium much more slowly than usual. Here's a mathematical version of the effect:

a) Obtain the analytical solution to $= -x^3$ for an arbitrary initial condition. Show that $x(t) \to 0$ as $t \to \infty$, but that the decay is not exponential. (You should find that the decay is a much slower algebraic function of t.)
b) To get some intuition about the slowness of the decay, make a numerically accurate plot of the solution for the initial condition $x_0 = 10$, for $0 \le t \le 10$. Then, on the same graph, plot the solution to $= -x$ for the same initial condition.

2.5 Existence and Uniqueness

2.5.1 (Reaching a fixed point in a finite time) A particle travels on the half-line $x \ge 0$ with a velocity given by $= -x^c$, where c is real and constant.

a) Find all values of c such that the origin $x = 0$ is a stable fixed point.

b) Now assume that c is chosen such that $x = 0$ is stable. Can the particle ever reach the origin in *a finite* time? Specifically, how long does it take for the particle to travel from $x = 1$ to $x = 0$, as a function of c?

2.5.2 ("Blow-up": Reaching infinity in a finite time) Show that the solution to $= 1 + x^{10}$ escapes to $+\infty$ in a finite time, starting from any initial condition. (Hint: Don't try to find an exact solution; instead, compare the solutions to those of $= 1 + x^2$.)

2.5.3 Consider the equation $= rx + x^3$, where $r > 0$ is fixed. Show that $x(t) \rightarrow \pm\infty$ in finite time, starting from any initial condition $x_0 \neq 0$.

2.5.4 (Infinitely many solutions with the same initial condition) Show that the initial value problem $=$ $x^{1/3}$, $(0) = 0$, has an infinite number of solutions. (Hint: Construct a solution that stays at $x = 0$ until some arbitrary time t_0, after which it takes off.)

2.5.5 (A general example of non-uniqueness) Consider the initial value problem $= |x|^{p/q}$, $x(0) = 0$, where p and q are positive integers with no common factors.

a) Show that there are an infinite number of solutions for $x(t)$ if $p < q$.

b) Show that there is a unique solution if $p > q$.

2.5.6 (The leaky bucket) The following example (Hubbard and West 1991, p. 159) shows that in some physical situations, non-uniqueness is natural and obvious, not pathological.

Consider a water bucket with a hole in the bottom. If you see an empty bucket with a puddle beneath it, can you figure out when the bucket was full? No, of course not! It could have finished emptying a minute ago, ten minutes ago, or whatever. The solution to the corresponding differential equation must be non-unique when integrated backwards in time.

Here's a crude model of the situation. Let $h(t)$ = height of the water remaining in the bucket at time t; a = area of the hole; A = cross-sectional area of the bucket (assumed constant); $v(t)$ = velocity of the water passing through the hole.

a) Show that $av(t) = A\dot{h}(t)$. What physical law are you invoking?

b) To derive an additional equation, use conservation of energy. First, find the change in potential energy in the system, assuming that the height of the water in the bucket decreases by an amount Δh and that the water has density . Then find the kinetic energy transported out of the bucket by the escaping water. Finally, assuming all the potential energy is converted into kinetic energy, derive the equation $v^2 = 2gh$.

c) Combining (b) and (c), show h˙=–Ch, where C=2g(aA).

d) Given $h(0) = 0$ (bucket empty at $t = 0$), show that the solution for $h(t)$ is non-unique *in backwards time*, i.e., for $t < 0$.

2.6 Impossibility of Oscillations

2.6.1 Explain this paradox: a simple harmonic oscillator m $=-kx$ is a system that oscillates in one dimension (along the x-axis). But the text says one-dimensional systems can't oscillate.

2.6.2 (No periodic solutions to $= f(x)$) Here's an analytic proof that periodic solutions are impossible for a vector field on a line. Suppose on the contrary that $x(t)$ is a nontrivial periodic solution, i.e., $x(t) = x(t + T)$ for some $T > 0$, and $x(t) \neq x(t + s)$ for all $0 < s < T$ Derive a contradiction by considering ∫tt+Tf(x)dxdtdt.

2.7 Potentials

For each of the following vector fields, plot the potential function $V(x)$ and identify all the equilibrium points and their stability.

2.7.1 x˙=x(1−x)

2.7.2 x˙=3

2.7.3 x˙=sinx

2.7.4 x˙=2+sinx

2.7.5 x˙=−sinhx

2.7.6 x˙=r+x−x3, for various values of r.

2.7.7 (Another proof that solutions to $= f(x)$ can't oscillate) Let $= f(x)$ be a vector field on the line. Use the existence of a potential function $V(x)$ to show that solutions $x(t)$ cannot oscillate.

2.8 Solving Equations on the Computer

2.8.1 (Slope field) The slope is constant along horizontal lines in Figure 2.8.2. Why should we have expected this?

2.8.2 Sketch the slope field for the following differential equations. Then "integrate" the equation manually by drawing trajectories that are everywhere parallel to the local slope.

a)x˙=xb)x˙=1−x2c)x˙=1−4x(1−x)d)x˙=sinx

2.8.3 (Calibrating the Euler method) The goal of this problem is to test the Euler method on the initial value problem $= -x$, $x(0) = 1$.

a) Solve the problem analytically. What is the exact value of $x(1)$?

b) Using the Euler method with step size $\Delta t = 1$, estimate $x(1)$ numerically—call the result x^(1). Then repeat, using $\Delta t = 10^{-n}$, for $n = 1, 2, 3, 4$.

c) Plot the error E=| x^(1)−x(1) | as a function of At. Then plot ln E vs. ln t. Explain the results.

2.8.4 Redo Exercise 2.8.3, using the improved Euler method.

2.8.5 Redo Exercise 2.8.3, using the Runge–Kutta method.

2.8.6 (Analytically intractable problem) Consider the initial value problem $= x + e^{-x}$, $x(0) = 0$. In contrast to Exercise 2.8.3, this problem can't be solved analytically.

a) Sketch the solution $x(t)$ for $t \geq 0$.

b) Using some analytical arguments, obtain rigorous bounds on the value of x at $t = 1$. In other words, prove that $a < x(1) <$ b, for a, b to be determined. By being clever, try to make a and b as close together as possible. (Hint: Bound the given vector field by approximate vector fields that can be integrated analytically.)

c) Now for the numerical part: Using the Euler method, compute x at $t = 1$, correct to three decimal places. How small does the stepsize need to be to obtain the desired accuracy? (Give the order of magnitude, not the exact number.)

d) Repeat part (b), now using the Runge–Kutta method. Compare the results for stepsizes $\Delta t = 1$, $\Delta t = 0.1$, and $\Delta t = 0.01$.

2.8.7 (Error estimate for Euler method) In this question you'll use Taylor series expansions to estimate the error in taking one step by the Euler method. The exact solution and the Euler approximation both start at $x = x_0$ when $t = t_0$. We want to compare the exact value $x(t_1) = x(t_0 + \Delta t)$ with the Euler approximation $x_1 = x_0 + f(x_0)\Delta t$.

a) Expand $x(t_1) = x(t_0 + \Delta t)$ as a Taylor series in Δt, through terms of $O(\Delta t^2)$. Express your answer solely in terms of x_0, Δt, and f and its derivatives at x_0.

b) Show that the local error $|x(t_1) - x_1| \sim C(\Delta t)^2$ and give an explicit expression for the constant C. (Generally one is more interested in the global error incurred after integrating over a time interval of fixed length $T = n\Delta t$. Since each step produces an $O(\Delta t)^2$ error, and we take $n = T/\Delta t = O(\Delta t^{-1})$ steps, the global error $|x(t_n) - x_n|$ is $O(\Delta t)$, as claimed in the text.)

2.8.8 (Error estimate for the improved Euler method) Use the Taylor series arguments of Exercise 2.8.7 to show that the local error for the improved Euler method is $O(\Delta t^3)$.

2.8.9 (Error estimate for Runge–Kutta) Show that the Runge–Kutta method produces a local error of size $O(\Delta t^5)$. (Warning: This calculation involves massive amounts of algebra, but if you do it correctly, you'll be rewarded by seeing many wonderful cancellations. Teach yourself *Mathematica, Maple*, or some other symbolic manipulation language, and do the problem on the computer.)

3

BIFURCATIONS

3.0 Introduction

As we've seen in Chapter 2, the dynamics of vector fields on the line is very limited: all solutions either settle down to equilibrium or head out to $\pm\infty$. Given the triviality of the dynamics, what's interesting about one-dimensional systems? Answer: *Dependence on parameters.* The qualitative structure of the flow can change as parameters are varied. In particular, fixed points can be created or destroyed, or their stability can change. These qualitative changes in the dynamics are called ***bifurcations***, and the parameter values at which they occur are called ***bifurcation points.***

Bifurcations are important scientifically—they provide models of transitions and instabilities as some *control parameter* is varied. For example, consider the buckling of a beam. If a small weight is placed on top of the beam in Figure 3.0.1, the beam can support the load and remain vertical. But if the load is too heavy, the vertical position becomes unstable, and the beam may buckle.

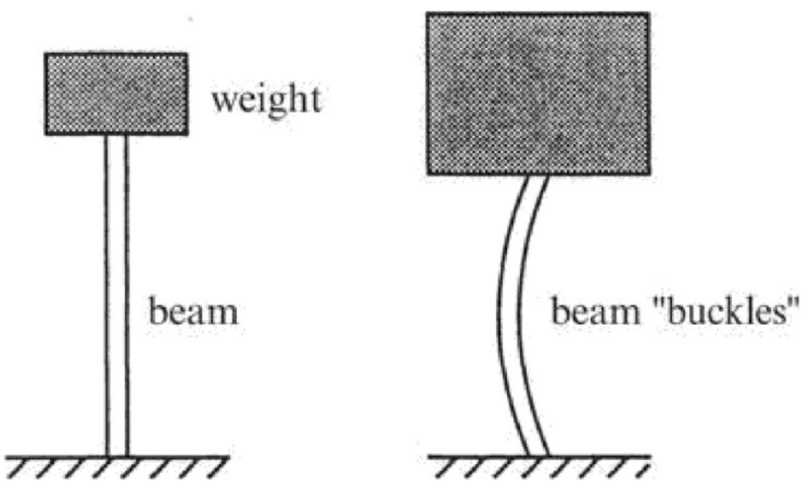

Figure 3.0.1

Here the weight plays the role of the control parameter, and the deflection of the beam from vertical plays the role of the dynamical variable x.

One of the main goals of this book is to help you develop a solid and practical understanding of bifurcations. This chapter introduces the simplest examples: bifurcations of fixed points for flows on the line. We'll use these bifurcations to model such dramatic phenomena as the onset of coherent radiation in a laser and the outbreak of an insect population. (In later chapters, when we step up to two-and three-dimensional phase spaces, we'll explore additional types of bifurcations and their scientific applications.)

We begin with the most fundamental bifurcation of all.

3.1 Saddle-Node Bifurcation

The saddle-node bifurcation is the basic mechanism by which fixed points are *created and destroyed.* As a parameter is varied, two fixed points move toward each other, collide, and mutually annihilate.

The prototypical example of a saddle-node bifurcation is given by the first-order system

$$\dot{x}=r+x^2 \tag{1}$$

where r is a parameter, which may be positive, negative, or zero. When r is negative, there are two fixed points, one stable and one unstable (Figure 3.1.1a).

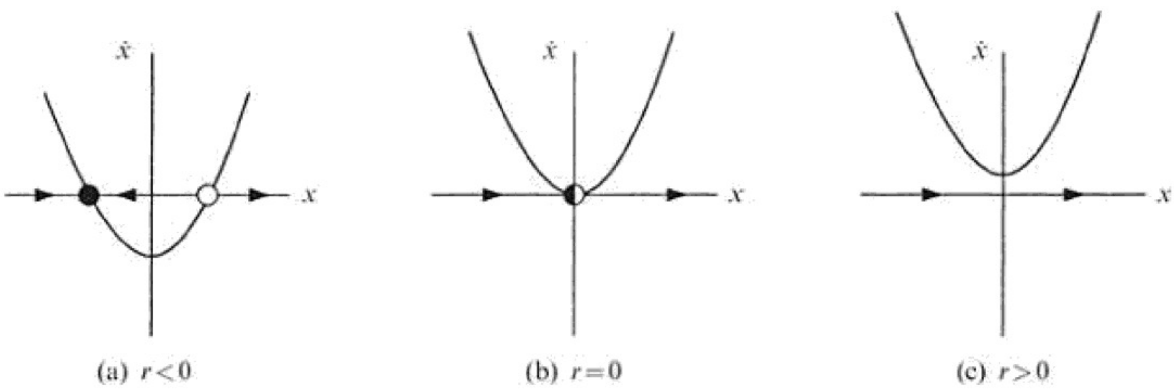

Figure 3.1.1

As r approaches 0 from below, the parabola moves up and the two fixed points move toward each other. When $r = 0$, the fixed points coalesce into a half-stable fixed point at $x^* = 0$ (Figure 3.1.1b). This type of fixed point is extremely delicate—it vanishes as soon as $r > 0$, and now there are no fixed points at all (Figure 3.1.1c).

In this example, we say that a *bifurcation* occurred at $r = 0$, since the vector fields for $r < 0$ and $r > 0$ are qualitatively different.

Graphical Conventions

There are several other ways to depict a saddle-node bifurcation. We can show a stack of vector fields for discrete values of r (Figure 3.1.2).

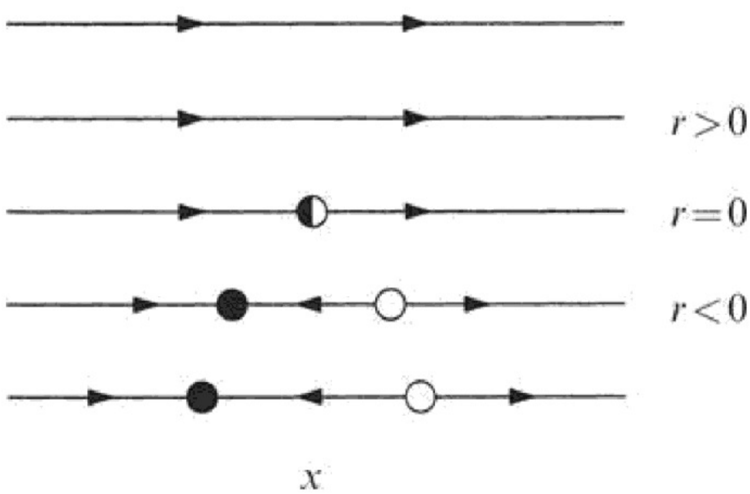

Figure 3.1.2

This representation emphasizes the dependence of the fixed points on r. In the limit of a *continuous* stack of vector fields, we have a picture like Figure 3.1.3. The curve shown is $r = -x^2$, i.e., $\dot{x}=0$, which gives the fixed points for different r. To distinguish between stable and unstable fixed points, we use a solid line for stable points and a broken line for unstable ones.

However, the most common way to depict the bifurcation is to invert the axes of Figure 3.1.3. The rationale is that r plays the role of an independent variable, and so should be plotted horizontally (Figure 3.1.4). The drawback is that now the x-axis has to be plotted vertically, which looks strange at first. Arrows are sometimes included in the picture, but not always. This picture is called the ***bifurcation diagram*** for the saddle-node bifurcation.

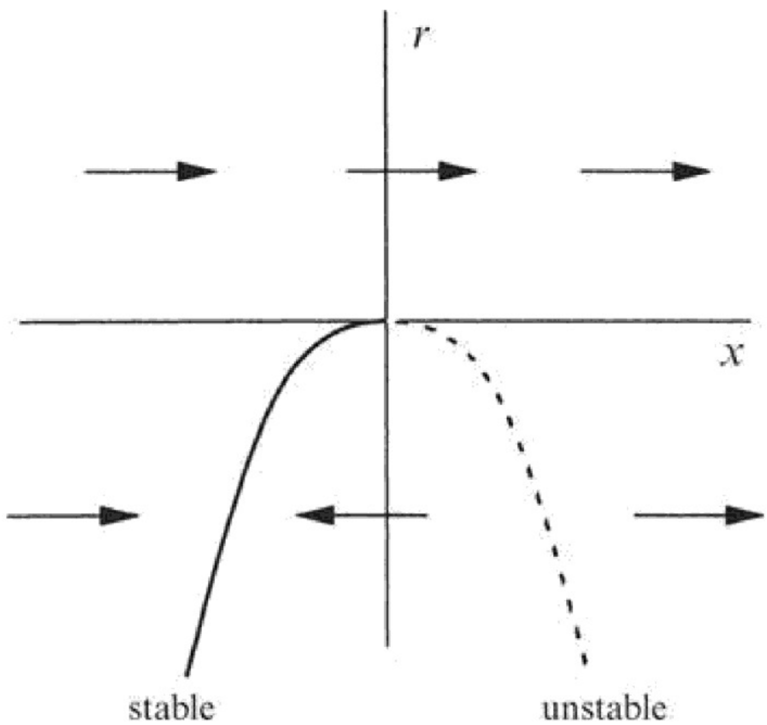

Figure 3.1.3

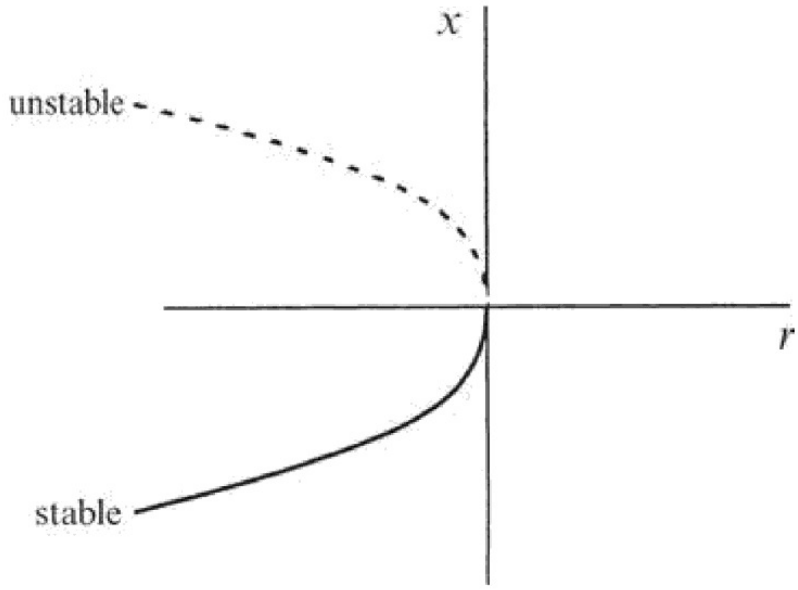

Figure 3.1.4

Terminology

Bifurcation theory is rife with conflicting terminology. The subject really hasn't settled down yet, and different people use different words for the same thing. For example, the saddle-node bifurcation is sometimes called *a fold bifurcation* (because the curve in Figure 3.1.4 has a fold in it) or a *turning-point bifurcation* (because the point $(x,r) = (0,0)$ is a "turning point.") Admittedly, the term *saddle-node* doesn't make much sense for vector fields on the line. The name derives from a completely analogous bifurcation seen in a higher-dimensional context, such as vector fields on the plane, where fixed points known as saddles and nodes can collide and annihilate (see Section 8.1).

The prize for most inventive terminology must go to Abraham and Shaw (1988), who write of a ***blue sky bifurcation.*** This term comes from viewing a saddle-node bifurcation in the other direction: a pair of fixed points appears "out of the clear blue sky" as a parameter is varied. For example, the vector field

$$\dot{x}=r-x^2 \qquad (2)$$

has no fixed points for $r < 0$, but then one materializes when $r = 0$ and splits into two when $r > 0$ (Figure 3.1.5). Incidentally, this example also explains why we use the word "bifurcation": it means "splitting into two branches."

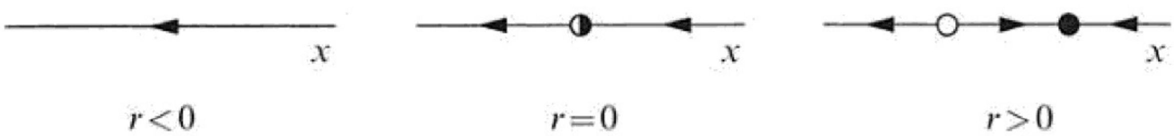

Figure 3.1.5

EXAMPLE 3.1.1:

Give a linear stability analysis of the fixed points in Figure 3.1.5.

Solution: The fixed points for $= f(x) = r - x^2$ are given by x*=±r. There are two fixed points for $r > 0$, and none for $r < 0$. To determine linear stability, we compute $f'(x^*) = -2x^*$. Thus x*=+r is stable, since $f'(x^*) < 0$. Similarly x*=−r is unstable. At the bifurcation point $r = 0$, we find $f'(x^*) = 0$; the linearization vanishes when the fixed points coalesce.

EXAMPLE 3.1.2:

Show that the first-order system $= r - x - e^{-x}$ undergoes a saddle-node bifurcation as r is varied, and find the value of r at the bifurcation point.

Solution: The fixed points satisfy $f(x) = r - x - e^{-x} = 0$. But now we run into a difficulty—in contrast to Example 3.1.1, we can't find the fixed points explicitly as a function of r. Instead we adopt a geometric approach. One method would be to graph the function $f(x) = r - x - e^{-x}$ for different values of r, look for its roots x^*, and then sketch the vector field on the x-axis. This method is fine, but there's an easier way. The point is that the two functions $r - x$ and e^{-x} have much more familiar graphs than their difference $r - x - e^{-x}$. So we plot $r - x$ and e^{-x} on the same picture (Figure 3.1.6a). Where the line $r - x$ intersects the curve e^{-x}, we have $r - x = e^{-x}$ and so $f(x) = 0$. *Thus, intersections of the line and the curve correspond to fixed points for the system.* This picture also allows us to read off the direction of flow on the x-axis: the flow is to the right where the line lies above the curve, since $r - x > e^{-x}$ and therefore > 0. Hence, the fixed point on the right is stable, and the one on the left is unstable.

Now imagine we start decreasing the parameter r. The line $r - x$ slides down and the fixed points approach each other. At some critical value $r = r_c$, the line becomes *tangent* to the curve and the fixed points coalesce in a saddle-node bifurcation (Figure 3.1.6b). For r below this critical value, the line lies below the curve and there are no fixed points (Figure 3.1.6c).

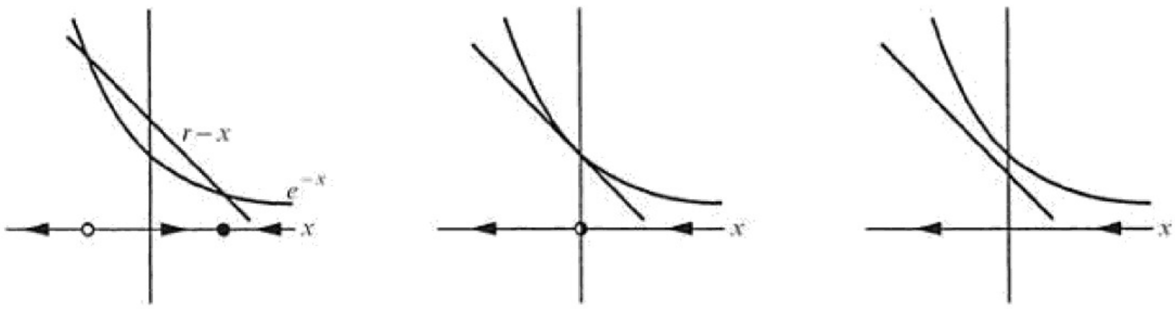

(a) (b) (c)

Figure 3.1.6

To find the bifurcation point r_c, we impose the condition that the graphs of $r - x$ and e^{-x} intersect *tangentially.* Thus we demand equality of the functions *and* their derivatives:

$$e^{-x} = r - x$$

and

$$\frac{d}{dx}e^{-x} = \frac{d}{dx}(r - x).$$

The second equation implies $-e^{-x} = -1$, so $x = 0$. Then the first equation yields $r = 1$. Hence the bifurcation point is $r_c = 1$, and the bifurcation occurs at $x = 0$. ■

Normal Forms

In a certain sense, the examples $= r - x^2$ or $= r + x^2$ are representative of *all* saddle-node bifurcations; that's why we called them "prototypical." The idea is that, close to a saddle-node bifurcation, the dynamics typically look like $= r - x^2$ or $= r + x^2$.

For instance, consider Example 3.1.2 near the bifurcation at $x = 0$ and $r = 1$. Using the Taylor expansion for e^{-x} about $x = 0$, we find

$$\dot{x} = r - x - e^{-x} = r - x - \left[1 - x + \frac{x^2}{2!} + \cdots\right] = (r - 1) - \frac{x^2}{2} + \cdots$$

to leading order in x. This has the same algebraic form as $= r - x^2$, and can be made to agree exactly by appropriate rescalings of x and r.

It's easy to understand why saddle-node bifurcations typically have this algebraic form. We just ask ourselves: how can two fixed points of $\dot{x} = f(x)$ collide and disappear as a parameter r is varied? Graphically, fixed points occur where the graph of $f(x)$ intersects the x-axis. For a saddle-node bifurcation to be possible, we need two nearby roots of $f(x)$; this means $f(x)$ must look locally "bowl-shaped" or parabolic (Figure 3.1.7).

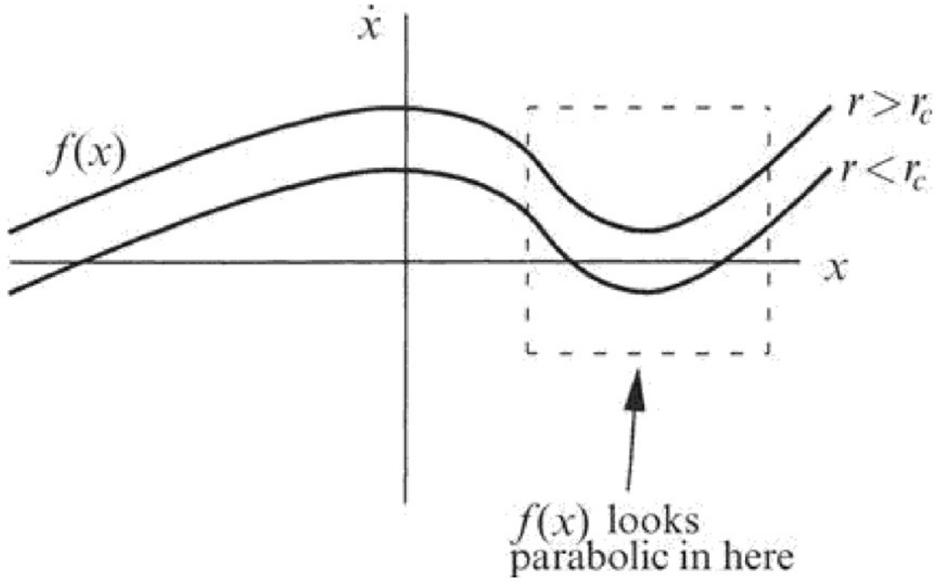

Figure 3.1.7

Now we use a microscope to zoom in on the behavior near the bifurcation. As r varies, we see a parabola intersecting the x-axis, then becoming tangent to it, and then failing to intersect. This is exactly the scenario in the prototypical Figure 3.1.1.

Here's a more algebraic version of the same argument. We regard f as a function of both x and r, and examine the behavior of $= f(x, r)$ near the bifurcation at $x = x^*$ and $r = r_c$. Taylor's expansion yields

$$\dot{x} = f(x, r) = f(x^*, r_c) + (x - x^*)\left.\frac{\partial f}{\partial x}\right|_{(x^*, r_c)} + (r - r_c)\left.\frac{\partial f}{\partial r}\right|_{(x^*, r_c)} + \frac{1}{2}(x - x^*)^2\left.\frac{\partial^2 f}{\partial x^2}\right|_{(x^*, r_c)} + \cdots$$

where we have neglected quadratic terms in $(r - r_c)$ and cubic terms in $(x - x^*)$. Two of the terms in this equation vanish: $f(x^*, r_c) = 0$ since x^* is a fixed point, and $\partial f/\partial x|_{(x^*, r_c)} = 0$ by the tangency condition of a saddle-node bifurcation. Thus

$$\dot{x} = a(r - r_c) + b(x - x^*)^2 + \cdots \qquad (3)$$

where $a = \partial f/\partial x|_{(x^*, r_c)}$ and $b = \frac{1}{2}\partial^2 f/\partial x^2|_{(x^*, r_c)}$ Equation (3) agrees with the form of our prototypical examples. (We are assuming that $a, b \neq 0$, which is the typical case; for instance, it would be a very special situation if the second derivative $\partial^2 f/\partial x^2$ also happened to vanish at the fixed point.)

What we have been calling prototypical examples are more conventionally known as ***normal forms*** for the saddle-node bifurcation. There is much, much more to normal forms than we have indicated here. We will be seeing their importance throughout this book. For a more detailed and precise discussion, see Guckenheimer and Holmes (1983) or Wiggins (1990).

3.2 Transcritical Bifurcation

There are certain scientific situations where a fixed point must exist for all values of a parameter and can never be destroyed. For example, in the logistic equation and other simple models for the growth of a single species, there is a fixed point at zero population, regardless of the value of the growth rate. However, such a fixed point may *change its stability* as the parameter is varied. The transcritical bifurcation is the standard mechanism for such changes in stability.

The normal form for a transcritical bifurcation is

$$\dot{x} = rx - x^2. \quad (1)$$

This looks like the logistic equation of Section 2.3, but now we allow x and r to be either positive or negative.

Figure 3.2.1 shows the vector field as r varies. Note that there is a fixed point at $x^* = 0$ for *all* values of r.

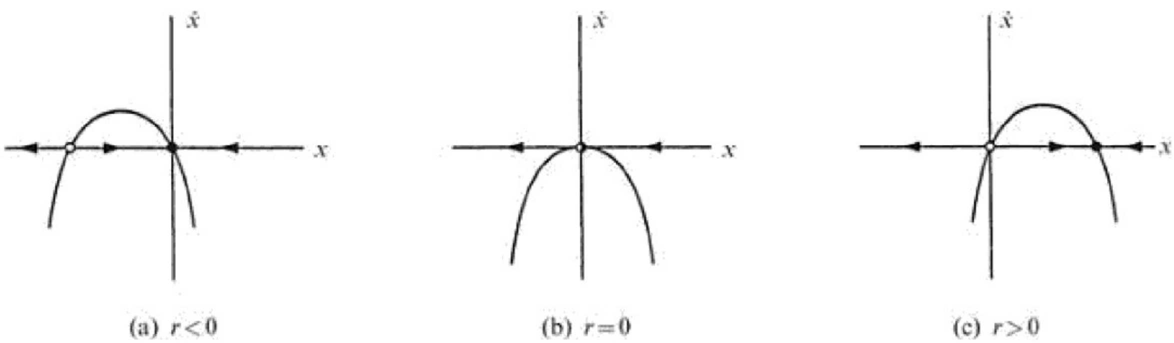

Figure 3.2.1

For $r < 0$, there is an unstable fixed point at $x^* = r$ and a stable fixed point at $x^* = 0$. As r increases, the unstable fixed point approaches the origin, and coalesces with it when $r = 0$. Finally, when $r > 0$, the origin has become unstable, and $x^* = r$ is now stable. Some people say that an ***exchange of stabilities*** has taken place between the two fixed points.

Please note the important difference between the saddle-node and transcritical bifurcations: in the transcritical case, the two fixed points don't disappear after the bifurcation—instead they just switch their stability.

Figure 3.2.2 shows the bifurcation diagram for the transcritical bifurcation. As in Figure 3.1.4, the parameter r is regarded as the independent variable, and the fixed points $x^* = 0$ and $x^* = r$ are shown as dependent variables.

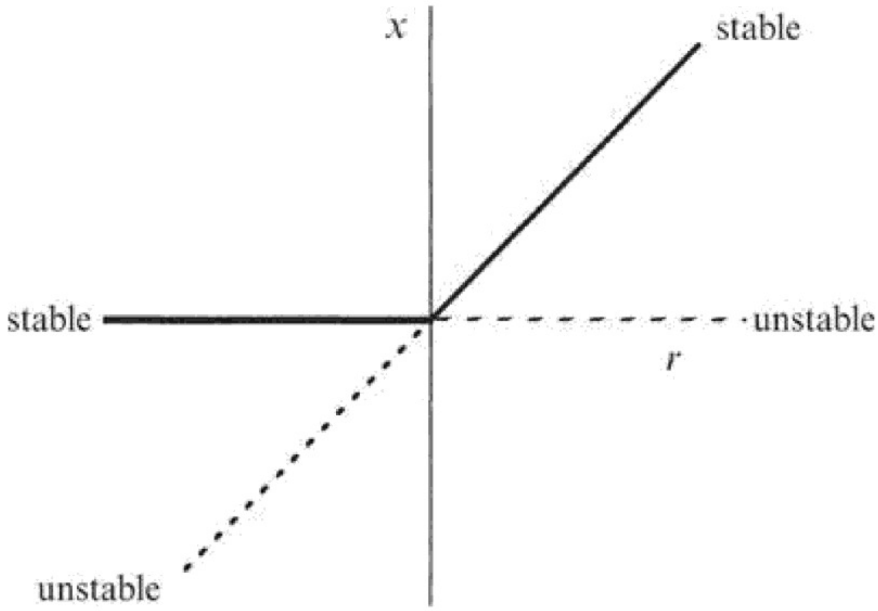

Figure 3.2.2

EXAMPLE 3.2.1:

Show that the first-order system $= x(1-x^2) - a(1-e^{-bx})$ undergoes a transcritical bifurcation at $x = 0$ when the parameters *a, b* satisfy a certain equation, to be determined. (This equation defines a ***bifurcation curve*** in the (*a*, *b*) parameter space.) Then find an approximate formula for the fixed point that bifurcates from $x = 0$, assuming that the parameters are close to the bifurcation curve.

Solution: Note that $x = 0$ is a fixed point for all (*a*, *b*). This makes it plausible that the fixed point will bifurcate transcritically, if it bifurcates at all. For small *x*, we find

$$1-e^{-bx} = 1-\left[1-bx+\frac{1}{2}b^2x^2+O(x^3)\right] = bx-\frac{1}{2}b^2x^2+O(x^3)$$

and so

$$\dot{x} = x-a\left(bx-\frac{1}{2}b^2x^2\right)+O(x^3) = (1-ab)x+\left(\frac{1}{2}ab^2\right)x^2+O(x^3).$$

Hence a transcritical bifurcation occurs when $ab = 1$; this is the equation for the bifurcation curve. The nonzero fixed point is given by the solution of $1-ab+\left(\frac{1}{2}ab^2\right)x \approx 0$, i.e.,

$$x^* \approx \frac{2(ab-1)}{ab^2}.$$

This formula is approximately correct only if x^* is small, since our series expansions are based on the assumption of small *x*. Thus the formula holds only when *ab* is close to 1, which means that the parameters must be close to the bifurcation curve. ■

EXAMPLE 3.2.2:

Analyze the dynamics *of* $= r \ln x + x - 1$ near $x = 1$, and show that the system undergoes a transcritical bifurcation at a certain value of *r*. Then find new variables *X* and *R* such that the system reduces to the approximate normal form $\approx RX - X^2$ near the bifurcation.

Solution: First note that $x = 1$ is a fixed point for all values of *r*. Since we are interested in the dynamics near this fixed point, we introduce a new variable $u = x - 1$, where *u* is small. Then

$$\dot{u} = \dot{x} = r\ln(1+u) + u = r\left[u - \tfrac{1}{2}u^2 + O(u^3) \right] + u \approx (r+1)u - \tfrac{1}{2}ru^2 + O(u^3).$$

Hence a transcritical bifurcation occurs at $r_c = -1$.

To put this equation into normal form, we first need to get rid of the coefficient of u^2. Let $u = av$, where a will be chosen later. Then the equation for v is

$$\dot{v} = (r+1)v - \left(\tfrac{1}{2}ra\right)v^2 + O(v^3).$$

So if we choose $a = 2/r$, the equation becomes

$$\dot{v} = (r+1)v - v^2 + O(v^3).$$

Now if we let $R = r + 1$ and $X = v$, we have achieved the approximate normal form $\approx RX - X^2$, where cubic terms of order $O(X^3)$ have been neglected. In terms of the original variables, $X = v = u/a = \tfrac{1}{2}r(x-1)$. ■

To be a bit more accurate, the theory of normal forms assures us that we can find a change of variables such that the system becomes $= RX - X^2$, with *strict*, rather than approximate, equality. Our solution above gives an approximation to the necessary change of variables. For careful treatments of normal form theory, see the books of Guckenheimer and Holmes (1983), Wiggins (1990), or Manneville (1990).

3.3 Laser Threshold

Now it's time to apply our mathematics to a scientific example. We analyze an extremely simplified model for a laser, following the treatment given by Haken (1983).

Physical Background

We are going to consider a particular type of laser known as a solid-state laser, which consists of a collection of special "laser-active" atoms embedded in a solid-state matrix, bounded by partially reflecting mirrors at either end. An external energy source is used to excite or "pump" the atoms out of their ground states (Figure 3.3.1).

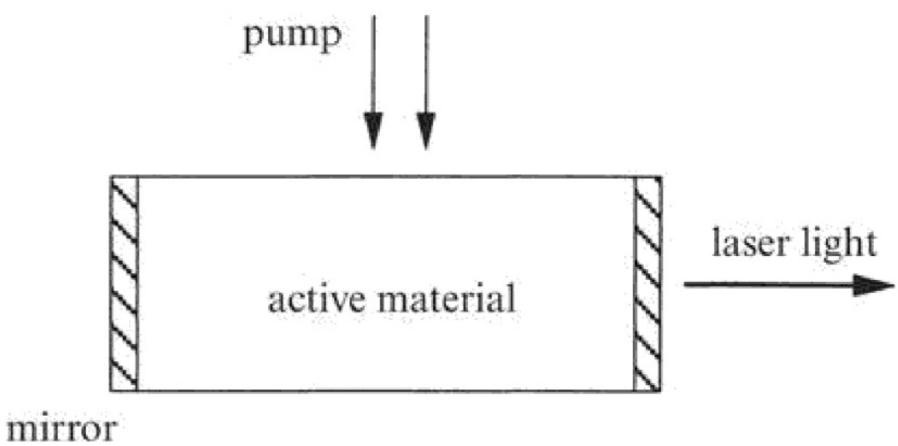

Figure 3.3.1

Each atom can be thought of as a little antenna radiating energy. When the pumping is relatively weak, the laser acts just like an ordinary *lamp:* the excited atoms oscillate independently of one another and emit randomly phased light waves.

Now suppose we increase the strength of the pumping. At first nothing different happens, but then suddenly, when the pump strength exceeds a certain threshold, the atoms begin to oscillate in phase—the lamp has turned into a ***laser.*** Now the trillions of little antennas act like one giant antenna and produce a beam of radiation that is much more coherent and intense than that produced below the laser threshold.

This sudden onset of coherence is amazing, considering that the atoms are being excited completely at random by the pump! Hence the process is *self-organizing:* the coherence develops because of a cooperative interaction among the atoms themselves.

Model

A proper explanation of the laser phenomenon would require us to delve into quantum mechanics. See Milonni and Eberly (1988) for an intuitive discussion.

Instead we consider a simplified model of the essential physics (Haken 1983, p. 127). The dynamical variable is the number of photons $n(t)$ in the laser field. Its rate of change is given by

$$\dot{n} = \text{gain} - \text{loss} = GnN - kn.$$

The gain term comes from the process of *stimulated emission*, in which photons stimulate excited atoms to emit additional photons. Because this process occurs via random encounters between photons and excited atoms, it occurs at a rate proportional to n and to the number of excited atoms, denoted by $N(t)$. The parameter $G > 0$ is known as the gain coefficient. The loss term models the escape of photons through the endfaces of the laser. The parameter $k > 0$ is a rate constant; its reciprocal $= 1/k$ represents the typical lifetime of a photon in the laser.

Now comes the key physical idea: after an excited atom emits a photon, it drops down to a lower energy level and is no longer excited. Thus N decreases by the emission of photons. To capture this effect, we need to write an equation relating N to n. Suppose that in the absence of laser action, the pump keeps the number of excited atoms fixed at N_0. Then the *actual* number of excited atoms will be reduced by the laser process. Specifically, we assume

$$N(t) = N_0 - \alpha n,$$

where > 0 is the rate at which atoms drop back to their ground states. Then

$$\dot{n} = Gn(N_0 - \alpha n) - kn = (GN - k)n - (\alpha G)n^2.$$

We're finally on familiar ground—this is a first-order system for $n(t)$. Figure 3.3.2 shows the corresponding vector field for different values of the pump strength N_0. Note that only positive values of n are physically meaningful.

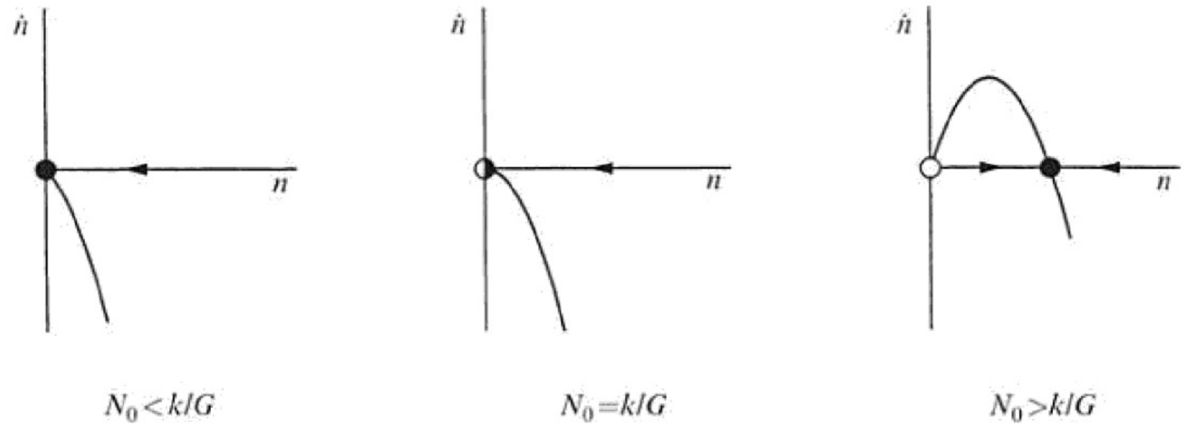

Figure 3.3.2

When $N_0 < k/G$, the fixed point at $n^* = 0$ is stable. This means that there is no stimulated emission and the laser acts like a lamp. As the pump strength N_0 is increased, the system undergoes a transcritical bifurcation when $N_0 = k/G$. For $N_0 > k/G$, the origin loses stability and a stable fixed point appears at $n^* = (GN_0 - k)/G > 0$, corresponding to spontaneous laser action. Thus $N_0 = k/G$ can be interpreted as the ***laser threshold*** in this model. Figure 3.3.3 summarizes our results.

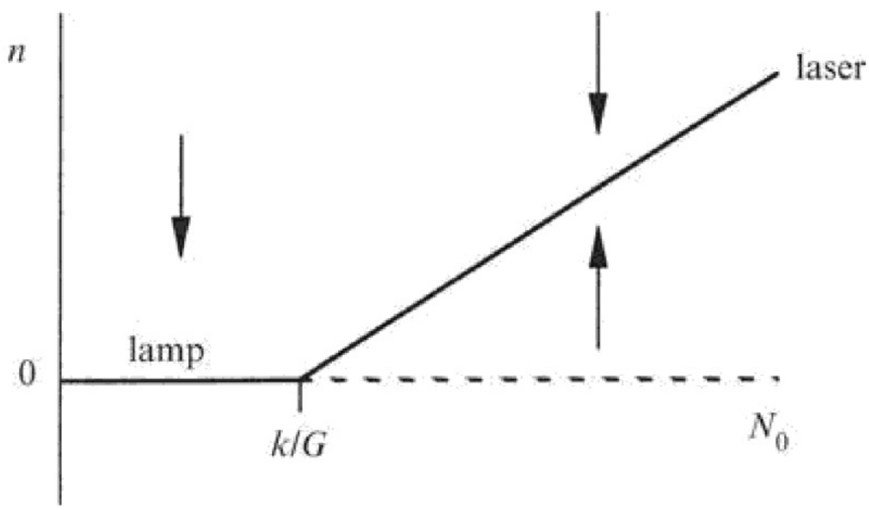

Figure 3.3.3

Although this model correctly predicts the existence of a threshold, it ignores the dynamics of the excited atoms, the existence of spontaneous emission, and several other complications. See Exercises 3.3.1 and 3.3.2 for improved models.

3.4 Pitchfork Bifurcation

We turn now to a third kind of bifurcation, the so-called pitchfork bifurcation. This bifurcation is common in physical problems that have a ***symmetry.*** For example, many problems have a spatial symmetry between left and right. In such cases, fixed points tend to appear and disappear in symmetrical pairs. In the buckling example of Figure 3.0.1, the beam is stable in the vertical position if the load is small. In this case there is a stable fixed point corresponding to zero deflection. But if the load exceeds the buckling threshold, the beam may buckle to *either* the left or the right. The vertical position has gone unstable, and two new symmetrical fixed points, corresponding to left- and rightbuckled configurations, have been born.

There are two very different types of pitchfork bifurcation. The simpler type is called *supercritical*, and will be discussed first.

Supercritical Pitchfork Bifurcation

The normal form of the supercritical pitchfork bifurcation is

$$\dot{x}=rx-x^3. \qquad (1)$$

Note that this equation is ***invariant*** under the change of variables $x \to -x$. That is, if we replace x by $-x$ and then cancel the resulting minus signs on both sides of the equation, we get (1) back again. This invariance is the mathematical expression of the left-right symmetry mentioned earlier. (More technically, one says that the vector field is *equivariant*, but we'll use the more familiar language.)

Figure 3.4.1 shows the vector field for different values of r.

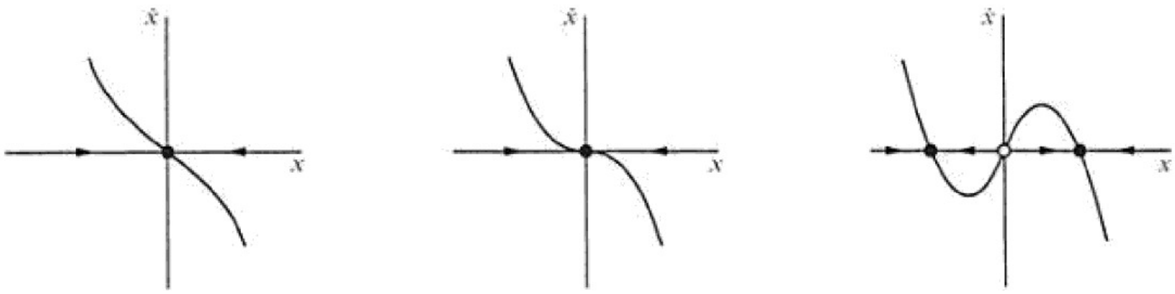

(a) $r<0$ (b) $r=0$ (c) $r>0$

Figure 3.4.1

When $r < 0$, the origin is the only fixed point, and it is stable. When $r = 0$, the origin is still stable, but much more weakly so, since the linearization vanishes. Now solutions no longer decay exponentially fast—instead the decay is a much slower algebraic function of time (recall Exercise 2.4.9). This lethargic decay is called ***critical slowing down*** in the physics literature. Finally, when $r > 0$, the origin has become unstable. Two new stable fixed points appear on either side of the origin, symmetrically located at $x^*=\pm r$.

The reason for the term "pitchfork" becomes clear when we plot the bifurcation diagram (Figure 3.4.2). Actually, pitchfork trifurcation might be a better word!

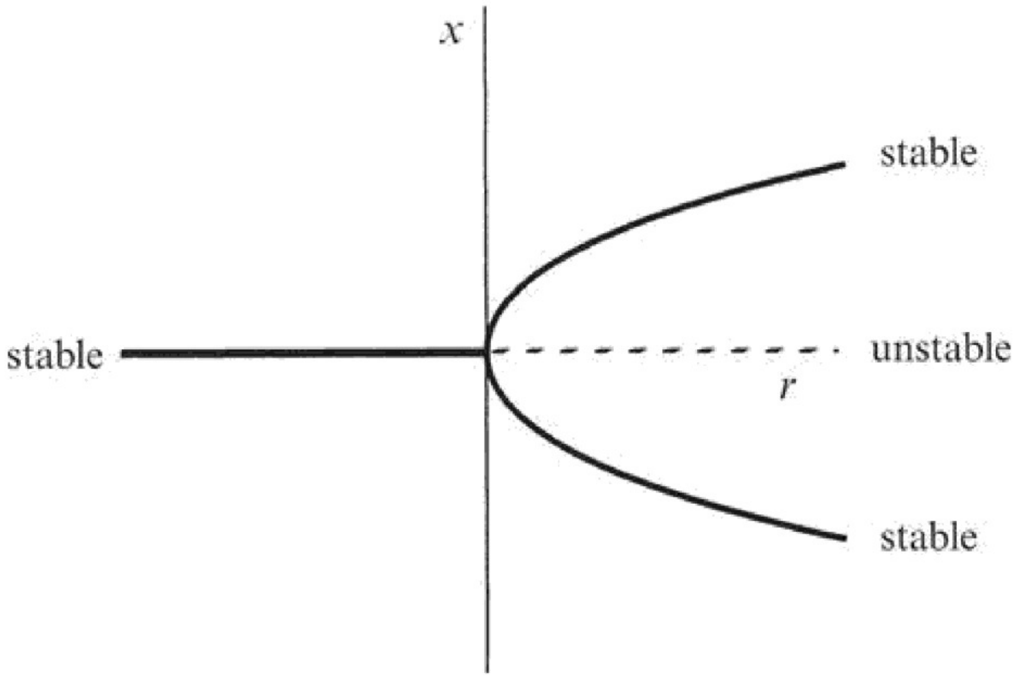

Figure 3.4.2

EXAMPLE 3.4.1:

Equations similar to $= -x + \beta \tanh x$ arise in statistical mechanical models of magnets and neural networks (see Exercise 3.6.7 and Palmer 1989). Show that this equation undergoes a supercritical pitchfork bifurcation as β is varied. Then give a *numerically accurate* plot of the fixed points for each β.

Solution: We use the strategy of Example 3.1.2 to find the fixed points. The graphs of $y = x$ and $y = \beta \tanh x$ are shown in Figure 3.4.3; their intersections correspond to fixed points. The key thing to realize is that as β increases, the tanh curve becomes steeper at the origin (its slope there is β). Hence for $\beta < 1$ the origin is the only fixed point. A pitchfork bifurcation occurs at $\beta = 1$, $x^* = 0$, when the tanh curve develops a slope of 1 at the origin. Finally, when ß > 1, two new stable fixed points appear, and the origin becomes unstable.

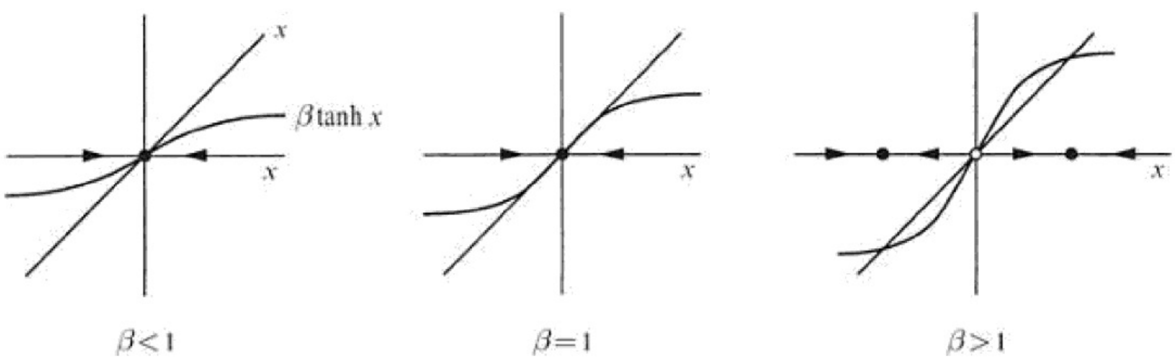

Figure 3.4.3

Now we want to compute the fixed points x^* for each β. Of course, one fixed point always occurs at $x^* = 0$; we are looking for the other, nontrivial fixed points. One approach is to solve the equation $x^* = \beta \tanh x^*$ numerically, using the Newton-Raphson method or some other root-finding scheme. (See Press et al. (2007) for a friendly and informative discussion of numerical methods.) But there's an easier way, which comes from changing our point of view. Instead of studying the dependence of x^* on β, we think of x^* as the *independent* variable, and then compute $\beta = x^*/\tanh x^*$. This gives us a table of pairs (x^*, For each pair, we plot β horizontally and x^* vertically. This yields the bifurcation diagram (Figure 3.4.4).

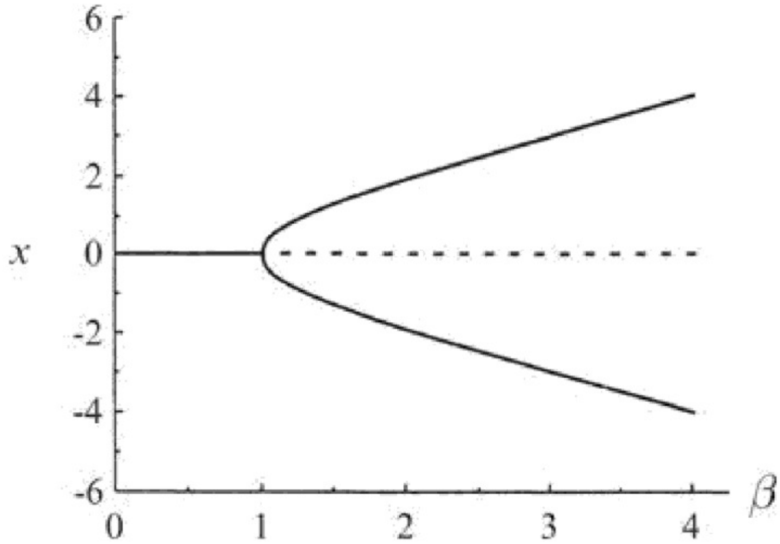

Figure 3.4.4

The shortcut used here exploits the fact that $f(x, \beta) = -x + \beta \tanh x$ depends more simply on β than on x. This is frequently the case in bifurcation problems—the dependence on the control parameter is usually simpler than the dependence on x. ■

EXAMPLE 3.4.2:

Plot the potential $V(x)$ for the system $\quad = rx - x^3$, for the cases $r < 0$, $r = 0$, and $r > 0$.

Solution: Recall from Section 2.7 that the potential for $x = f(x)$ is defined by $f(x) = -dV/dx$. Hence we need to solve $-dV/dx = rx - x^3$. Integration yields V(x)=−12rx2+14x4, where we neglect the arbitrary constant of integration. The corresponding graphs are shown in Figure 3.4.5.

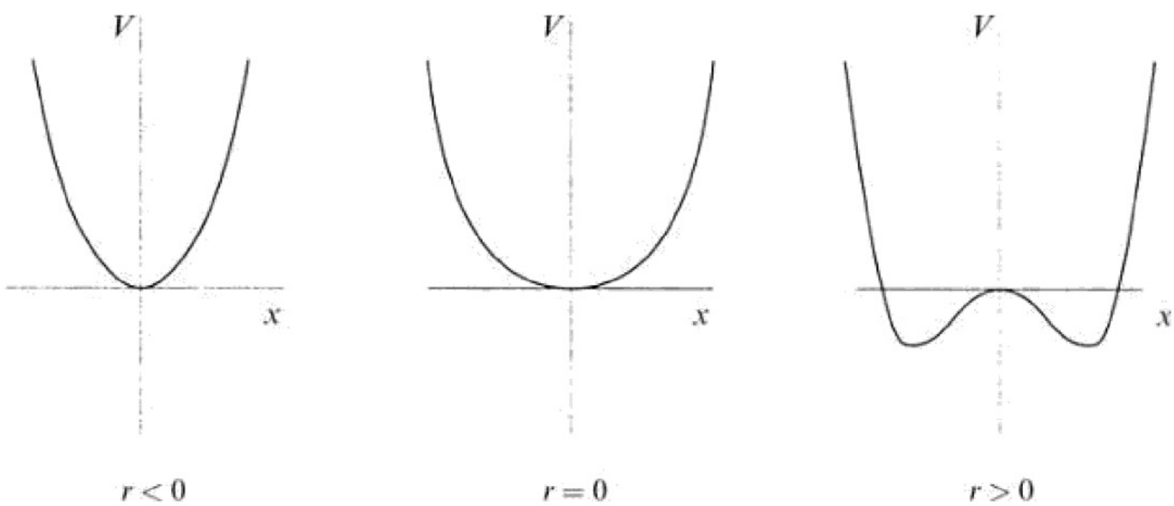

Figure 3.4.5

When $r < 0$, there is a quadratic minimum at the origin. At the bifurcation value $r = 0$, the minimum becomes a much flatter quartic. For $r > 0$, a local *maximum* appears at the origin, and a symmetric pair of minima occur to either side of it. ■

Subcritical Pitchfork Bifurcation

In the supercritical case $= rx - x^3$ discussed above, the cubic term is *stabilizing:* it acts as a restoring force that pulls $x(t)$ back toward $x = 0$. If instead the cubic term were *destabilizing*, as in

x˙=rx+x3, *(2)*

then we'd have a ***subcritical*** pitchfork bifurcation. Figure 3.4.6 shows the bifurcation diagram.

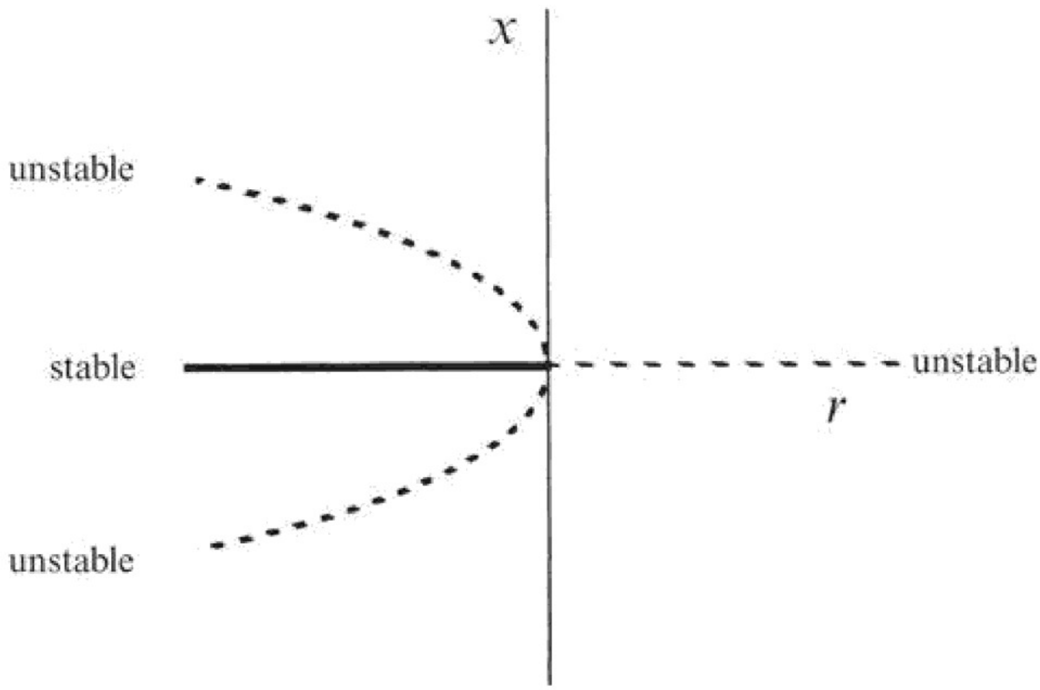

Figure 3.4.6

Compared to Figure 3.4.2, the pitchfork is inverted. The nonzero fixed points x*=±−r are *unstable*, and exist only *below* the bifurcation ($r < 0$), which motivates the term "subcritical." More importantly, the origin is stable for $r < 0$ and unstable for $r > 0$, as in the supercritical case, but now the instability for $r > 0$ is not opposed by the cubic term—in fact the cubic term lends a helping hand in driving the trajectories out to infinity! This effect leads to *blow-up:* one can show that x(t) →±∞ in finite time, starting from any initial condition $x_0 \neq 0$ (Exercise 2.5.3).

In real physical systems, such an explosive instability is usually opposed by the stabilizing influence of higher-order terms. Assuming that the system is still symmetric under $x \to -x$, the first stabilizing term must be x^5. Thus the canonical example of a system with a subcritical pitchfork bifurcation is

x˙=rx+x3−x5. *(3)*

There's no loss in generality in assuming that the coefficients of x^3 and x^5 are 1 (Exercise 3.5.8).

The detailed analysis of (3) is left to you (Exercises 3.4.14 and 3.4.15). But we will summarize the main results here. Figure 3.4.7 shows the bifurcation diagram for (3). For small x, the picture looks just like Figure 3.4.6: the origin is locally stable for $r < 0$, and two backward-bending branches of unstable fixed points bifurcate from the origin when $r = 0$. The new feature, due to the x^5 term, is that the unstable branches turn around and become stable at $r = r$, where $r < 0$. These stable ***large-amplitude*** branches exist for all $r > r_s$. There are several things to note about Figure 3.4.7:

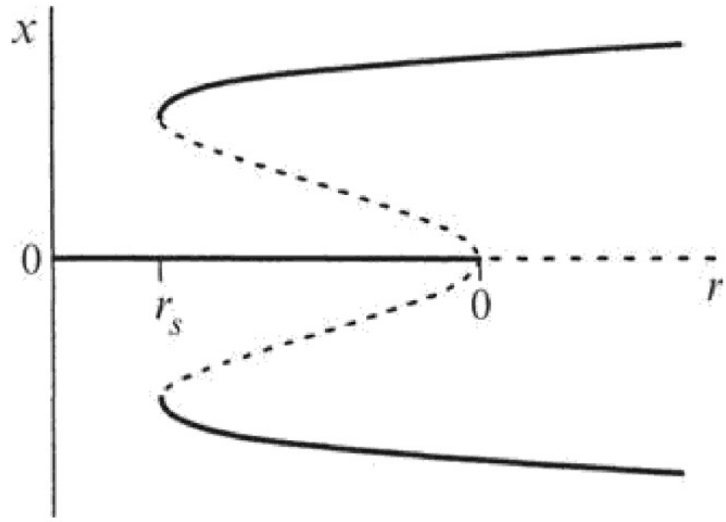

Figure 3.4.7

1. In the range $r_s < r < 0$, two qualitatively different stable states coexist, namely the origin and the large-amplitude fixed points. The initial condition x_0 determines which fixed point is approached as $t \to \infty$. One consequence is that the origin is stable to small perturbations, but not to large ones—in this sense the origin is *locally* stable, but not *globally* stable.
2. The existence of different stable states allows for the possibility of ***jumps*** and ***hysteresis*** as r is varied. Suppose we start the system in the state $x^* = 0$, and then slowly increase the parameter r (indicated by an arrow along the r-axis of Figure 3.4.8).

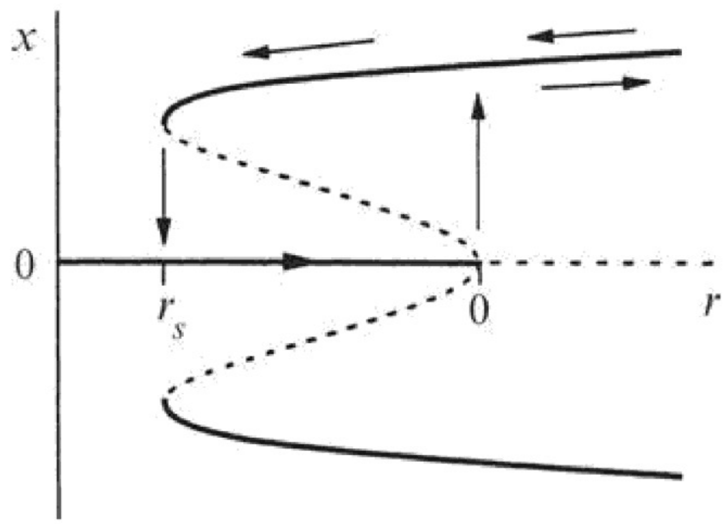

Figure 3.4.8

Then the state remains at the origin until $r = 0$, when the origin loses stability. Now the slightest nudge will cause the state to *jump* to one of the large-amplitude branches. With further increases of r, the state moves out along the large-amplitude branch. If r is now decreased, the state remains on the large-amplitude branch, even when r is decreased below 0! We have to lower r even further (down past r_s) to get the state to jump back to the origin. This lack of reversibility as a parameter is varied is called *hysteresis*.

3. The bifurcation at r is a saddle-node bifurcation, in which stable and unstable fixed points are born "out of the clear blue sky" as r is increased (see Section 3.1).

Terminology

As usual in bifurcation theory, there are several other names for the bifurcations discussed here. The supercritical pitchfork is sometimes called a forward bifurcation, and is closely related to a continuous or second-order phase transition in statistical mechanics. The subcritical bifurcation is sometimes called an inverted or backward bifurcation, and is related to discontinuous or first-order phase transitions. In the engineering literature, the supercritical bifurcation is sometimes called soft or safe, because the nonzero fixed points are born at small amplitude; in contrast, the subcritical bifurcation is hard or dangerous, because of the jump from zero to large amplitude.

3.5 Overdamped Bead on a Rotating Hoop

In this section we analyze a classic problem from first-year physics, the bead on a rotating hoop. This problem provides an example of a bifurcation in a mechanical system. It also illustrates the subtleties involved in replacing Newton's law, which is a second-order equation, by a simpler first-order equation.

The mechanical system is shown in Figure 3.5.1. A bead of mass m slides along a wire hoop of radius r. The hoop is constrained to rotate at a constant angular velocity u about its vertical axis. The problem is to analyze the motion of the bead, given that it is acted on by both gravitational and centrifugal forces. This is the usual statement of the problem, but now we want to add a new twist: suppose that there's also a frictional force on the bead that opposes its motion. To be specific, imagine that the whole system is immersed in a vat of molasses or some other very viscous fluid, and that the friction is due to viscous damping.

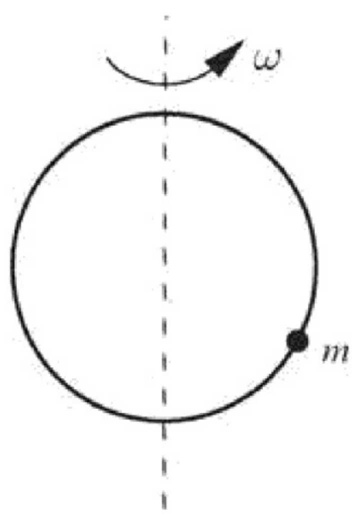

Figure 3.5.1

Let ϕ be the angle between the bead and the downward vertical direction. By convention, we restrict ϕ to the range $-\pi<\phi\leq\pi$, so there's only one angle for each point on the hoop. Also, let $\quad = r \sin\phi$ denote the distance of the bead from the vertical axis. Then the coordinates are as shown in Figure 3.5.2.

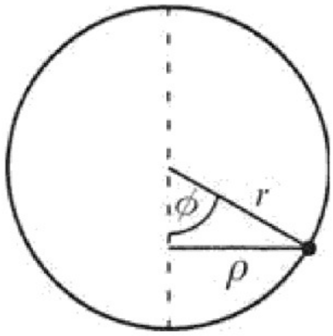

Figure 3.5.2

Now we write Newton's law for the bead. There's a downward gravitational force *mg*, a sideways centrifugal force *m* ^{2}, and a tangential damping force $b\dot{\phi}$. (The constants *g* and *b* are taken to be positive; negative signs will be added later as needed.) The hoop is assumed to be rigid, so we only have to resolve the forces along the tangential direction, as shown in Figure 3.5.3. After substituting = *r* sin in the centrifugal term, and recalling that the tangential acceleration is $r\ddot{\phi}$, we obtain the governing equation

$$mr\ddot{\phi} = -b\dot{\phi} - mg\sin\phi + mr\omega^2\sin\phi\cos\phi. \qquad (1)$$

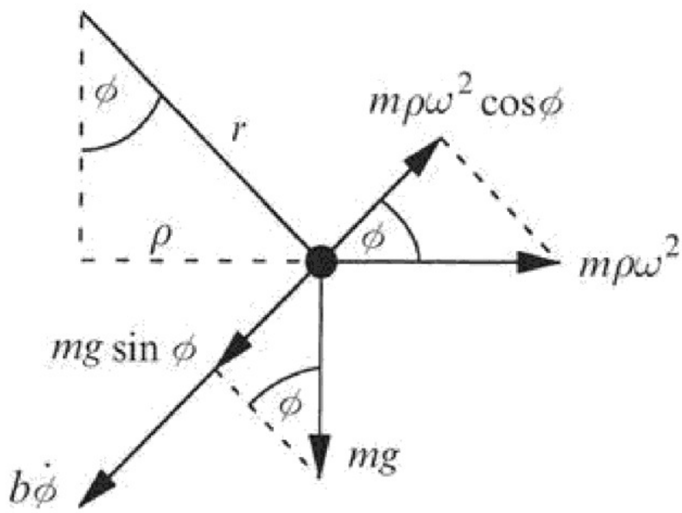

Figure 3.5.3

This is a *second-order* differential equation, since the second derivative $\ddot{\phi}$ is the highest one that appears. We are not yet equipped to analyze second-order equations, so we would like to find some conditions under which we can safely neglect the $mr\ddot{\phi}$ term. Then (1) reduces to a first-order equation, and we can apply our machinery to it.

Of course, this is a dicey business: we can't just neglect terms because we feel like it! But we will for now, and then at the end of this section we'll try to find a regime where our approximation is valid.

Analysis of the First-Order System

Our concern now is with the first-order system

$$b\dot{\phi} = -mg\sin\phi + mr\omega^2\sin\phi\cos\phi. = mg\sin\phi\left(\frac{r\omega^2}{g}\cos\phi - 1\right). \qquad (2)$$

The fixed points of (2) correspond to equilibrium positions for the bead. What's your intuition about where such equilibria can occur? We would expect the bead to remain at rest if placed at the top or the bottom of the hoop. Can other fixed points occur? And what about stability? Is the bottom always stable?

Equation (2) shows that there are always fixed points where sin = 0, namely * = 0 (the bottom of the hoop) and * = π (the top). The more interesting result is that there are two *additional* fixed points if

rω2g>1,

that is, *if the hoop is spinning fast enough.* These fixed points satisfy * = ±$\cos^{-1}$ (*g*/*r* ^{2}). To visualize them, we introduce a parameter

γ=rω2g

and solve cos * = 1/ graphically. We plot cos vs. , and look for intersections with the constant function l/ , shown as a horizontal line in Figure 3.5.4. For < 1 there are no intersections, whereas for > 1 there is a symmetrical pair of intersections to either side of * = 0.

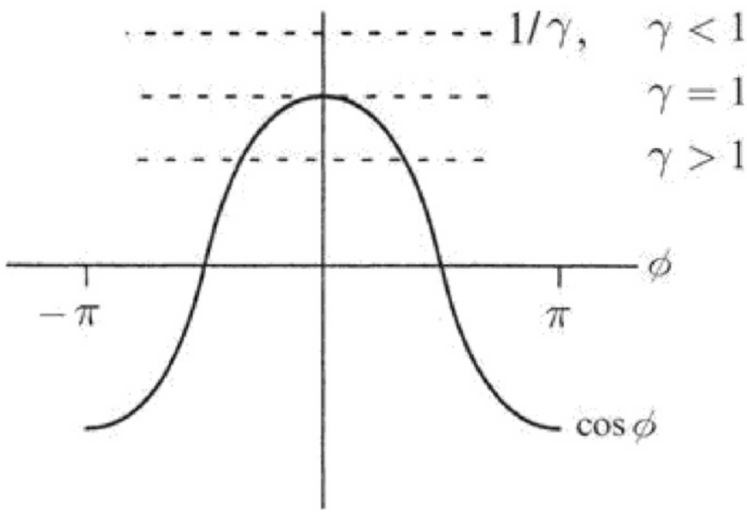

Figure 3.5.4

As → ∞, these intersections approach ±π/2. Figure 3.5.5 plots the fixed points on the hoop for the cases γ < 1 and γ > 1.

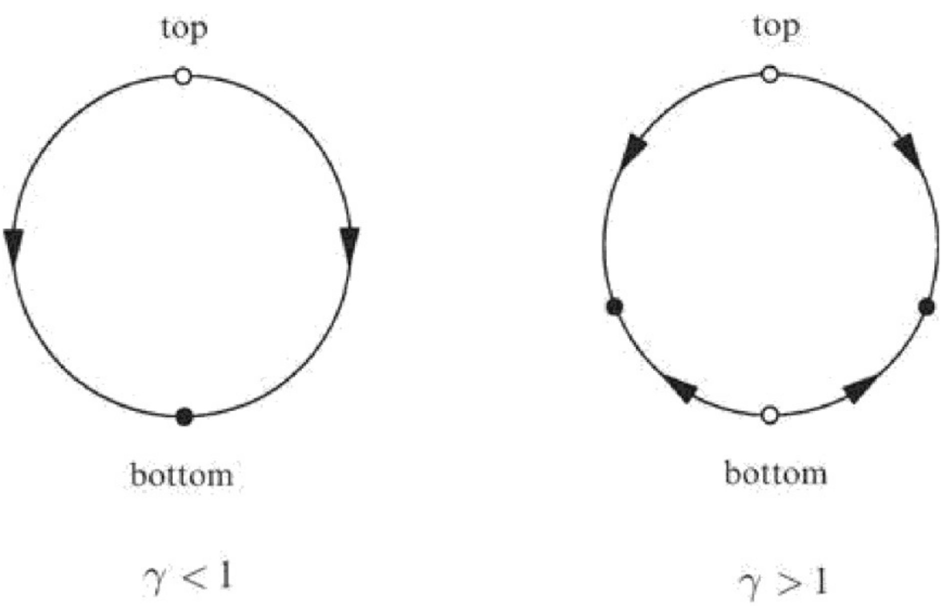

Figure 3.5.5

To summarize our results so far, let's plot *all* the fixed points as a function of the parameter γ (Figure 3.5.6). As usual, solid lines denote stable fixed points and broken lines denote unstable fixed points.

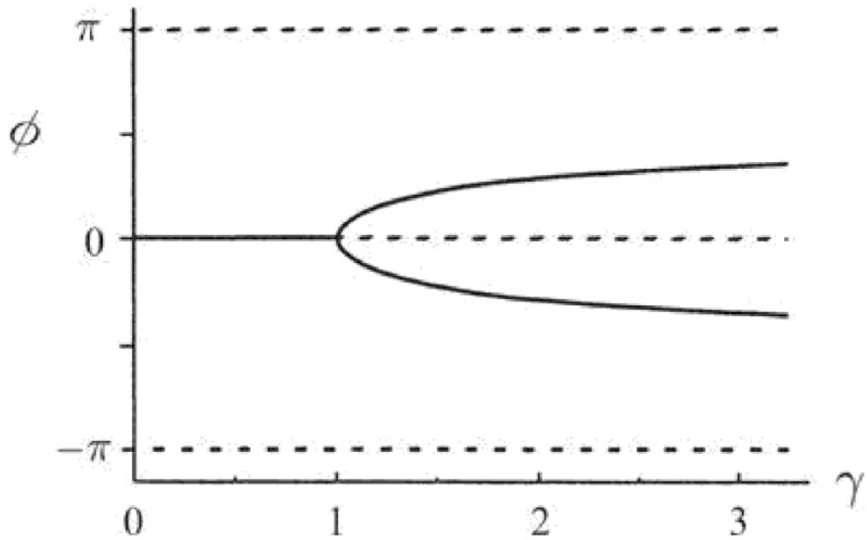

Figure 3.5.6

We now see that a ***supercritical pitchfork bifurcation*** occurs at $\gamma = 1$. It's left to you to check the stability of the fixed points, using linear stability analysis or graphical methods (Exercise 3.5.2).

Here's the physical interpretation of the results: When $\gamma < 1$, the hoop is rotating slowly and the centrifugal force is too weak to balance the force of gravity. Thus the bead slides down to the bottom and stays there. But if $\gamma > 1$, the hoop is spinning fast enough that the bottom becomes unstable. Since the centrifugal force *grows* as the bead moves farther from the bottom, any slight displacement of the bead will be *amplified.* The bead is therefore pushed up the hoop until gravity balances the centrifugal force; this balance occurs at $\;^* = \pm\cos^{-1}(g/r\quad^2)$. Which of these two fixed points is actually selected depends on the initial disturbance. Even though the two fixed points are entirely symmetrical, an asymmetry in the initial conditions will lead to one of them being chosen—physicists sometimes refer to these as ***symmetry-broken*** solutions. In other words, the solution has less symmetry than the governing equation.

What *is* the symmetry of the governing equation? Clearly the left and right halves of the hoop are physically equivalent—this is reflected by the invariance of (1) and (2) under the change of variables $\rightarrow -$. As we mentioned in Section 3.4, pitchfork bifurcations are to be expected in situations where such a symmetry exists.

Dimensional Analysis and Scaling

Now we need to address the question: When is it valid to neglect the inertia term $mr\ddot{\phi}$ in (1)? At first sight the limit $m \rightarrow 0$ looks promising, but then we notice that we're throwing out the baby with the bathwater: the centrifugal and gravitational terms vanish in this limit too! So we have to be more careful.

In problems like this, it is helpful to express the equation in ***dimensionless*** form (at present, all the terms in (1) have the dimensions of force.) The advantage of a dimensionless formulation is that we know how to define *small*—it means "much less than 1." Furthermore, nondimensionalizing the equation reduces the number of parameters by lumping them together into ***dimensionless groups.*** This reduction always simplifies the analysis. For an excellent introduction to dimensional analysis, see Lin and Segel (1988).

There are often several ways to nondimensionalize an equation, and the best choice might not be clear at first. Therefore we proceed in a flexible fashion. We define a dimensionless time τ by

$$\tau = \frac{t}{T}$$

where T is a ***characteristic time scale*** to be chosen later. When T is chosen correctly, the new derivatives $d\phi/d\tau$ and $d^2\phi/d\tau^2$ should be $O(1)$, i.e., of order unity. To express these new derivatives in terms of the old ones, we use the chain rule:

$$\dot{\phi} = \frac{d\phi}{dt} = \frac{d\phi}{d\tau}\frac{d\tau}{dt} = \frac{1}{T}\frac{d\phi}{d\tau}$$

and similarly

$$\ddot{\phi} = \frac{1}{T^2}\frac{d^2\phi}{d\tau^2}.$$

(The easy way to remember these formulas is to formally substitute $T\tau$ for t.) Hence (1) becomes

$$\frac{mr}{T^2}\frac{d^2\phi}{d\tau^2} = -\frac{b}{T}\frac{d\phi}{d\tau} - mg\sin\phi + mr\omega^2\sin\phi\cos\phi.$$

Now since this equation is a balance of forces, we nondimensionalize it by dividing by a force mg. This yields the dimensionless equation

$$\left(\frac{r}{gT^2}\right)\frac{d^2\phi}{d\tau^2} = -\left(\frac{b}{mgT}\right)\frac{d\phi}{d\tau} - \sin\phi + \left(\frac{r\omega^2}{g}\right)\sin\phi\cos\phi. \qquad (3)$$

Each of the terms in parentheses is a dimensionless group. We recognize the group $r\omega^2/g$ in the last term—that's our old friend γ from earlier in the section.

We are interested in the regime where the left-hand side of (3) is negligible compared to all the other terms, and where all the terms on the right-hand side are of comparable size. Since the derivatives are $O(1)$ by assumption, and $\sin\phi \approx O(1)$, we see that we need

$$\frac{b}{mgT} \approx O(1), \quad \text{and} \quad \frac{r}{gT^2} \ll 1.$$

The first of these requirements sets the time scale T: a natural choice is

$$T = \frac{b}{mg}.$$

Then the condition $r/gT^2 \ll 1$ becomes

$$\frac{r}{g}\left(\frac{mg}{b}\right)^2 \ll 1, \qquad (4)$$

or equivalently,

$$b^2 \gg m^2 g r.$$

This can be interpreted as saying that the *damping is very strong*, or that the mass is very small, now in a precise sense.

The condition (4) motivates us to introduce a dimensionless group

$$\varepsilon = \frac{m^2 g r}{b^2}. \qquad (5)$$

Then (3) becomes

$$\varepsilon \frac{d^2\phi}{d\tau^2} = -\frac{d\phi}{d\tau} - \sin\phi + \gamma \sin\phi\cos\phi. \qquad (6)$$

As advertised, the dimensionless Equation (6) is simpler than (1): the five parameters m, g, r, , and *b* have been replaced by two dimensionless groups γ and .

In summary, our dimensional analysis suggests that in the ***overdamped*** limit → 0, (6) should be well approximated by the first-order system

$$\frac{d\phi}{d\tau} = f(\phi) \qquad (7)$$

where

$$f(\phi) = -\sin\phi + \gamma \sin\phi\cos\phi = \sin\phi(\gamma\cos\phi - 1).$$

A Paradox

Unfortunately, *there is something fundamentally wrong with our idea of replacing a second-order equation by a first-order equation.* The trouble is that a second-order equation requires *two* initial conditions, whereas a first-order equation has only *one.* In our case, the bead's motion is determined by its initial position and velocity. These two quantities can be chosen completely independent of each other. But that's not true for the first-order system: given the initial position, the initial velocity is dictated by the equation *d /d* = *f* ( ). Thus the solution to the first-order system will not, in general, be able to satisfy *both* initial conditions.

We seem to have run into a paradox. Is (7) valid in the overdamped limit or not? If it is valid, how can we satisfy the two arbitrary initial conditions demanded by (6)?

The resolution of the paradox requires us to analyze the second-order system (6). We haven't dealt with second-order systems before—that's the subject of Chapter 5. But read on if you're curious; some simple ideas are all we need to finish the problem.

Phase Plane Analysis

Throughout Chapters 2 and 3, we have exploited the idea that a first-order system = *f* (*x*) can be regarded as a vector field on a line. By analogy, the *second*-order system (6) can be regarded as a vector field on a *plane*, the so-called ***phase plane.***

The plane is spanned by two axes, one for the angle and one for the angular velocity *d /d* . To simplify the notation, let

$$\Omega = \phi' \equiv d\phi/d\tau$$

where prime denotes differentiation with respect to . Then an initial condition for (6) corresponds to a point ($_0$, Ω_0) in the phase plane (Figure 3.5.7). As time evolves, the phase point ((*t*), $\Omega(t)$) moves around in the phase plane along a ***trajectory*** determined by the solution to (6).

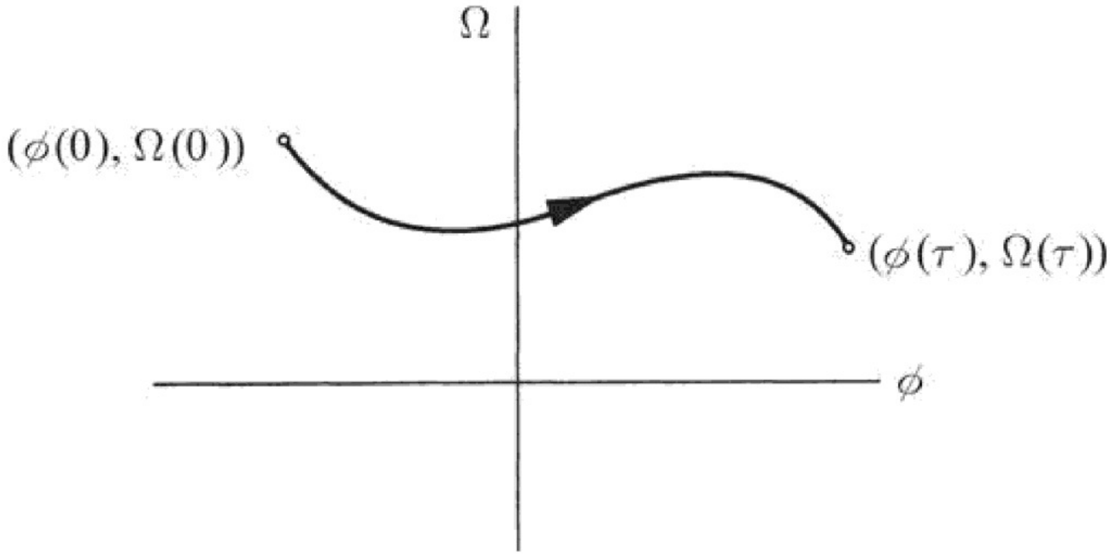

Figure 3.5.7

Our goal now is to see what those trajectories actually look like. As before, the key idea is that *the differential equation can be interpreted as a vector field on the phase space.* To convert (6) into a vector field, we first rewrite it as

$$\varepsilon\,\Omega' = f(\phi) - \Omega.$$

Along with the definition $\;'$ = Ω, this yields the ***vector field***

$$\phi' = \Omega \tag{8a}$$

$$\Omega' = 1\,\varepsilon\,(f(\phi) - \Omega). \tag{8b}$$

We interpret the vector (, Ω') at the point (, Ω) as the local velocity of a phase fluid flowing steadily on the plane. Note that the velocity vector now has two components, one in the -direction and one in the Ω-direction. To visualize the trajectories, we just imagine how the phase point would move as it is carried along by the phase fluid.

In general, the pattern of trajectories would be difficult to picture, but the present case is simple because we are only interested in the limit $\;\to 0$. In this limit, *all trajectories slam straight up or down onto the curve C defined by f* ( ) = Ω, *and then slowly ooze along this curve until they reach a fixed point* (Figure 3.5.8).

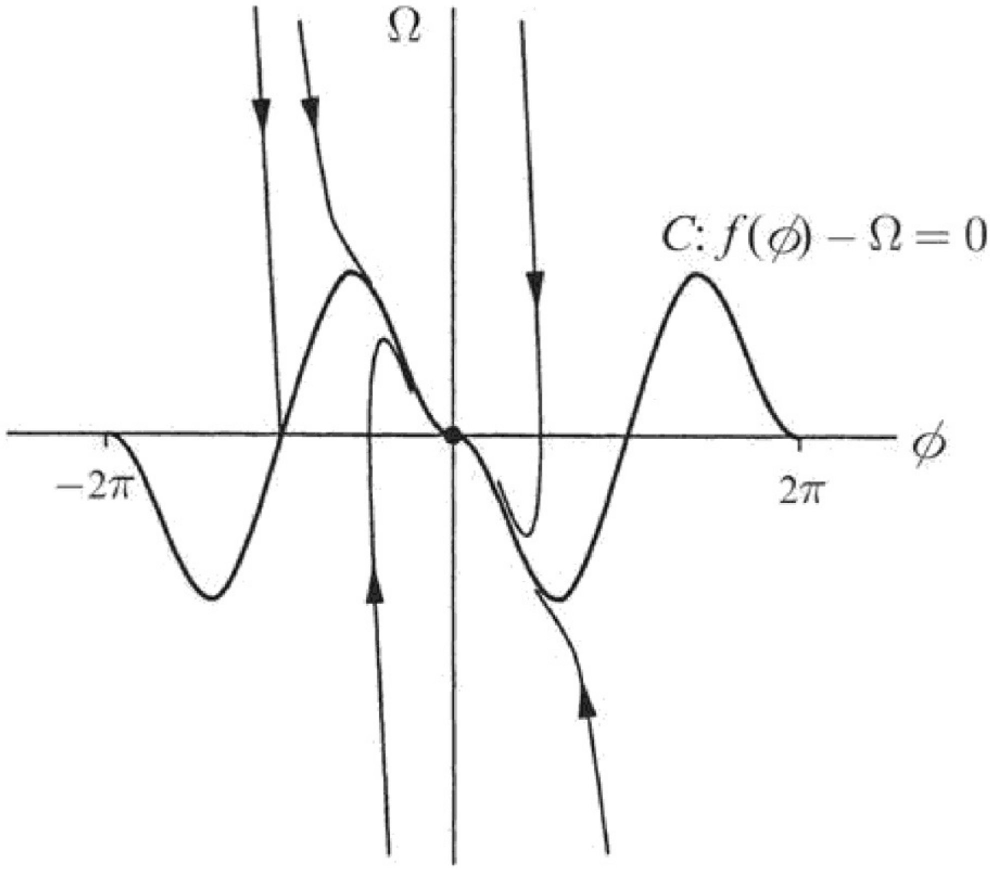

Figure 3.5.8

To arrive at this striking conclusion, let's do an order-of-magnitude calculation. Suppose that the phase point lies off the curve C. For instance, suppose (, Ω) lies an $O(1)$ distance below the curve C, i.e., $\Omega < f(\ )$ and $f(\ ) - \Omega \approx O(1)$. Then (8b) shows that Ω' is enormously positive: $\Omega' \approx O(1/\ ) \gg 1$. Thus the phase point zaps like lightning up to the region where $f(\ ) - \Omega \approx O(\ )$. In the limit $\ \to 0$, this region is indistinguishable from C. Once the phase point is on C, it evolves according to $\Omega \approx f(\ )$; that is, it approximately satisfies the first-order equation $' = f(\ )$.

Our conclusion is that a typical trajectory is made of two parts: a rapid initial ***transient,*** during which the phase point zaps onto the curve where $' = f(\ )$, followed by a much slower drift along this curve.

Now we see how the paradox is resolved: The second-order system (6) *does* behave like the first-order system (7), but only after a rapid initial transient. During this transient, it is *not* correct to neglect the term $d^2\ /d\ ^2$. The problem with our earlier approach is that we used only a single time scale $T = b/mg$; this time scale is characteristic of the slow drift process, but not of the rapid transient (Exercise 3.5.5).

A Singular Limit

The difficulty we have encountered here occurs throughout science and engineering. In some limit of interest (here, the limit of strong damping), the term containing the highest order derivative drops out of the governing equation. Then the initial conditions or boundary conditions can't be satisfied. Such a limit is often called *singular*. For example, in fluid mechanics, the limit of high Reynolds number is a singular limit; it accounts for the presence of extremely thin "boundary layers" in the flow over airplane wings. In our problem, the rapid transient played the role of a boundary layer—it is a thin layer of *time* that occurs near the boundary $t = 0$.

The branch of mathematics that deals with singular limits is called *singular perturbation theory.* See Jordan and Smith (1987) or Lin and Segel (1988) for an introduction. Another problem with a singular limit will be discussed briefly in Section 7.5.

3.6 Imperfect Bifurcations and Catastrophes

As we mentioned earlier, pitchfork bifurcations are common in problems that have a symmetry. For example, in the problem of the bead on a rotating hoop (Section 3.5), there was a perfect symmetry between the left and right sides of the hoop. But in many real-world circumstances, the symmetry is only approximate—an imperfection leads to a slight difference between left and right. We now want to see what happens when such imperfections are present.

For example, consider the system

$$\dot{x} = h + rx - x^3. \qquad (1)$$

If $h = 0$, we have the normal form for a supercritical pitchfork bifurcation, and there's a perfect symmetry between x and $-x$. But this symmetry is broken when $h \neq 0$; for this reason we refer to h as an ***imperfection parameter.***

Equation (1) is a bit harder to analyze than other bifurcation problems we've considered previously, because we have *two* independent parameters to worry about (h and r). To keep things straight, we'll think of r as fixed, and then examine the effects of varying h. The first step is to analyze the fixed points of (1). These can be found explicitly, but we'd have to invoke the messy formula for the roots of a cubic equation. It's clearer to use a graphical approach, as in Example 3.1.2. We plot the graphs of $y = rx - x^3$ and $y = -h$ on the same axes, and look for intersections (Figure 3.6.1). These intersections occur at the fixed points of (1). When $r \leq 0$, the cubic is monotonically decreasing, and so it intersects the horizontal line $y = -h$ in exactly one point (Figure 3.6.1a). The more interesting case is $r > 0$; then one, two, or three intersections are possible, depending on the value of h (Figure 3.6.1b).

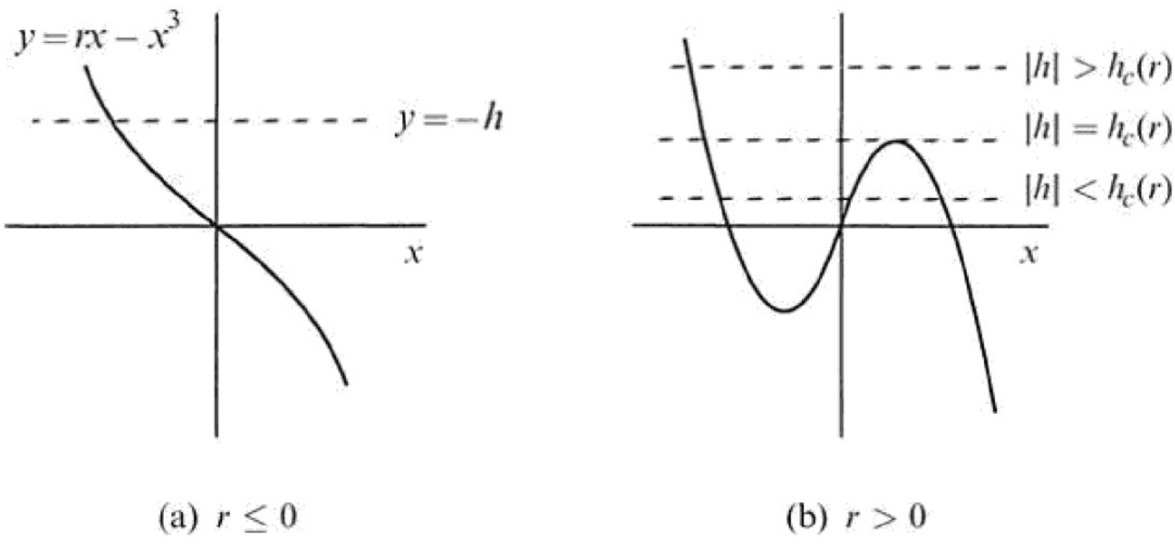

Figure 3.6.1

The critical case occurs when the horizontal line is just *tangent* to either the local minimum or maximum of the cubic; then we have a *saddle-node bifurcation.* To find the values of h at which this bifurcation occurs, note that the cubic has a local maximum when $\frac{d}{dx}(rx - x^3) = r - 3x^2 = 0$. Hence

$$x_{max} = \sqrt{\frac{r}{3}},$$

and the value of the cubic at the local maximum is

$$rxmax-(xmax)3=2r3r3.$$

Similarly, the value at the minimum is the negative of this quantity. Hence saddle-node bifurcations occur when $h = \pm h_c(r)$, where

$$hc(r)=2r2r3.$$

Equation (1) has three fixed points for $|h| < h_c(r)$ and one fixed point for $|h| > h_c(r)$.

To summarize the results so far, we plot the ***bifurcation curves*** $h = \pm h_c(r)$ in the (r, h) plane (Figure 3.6.2). Note that the two bifurcation curves meet tangentially at $(r, h) = (0,0)$; such a point is called a ***cusp point.*** We also label the regions that correspond to different numbers of fixed points. Saddle-node bifurcations occur all along the boundary of the regions, except at the cusp point, where we have a *codimension-2 bifurcation.* (This fancy terminology essentially means that we have had to tune *two* parameters, h and r, to achieve this type of bifurcation. Until now, all our bifurcations could be achieved by tuning a single parameter, and were therefore *codimension-1* bifurcations.)

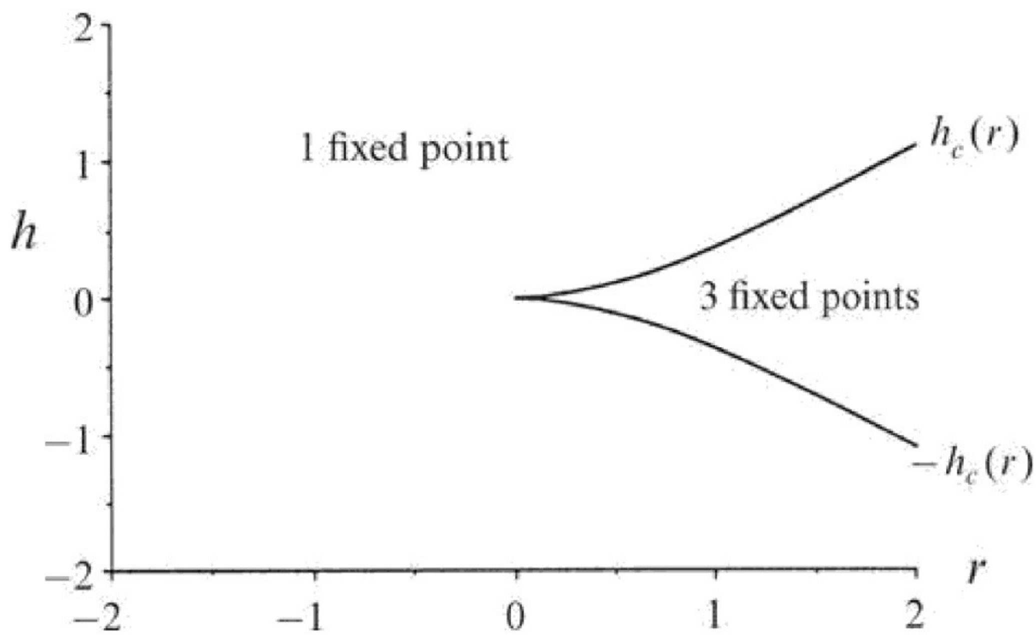

Figure 3.6.2

Pictures like Figure 3.6.2 will prove very useful in our future work. We will refer to such pictures as ***stability diagrams.*** They show the different types of behavior that occur as we move around in ***parameter space*** (here, the (r, h) plane).

Now let's present our results in a more familiar way by showing the bifurcation diagram of x^* vs. r, for fixed h (Figure 3.6.3).

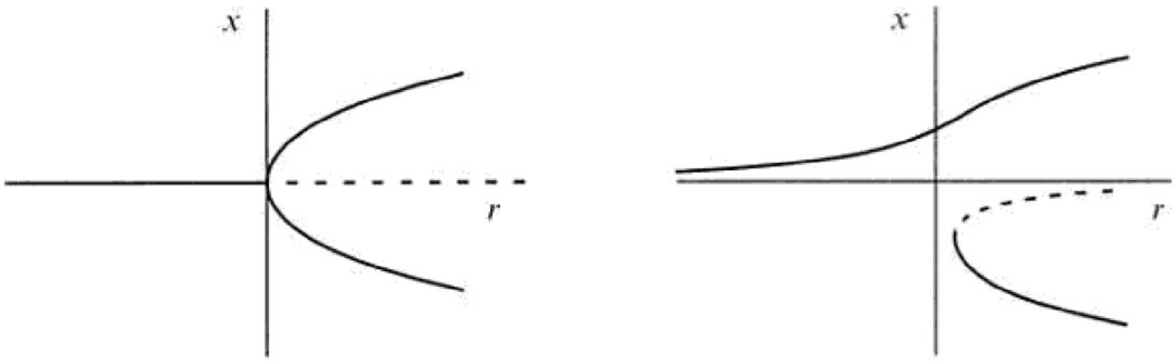

(a) $h = 0$ (b) $h \neq 0$

Figure 3.6.3

When $h = 0$ we have the usual pitchfork diagram (Figure 3.6.3a) but when $h \neq 0$, the pitchfork disconnects into two pieces (Figure 3.6.3b). The upper piece consists entirely of stable fixed points, whereas the lower piece has both stable and unstable branches. As we increase r from negative values, there's no longer a sharp transition at $r = 0$; the fixed point simply glides smoothly along the upper branch. Furthermore, the lower branch of stable points is not accessible unless we make a fairly large disturbance.

Alternatively, we could plot x^* vs. h, for fixed r (Figure 3.6.4).

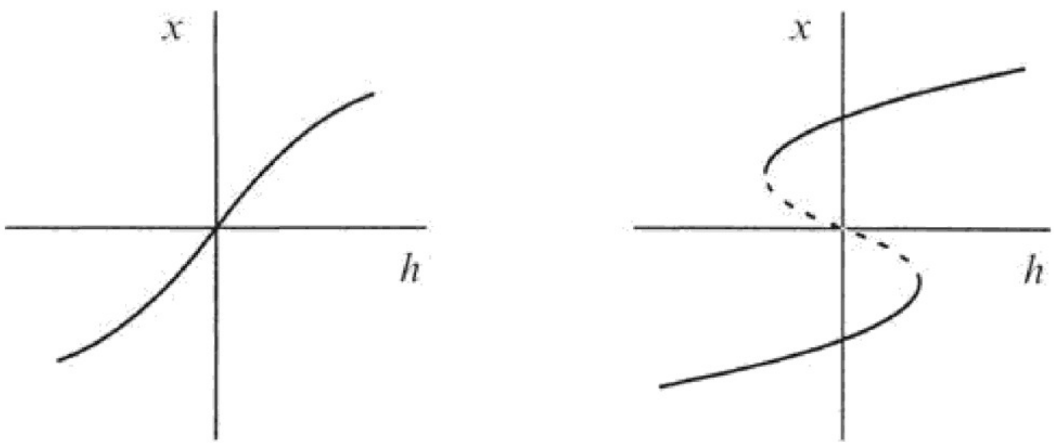

(a) $r \leq 0$ (b) $r > 0$

Figure 3.6.4

When $r \leq 0$ there's one stable fixed point for each h (Figure 3.6.4a). However, when $r > 0$ there are three fixed points when\h\ < $h_c(r)$, and one otherwise (Figure 3.6.4b). In the triple-valued region, the middle branch is unstable and the upper and lower branches are stable. Note that these graphs look like Figure 3.6.1 rotated by 90°.

There is one last way to plot the results, which may appeal to you if you like to picture things in three dimensions. This method of presentation contains all of the others as cross sections or projections. If we plot the fixed points x^* above the (r, h) plane, we get the ***cusp catastrophe*** surface shown in Figure 3.6.5. The surface folds over on itself in certain places. The projection of these folds onto the (r, h) plane yields the bifurcation curves shown in Figure 3.6.2. A cross section at fixed h yields Figure 3.6.3, and a cross section at fixed r yields Figure 3.6.4.

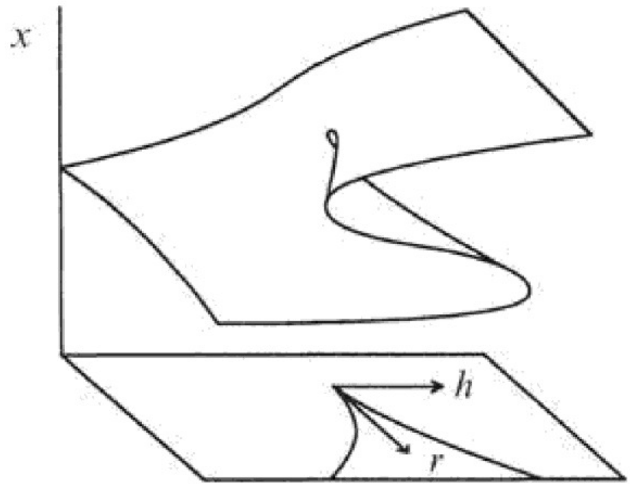

Figure 3.6.5

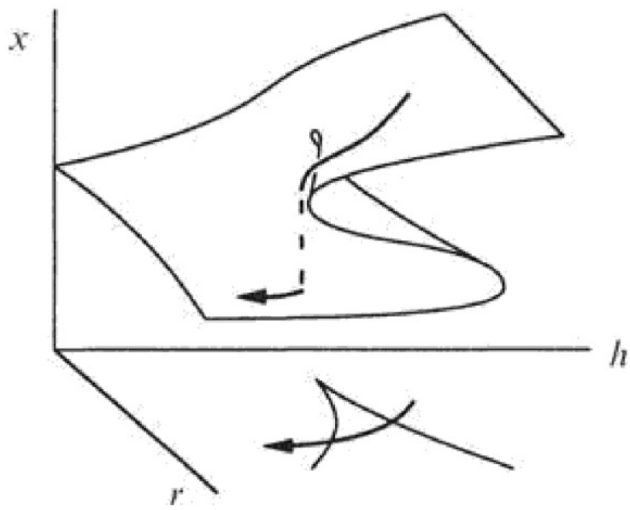

Figure 3.6.6

The term *catastrophe* is motivated by the fact that as parameters change, the state of the system can be carried over the edge of the upper surface, after which it drops discontinuously to the lower surface (Figure 3.6.6). This jump could be truly catastrophic for the equilibrium of a bridge or a building. We will see scientific examples of catastrophes in the context of insect outbreaks (Section 3.7) and in the following example from mechanics.

For more about catastrophe theory, see Zeeman (1977) or Poston and Stewart (1978). Incidentally, there was a violent controversy about this subject in the late 1970s. If you like watching fights, have a look at Zahler and Sussman (1977) and Kolata (1977).

Bead on a Tilted Wire

As a simple example of imperfect bifurcation and catastrophe, consider the following mechanical system (Figure 3.6.7).

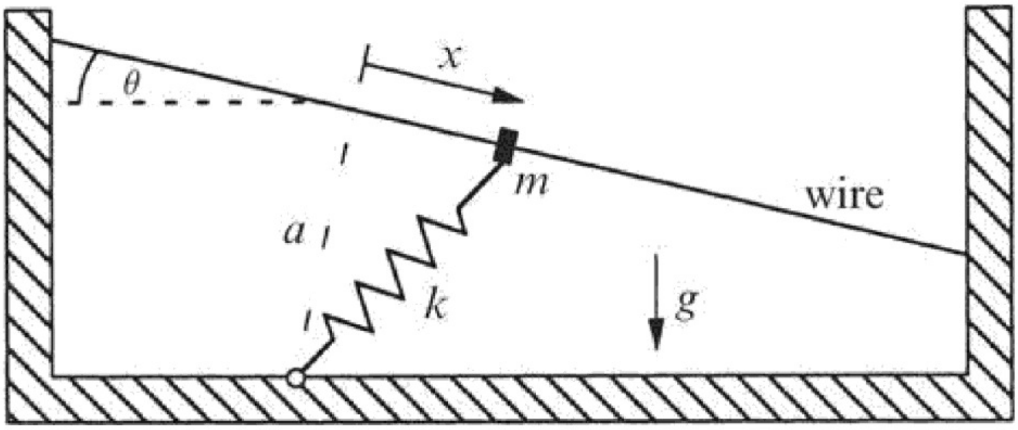

Figure 3.6.7

A bead of mass m is constrained to slide along a straight wire inclined at an angle   with respect to the horizontal. The mass is attached to a spring of stiffness k and relaxed length L_0, and is also acted on by gravity. We choose coordinates along the wire so that $x = 0$ occurs at the point closest to the support point of the spring; let a be the distance between this support point and the wire.

In Exercises 3.5.4 and 3.6.5, you are asked to analyze the equilibrium positions of the bead. But first let's get some physical intuition. When the wire is horizontal (= 0), there is perfect symmetry between the left and right sides of the wire, and $x = 0$ is always an equilibrium position. The stability of this equilibrium depends on the relative sizes of L_0 and a: if $L_0 < a$, the spring is in tension and so the equilibrium should be stable. But if $L_0 > a$, the spring is compressed and so we expect an *unstable* equilibrium at $x = 0$ and a pair of stable equilibria to either side of it. Exercise 3.5.4 deals with this simple case.

The problem becomes more interesting when we tilt the wire (≠ 0). For small tilting, we expect that there are still three equilibria if $L_0 > a$. However if the tilt becomes too steep, perhaps you can see intuitively that the uphill equilibrium might suddenly disappear, causing the bead to jump catastrophically to the downhill equilibrium. You might even want to build this mechanical system and try it. Exercise 3.6.5 asks you to work through the mathematical details.

3.7 Insect Outbreak

For a biological example of bifurcation and catastrophe, we turn now to a model for the sudden outbreak of an insect called the spruce budworm. This insect is a serious pest in eastern Canada, where it attacks the leaves of the balsam fir tree. When an outbreak occurs, the budworms can defoliate and kill most of the fir trees in the forest in about four years.

Ludwig et al. (1978) proposed and analyzed an elegant model of the interaction between budworms and the forest. They simplified the problem by exploiting a separation of time scales: the budworm population evolves on a *fast* time scale (they can increase their density fivefold in a year, so they have a characteristic time scale of months), whereas the trees grow and die on a *slow* time scale (they can completely replace their foliage in about 7–10 years, and their life span in the absence of budworms is 100–150 years.) Thus, as far as the budworm dynamics are concerned, the forest variables may be treated as constants. At the end of the analysis, we will allow the forest variables to drift very slowly—this drift ultimately triggers an outbreak.

Model

The proposed model for the budworm population dynamics is

$$\dot{N}=RN(1-NK)-p(N).$$

In the absence of predators, the budworm population $N(t)$ is assumed to grow logistically with growth rate R and carrying capacity K. The carrying capacity depends on the amount of foliage left on the trees, and so it is a slowly drifting parameter; at this stage we treat it as fixed. The term $p(N)$ represents the death rate due to *predation*, chiefly by birds, and is assumed to have the shape shown in Figure 3.7.1. There is almost no predation when budworms are scarce; the birds seek food elsewhere. However, once the population exceeds a certain critical level $N = A$, the predation turns on sharply and then saturates (the birds are eating as fast as they can). Ludwig et al. (1978) assumed the specific form

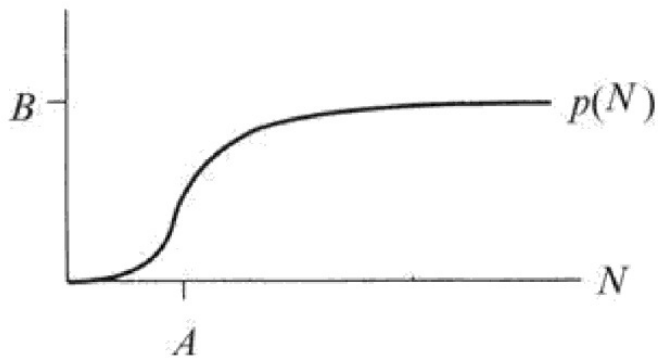

Figure 3.7.1

$$p(N)=BN2A2+N2$$

where $A, B > 0$. Thus the full model is

$$\dot{N}=RN(1-NK)-BN2A2+N2. \qquad (1)$$

We now have several questions to answer. What do we mean by an "outbreak" in the context of this model? The idea must be that, as parameters drift, the budworm population suddenly jumps from a low to a high level. But what do we mean by "low" and "high," and are there solutions with this character? To answer these questions, it is convenient to recast the model into a dimensionless form, as in Section 3.5.

Dimensionless Formulation

The model (1) has four parameters: R, K, A, and B. As usual, there are various ways to nondimensionalize the system. For example, both A and K have the same dimension as N, and so either N/A or N/K could serve as a dimensionless population level. It often takes some trial and error to find the best choice. In this case, our heuristic will be to scale the equation so that all the dimensionless groups are pushed into the *logistic* part of the dynamics, with none in the *predation* part. This turns out to ease the graphical analysis of the fixed points.

To get rid of the parameters in the predation term, we divide (1) by B and then let

$$x=N/A,$$

which yields

$$ABdxdt=RBAx(1-AxK)-x21+x2. \qquad (2)$$

Equation (2) suggests that we should introduce a dimensionless time and dimensionless groups r and k, as follows:

$$\tau = \frac{Bt}{A}, \quad r = \frac{RA}{B}, \quad k = \frac{K}{A}.$$

Then (2) becomes

$$\frac{dx}{d\tau} = rx\left(1 - \frac{x}{k}\right) - \frac{x^2}{1+x^2}, \tag{3}$$

which is our final dimensionless form. Here r and k are the dimensionless growth rate and carrying capacity, respectively.

Analysis of Fixed Points

Equation (3) has a fixed point at $x^* = 0$; it is *always unstable* (Exercise 3.7.1). The intuitive explanation is that the predation is extremely weak for small x, and so the budworm population grows exponentially for x near zero.

The other fixed points of (3) are given by the solutions of

$$r\left(1 - \frac{x}{k}\right) = \frac{x}{1+x^2}. \tag{4}$$

This equation is easy to analyze graphically—we simply graph the right- and left-hand sides of (4), and look for intersections (Figure 3.7.2). The left-hand side of (4) represents a straight line with x-intercept equal to k and a y-intercept equal to r, and the right-hand side represents a curve that is *independent of the parameters*! Hence, as we vary the parameters r and k, the line moves but the curve doesn't—this convenient property is what motivated our choice of nondimensionalization.

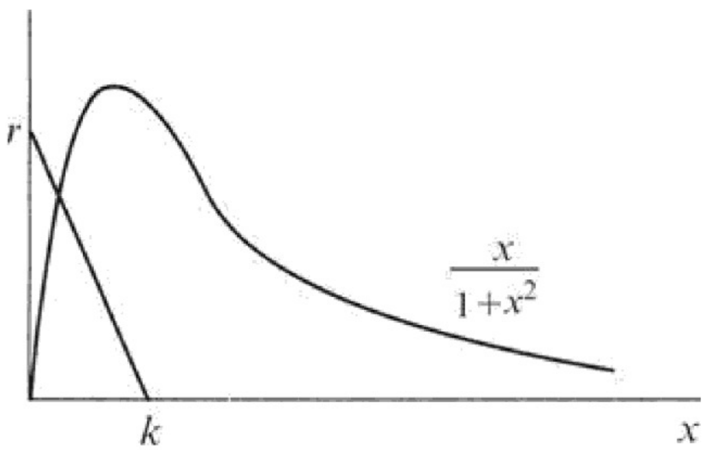

Figure 3.7.2

Figure 3.7.2 shows that if k is sufficiently small, there is exactly one intersection for any $r > 0$. However, for large k, we can have one, two, or three intersections, depending on the value of r (Figure 3.7.3). Let's suppose that there are three intersections *a, b*, and *c*. As we decrease r with k fixed, the line rotates counterclockwise about k. Then the fixed points b and c approach each other and eventually coalesce in a *saddle-node bifurcation* when the line intersects the curve *tangentially* (dashed line in Figure 3.7.3). After the bifurcation, the only remaining fixed point is a (in addition to $x^* = 0$, of course). Similarly, a and b can collide and annihilate as r is *increased.*

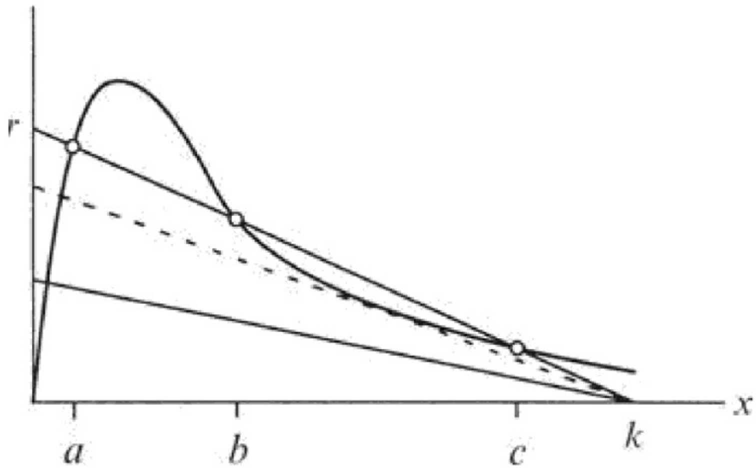

Figure 3.7.3

To determine the stability of the fixed points, we recall that $x^* = 0$ is unstable, and also observe that the stability type must alternate as we move along the x-axis.

Hence a is stable, b is unstable, and c is stable. Thus, for r and k in the range corresponding to three positive fixed points, the vector field is qualitatively like that shown in Figure 3.7.4. The smaller stable fixed point a is called the ***refuge*** level of the budworm population, while the larger stable point c is the ***outbreak*** level. From the point of view of pest control, one would like to keep the population at a and away from c. The fate of the system is determined by the initial condition x_0; an outbreak occurs if and only if $x_0 >$ b. In this sense the unstable equilibrium b plays the role of a ***threshold.***

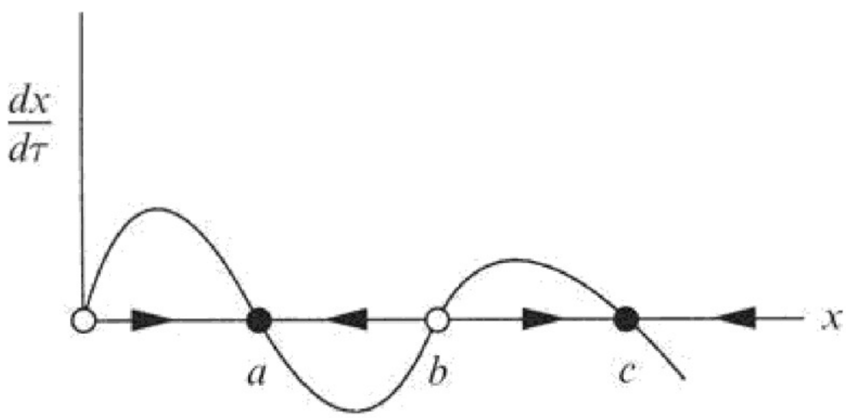

Figure 3.7.4

An outbreak can also be triggered by a saddle-node bifurcation. If the parameters r and k drift in such a way that the fixed point a disappears, then the population will jump suddenly to the outbreak level c. The situation is made worse by the hysteresis effect—even if the parameters are restored to their values before the outbreak, the population will not drop back to the refuge level.

Calculating the Bifurcation Curves

Now we compute the curves in (k, r) space where the system undergoes saddle-node bifurcations. The calculation is somewhat harder than that in Section 3.6: we will not be able to write r explicitly as a function of k, for example. Instead, the bifurcation curves will be written in the ***parametric form*** $(k\ (x), r\ (x))$, where x runs through all positive values. (Please don't be confused by this traditional terminology

—one would call x the “parameter” in these parametric equations, even though r and k are themselves parameters in a different sense.)

As discussed earlier, the condition for a saddle-node bifurcation is that the line $r(1 - x/k)$ intersects the curve $x/(1 + x^2)$ tangentially. Thus we require *both*

$$r(1-xk)=x1+x2 \qquad (5)$$

and

$$ddx[\ r(1-xk)\ ]=ddx[\ x1+x2\ ]. \qquad (6)$$

After differentiation, (6) reduces to

$$-rk=1-x2(1+x2)2. \qquad (7)$$

We substitute this expression for r/k into (5), which allows us to express r solely in terms of x. The result is

$$r=2x3(1+x2)2. \qquad (8)$$

Then inserting (8) into (7) yields

$$k=2x3x2-1. \qquad (9)$$

The condition $k > 0$ implies that x must be restricted to the range $x > 1$.

Together (8) and (9) define the bifurcation curves. For each $x > 1$, we plot the corresponding point $(k(x), r(x))$ in the (k, r) plane. The resulting curves are shown in Figure 3.7.5. (Exercise 3.7.2 deals with some of the analytical properties of these curves.)

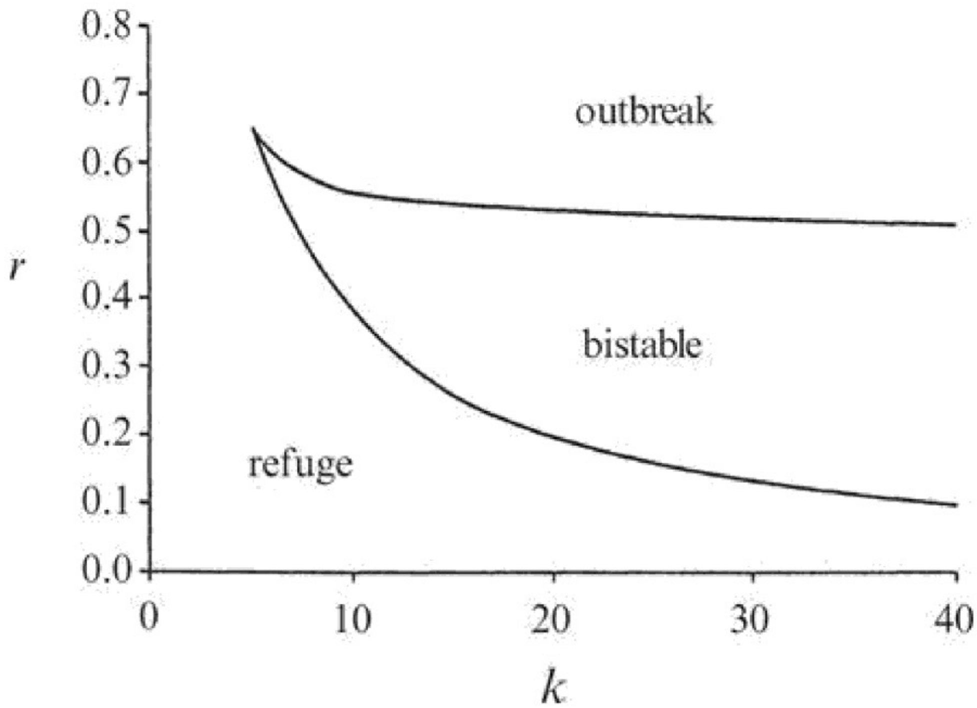

Figure 3.7.5

The different regions in Figure 3.7.5 are labeled according to the stable fixed points that exist. The refuge level a is the only stable state for low r, and the outbreak level c is the only stable state for large r. In the ***bistable*** region, both stable states exist.

The stability diagram is very similar to Figure 3.6.2. It too can be regarded as the projection of a cusp catastrophe surface, as schematically illustrated in Figure 3.7.6. You are hereby challenged to graph the surface accurately!

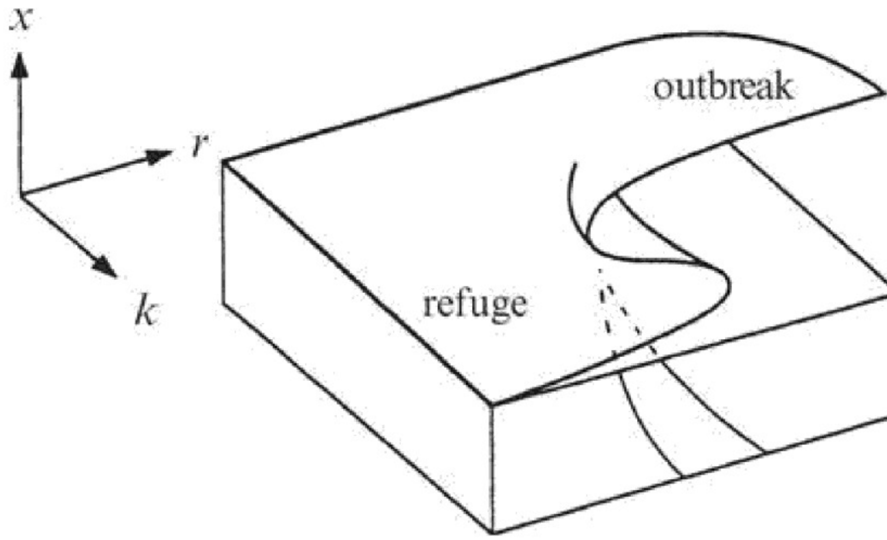

Figure 3.7.6

Comparison with Observations

Now we need to decide on biologically plausible values of the dimensionless groups $r = RA/B$ and $k = K/A$. A complication is that these parameters may drift slowly as the condition of the forest changes. According to Ludwig et al. (1978), r increases as the forest grows, while k remains fixed.

They reason as follows: let S denote the average size of the trees, interpreted as the total surface area of the branches in a stand. Then the carrying capacity K should be proportional to the available foliage, so $K = K' \ S$. Similarly, the half-saturation parameter A in the predation term should be proportional to S; predators such as birds search *units of foliage*, not acres of forest, and so the relevant quantity A must have the dimensions of budworms per unit of branch area. Hence $A = A' \ S$ and therefore

$$rRA' \ BS,k=K' \ A' . \qquad (10)$$

The experimental observations suggest that for a young forest, typically $k \approx 300$ and $r < 1/2$ so the parameters lie in the bistable region. The budworm population is kept down by the birds, which find it easy to search the small number of branches per acre. However, as the forest grows, S increases and therefore the point (k, r) drifts upward in parameter space toward the outbreak region of Figure 3.7.5. Ludwig et al. (1978) estimate that $r \approx 1$ for a fully mature forest, which lies dangerously in the outbreak region. After an outbreak occurs, the fir trees die and the forest is taken over by birch trees. But they are less efficient at using nutrients and eventually the fir trees come back—this recovery takes about 50–100 years (Murray 2002).

We conclude by mentioning some of the approximations in the model presented here. The tree dynamics have been neglected; see Ludwig et al. (1978) for a discussion of this longer time-scale behavior. We've also neglected the *spatial* distribution of budworms and their possible dispersal—see Ludwig et al. (1979) and Murray (2002) for treatments of this aspect of the problem.

EXERCISES FOR CHAPTER 3