

From the lead author of the international  
bestseller *Limits to Growth*

# Thinking in Systems

A Primer



Donella H.  
Meadows

edited by Diana Wright



# Thinking in Systems

OTHER BOOKS BY DONELLA H. MEADOWS:

*Harvesting One Hundredfold: Key Concepts and Case Studies in Environmental Education* (1989).

*The Global Citizen* (1991).

WITH DENNIS MEADOWS:

*Toward Global Equilibrium* (1973).

WITH DENNIS MEADOWS AND JØRGEN RANDERS:

*Beyond the Limits* (1992).

*Limits to Growth: The 30-Year Update* (2004).

WITH DENNIS MEADOWS, JØRGEN RANDERS, AND WILLIAM W. BEHRENS III:

*The Limits to Growth* (1972).

WITH DENNIS MEADOWS, ET AL.:

*The Dynamics of Growth in a Finite World* (1974).

WITH J. RICHARDSON AND G. BRUCKMANN:

*Groping in the Dark: The First Decade of Global Modeling* (1982).

WITH J. ROBINSON:

*The Electronic Oracle: Computer Models and Social Decisions* (1985).

# Thinking in Systems

— *A Primer* —

Donella H. Meadows

*Edited by Diana Wright,  
Sustainability Institute*



publishing for a sustainable future

LONDON • STERLING, VA

First published by Earthscan in the UK in 2009

Copyright © 2008 by Sustainability Institute.

**All rights reserved**

ISBN: 978-1-84407-726-7 (pb)

ISBN: 978-1-84407-725-0 (hb)

Typeset by Peter Holm, Sterling Hill Productions

Cover design by Dan Bramall

For a full list of publications please contact:

Earthscan

Dunstan House

14a St Cross St

London, EC1N 8XA, UK

Tel: +44 (0)20 7841 1930

Fax: +44 (0)20 7242 1474

Email: [earthinfo@earthscan.co.uk](mailto:earthinfo@earthscan.co.uk)

Web: [www.earthscan.co.uk](http://www.earthscan.co.uk)

22883 Quicksilver Drive, Sterling, VA 20166-2012, USA

Earthscan publishes in association with the International Institute for Environment and Development

A catalogue record for this book is available from the British Library

Library of Congress Cataloging-in-Publication Data has been applied for.

At Earthscan we strive to minimize our environmental impacts and carbon footprint through reducing waste, recycling and offsetting our CO<sub>2</sub> emissions, including those created through publication of this book. For more details of our environmental policy, see [www.earthscan.co.uk](http://www.earthscan.co.uk).

This book was printed in the UK by TJ International Ltd, an ISO 14001 accredited company. The paper used is FSC certified.



Part of this work has been adapted from an article originally published under the title “Whole Earth Models and Systems” in *Coevolution Quarterly* (Summer 1982). An early version of Chapter 6 appeared as “Places to Intervene in a System” in *Whole Earth Review* (Winter 1997) and later as an expanded paper published by the Sustainability Institute. Chapter 7, “Living in a World of Systems,” was originally published as “Dancing with Systems” in *Whole Earth Review* (Winter 2001).

FOR DANA

(1941–2001)

*and for all those who would learn from her*



# CONTENTS

A Note from the Author | ix  
A Note from the Editor | xi  
Introduction: The Systems Lens | 1

## **Part One: System Structures and Behavior**

ONE. The Basics | 11  
TWO. A Brief Visit to the Systems Zoo | 35

## **Part Two: Systems and Us**

THREE. Why Systems Work So Well | 75  
FOUR. Why Systems Surprise Us | 86  
FIVE. System Traps . . . and Opportunities | 111

## **Part Three: Creating Change—in Systems and in Our Philosophy**

SIX. Leverage Points—Places to Intervene in a System | 145  
SEVEN. Living in a World of Systems | 166

## **Appendix**

System Definitions: A Glossary | 187  
Summary of Systems Principles | 188  
Springing the System Traps | 191  
Places to Intervene in a System | 194  
Guidelines for Living in a World of Systems | 194  
Model Equations | 195

Notes | 204  
Bibliography of Systems Resources | 208  
Editor's Acknowledgments | 211  
About the Author | 213  
Index | 215





## A NOTE FROM THE AUTHOR

This book has been distilled out of the wisdom of thirty years of systems modeling and teaching carried out by dozens of creative people, most of them originally based at or influenced by the MIT System Dynamics group. Foremost among them is Jay Forrester, the founder of the group. My particular teachers (and students who have become my teachers) have been, in addition to Jay: Ed Roberts, Jack Pugh, Dennis Meadows, Hartmut Bossel, Barry Richmond, Peter Senge, John Sterman, and Peter Allen, but I have drawn here from the language, ideas, examples, quotes, books, and lore of a large intellectual community. I express my admiration and gratitude to all its members.

I also have drawn from thinkers in a variety of disciplines, who, as far as I know, never used a computer to simulate a system, but who are natural systems thinkers. They include Gregory Bateson, Kenneth Boulding, Herman Daly, Albert Einstein, Garrett Hardin, Václav Havel, Lewis Mumford, Gunnar Myrdal, E.F. Schumacher, a number of modern corporate executives, and many anonymous sources of ancient wisdom, from Native Americans to the Sufis of the Middle East. Strange bedfellows, but systems thinking transcends disciplines and cultures and, when it is done right, it overarches history as well.

Having spoken of transcendence, I need to acknowledge factionalism as well. Systems analysts use overarching concepts, but they have entirely human personalities, which means that they have formed many fractious schools of systems thought. I have used the language and symbols of system dynamics here, the school in which I was taught. And I present only the core of systems theory here, not the leading edge. I don't deal with the most abstract theories and am interested in analysis only when I can see how it helps solve real problems. When the abstract end of systems theory does that, which I believe it will some day, another book will have to be written.

Therefore, you should be warned that this book, like all books, is biased and incomplete. There is much, much more to systems thinking than is

presented here, for you to discover if you are interested. One of my purposes is to make you interested. Another of my purposes, the main one, is to give you a basic ability to understand and to deal with complex systems, even if your formal systems training begins and ends with this book.

—DONELLA MEADOWS, 1993

## A NOTE FROM THE EDITOR

In 1993, Donella (Dana) Meadows completed a draft of the book you now hold. The manuscript was not published at the time, but circulated informally for years. Dana died quite unexpectedly in 2001—before she completed this book. In the years since her death, it became clear that her writings have continued to be useful to a wide range of readers. Dana was a scientist and writer, and one of the best communicators in the world of systems modeling.

In 1972, Dana was lead author of *The Limits to Growth*—a best-selling and widely translated book. The cautions she and her fellow authors issued then are recognized today as the most accurate warnings of how unsustainable patterns could, if unchecked, wreak havoc across the globe. That book made headlines around the world for its observations that continual growth in population and consumption could severely damage the ecosystems and social systems that support life on earth, and that a drive for limitless economic growth could eventually disrupt many local, regional, and global systems. The findings in that book and its updates are, once again, making front-page news as we reach peak oil, face the realities of climate change, and watch a world of 6.6 billion people deal with the devastating consequences of physical growth.

In short, Dana helped usher in the notion that we have to make a major shift in the way we view the world and its systems in order to correct our course. Today, it is widely accepted that systems thinking is a critical tool in addressing the many environmental, political, social, and economic challenges we face around the world. Systems, big or small, can behave in similar ways, and understanding those ways is perhaps our best hope for making lasting change on many levels. Dana was writing this book to bring that concept to a wider audience, and that is why I and my colleagues at the Sustainability Institute decided it was time to publish her manuscript posthumously.

Will another book really help the world and help you, the reader? I think

so. Perhaps you are working in a company (or own a company) and are struggling to see how your business or organization can be part of a shift toward a better world. Or maybe you're a policy maker who is seeing others "push back" against your good ideas and good intentions. Perhaps you're a manager who has worked hard to fix some important problems in your company or community, only to see other challenges erupt in their wake. As one who advocates for changes in how a society (or a family) functions, what it values and protects, you may see years of progress easily undone in a few swift reactions. As a citizen of an increasingly global society, perhaps you are just plain frustrated with how hard it is to make a positive and lasting difference.

If so, I think that this book can help. Although one can find dozens of titles on "systems modeling" and "systems thinking," there remains a clear need for an approachable and inspiring book about systems and us—why we find them at times so baffling and how we can better learn to manage and redesign them.

At the time that Dana was writing *Thinking in Systems*, she had recently completed the twenty-year update to *Limits to Growth*, titled *Beyond the Limits*. She was a Pew Scholar in Conservation and the Environment, was serving on the Committee on Research and Exploration for the National Geographic Society, and she was teaching about systems, environment, and ethics at Dartmouth College. In all aspects of her work, she was immersed in the events of the day. She understood those events to be the outward behavior of often complex systems.

Although Dana's original manuscript has been edited and restructured, many of the examples you will find in this book are from her first draft in 1993. They may seem a bit dated to you, but in editing her work I chose to keep them because their teachings are as relevant now as they were then. The early 1990s were the time of the dissolution of the Soviet Union and great shifts in other socialist countries. The North American Free Trade Agreement was newly signed. Iraq's army invaded Kuwait and then retreated, burning oil fields on the way out. Nelson Mandela was freed from prison, and South Africa's apartheid laws were repealed. Labor leader Lech Walesa was elected president of Poland, and poet Václav Havel was elected president of Czechoslovakia. The International Panel on Climate Change issued its first assessment report, concluding that "emissions from human activities are substantially increasing the atmospheric concentra-

tions of greenhouse gases and that this will enhance the greenhouse effect and result in an additional warming of the Earth's surface." The UN held a conference in Rio de Janeiro on environment and development.

While traveling to meetings and conferences during this time, Dana read the *International Herald Tribune* and during a single week found many examples of systems in need of better management or complete redesign. She found them in the newspaper because they are all around us every day. Once you start to see the events of the day as parts of trends, and those trends as symptoms of underlying system structure, you will be able to consider new ways to manage and new ways to live in a world of complex systems. In publishing Dana's manuscript, I hope to increase the ability of readers to understand and talk about the systems around them and to act for positive change.

I hope this small approachable introduction to systems and how we think about them will be a useful tool in a world that rapidly needs to shift behaviors arising from very complex systems. This is a simple book for and about a complex world. It is a book for those who want to shape a better future.

—DIANA WRIGHT, 2008



If a factory is torn down but the rationality which produced it is left standing, then that rationality will simply produce another factory. If a revolution destroys a government, but the systematic patterns of thought that produced that government are left intact, then those patterns will repeat themselves. . . . There's so much talk about the system. And so little understanding.

—ROBERT PIRSIG, *Zen and the Art of Motorcycle Maintenance*





# Introduction: The System Lens

---

Managers are not confronted with problems that are independent of each other, but with dynamic situations that consist of complex systems of changing problems that interact with each other. I call such situations messes. . . . Managers do not solve problems, they manage messes.

—RUSSELL ACKOFF,<sup>1</sup> operations theorist

Early on in teaching about systems, I often bring out a Slinky. In case you grew up without one, a Slinky is a toy—a long, loose spring that can be made to bounce up and down, or pour back and forth from hand to hand, or walk itself downstairs.

I perch the Slinky on one upturned palm. With the fingers of the other hand, I grasp it from the top, partway down its coils. Then I pull the bottom hand away. The lower end of the Slinky drops, bounces back up again, yo-yos up and down, suspended from my fingers above.

“What made the Slinky bounce up and down like that?” I ask students.

“Your hand. You took away your hand,” they say.

So I pick up the box the Slinky came in and hold it the same way, poised on a flattened palm, held from above by the fingers of the other hand. With as much dramatic flourish as I can muster, I pull the lower hand away.

Nothing happens. The box just hangs there, of course.

“Now once again. What made the Slinky bounce up and down?”

The answer clearly lies within the Slinky itself. The hands that manipulate it suppress or release some behavior that is latent within the structure of the spring.

That is a central insight of systems theory.

Once we see the relationship between structure and behavior, we can begin to understand how systems work, what makes them produce poor results, and how to shift them into better behavior patterns. As our world

continues to change rapidly and become more complex, systems thinking will help us manage, adapt, and see the wide range of choices we have before us. It is a way of thinking that gives us the freedom to identify root causes of problems and see new opportunities.

So, what is a system? A system is a set of things—people, cells, molecules, or whatever—interconnected in such a way that they produce their own pattern of behavior over time. The system may be buffeted, constricted, triggered, or driven by outside forces. But the system’s response to these forces is characteristic of itself, and that response is seldom simple in the real world.

When it comes to Slinkies, this idea is easy enough to understand. When it comes to individuals, companies, cities, or economies, it can be heretical. The system, to a large extent, causes its own behavior! An outside event may unleash that behavior, but the same outside event applied to a different system is likely to produce a different result.

Think for a moment about the implications of that idea:

- Political leaders don’t cause recessions or economic booms. Ups and downs are inherent in the structure of the market economy.
- Competitors rarely cause a company to lose market share. They may be there to scoop up the advantage, but the losing company creates its losses at least in part through its own business policies.
- The oil-exporting nations are not solely responsible for oil-price rises. Their actions alone could not trigger global price rises and economic chaos if the oil consumption, pricing, and investment policies of the oil-importing nations had not built economies that are vulnerable to supply interruptions.
- The flu virus does not attack you; you set up the conditions for it to flourish within you.
- Drug addiction is not the failing of an individual and no one person, no matter how tough, no matter how loving, can cure a drug addict—not even the addict. It is only through understanding addiction as part of a larger set of influences and societal issues that one can begin to address it.

Something about statements like these is deeply unsettling. Something else is purest common sense. I submit that those two somethings—a resistance to and a recognition of systems principles—come from two kinds of human experience, both of which are familiar to everyone.

On the one hand, we have been taught to analyze, to use our rational ability, to trace direct paths from cause to effect, to look at things in small and understandable pieces, to solve problems by acting on or controlling the world around us. That training, the source of much personal and societal power, leads us to see presidents and competitors, OPEC and the flu and drugs as the causes of our problems.

On the other hand, long before we were educated in rational analysis, we all dealt with complex systems. We are complex systems—our own bodies are magnificent examples of integrated, interconnected, self-maintaining complexity. Every person we encounter, every organization, every animal, garden, tree, and forest is a complex system. We have built up intuitively, without analysis, often without words, a practical understanding of how these systems work, and how to work with them.

Modern systems theory, bound up with computers and equations, hides the fact that it traffics in truths known at some level by everyone. It is often possible, therefore, to make a direct translation from systems jargon to traditional wisdom.

Because of feedback delays within complex systems, by the time a problem becomes apparent it may be unnecessarily difficult to solve.

— *A stitch in time saves nine.*

According to the competitive exclusion principle, if a reinforcing feedback loop rewards the winner of a competition with the means to win further competitions, the result will be the elimination of all but a few competitors.

— *For he that hath, to him shall be given; and he that hath not, from him shall be taken even that which he hath (Mark 4:25)*

or

— *The rich get richer and the poor get poorer.*

A diverse system with multiple pathways and redundancies is

more stable and less vulnerable to external shock than a uniform system with little diversity.

— *Don't put all your eggs in one basket.*

Ever since the Industrial Revolution, Western society has benefited from science, logic, and reductionism over intuition and holism. Psychologically and politically we would much rather assume that the cause of a problem is “out there,” rather than “in here.” It’s almost irresistible to blame something or someone else, to shift responsibility away from ourselves, and to look for the control knob, the product, the pill, the technical fix that will make a problem go away.

Serious problems have been solved by focusing on external agents—preventing smallpox, increasing food production, moving large weights and many people rapidly over long distances. Because they are embedded in larger systems, however, some of our “solutions” have created further problems. And some problems, those most rooted in the internal structure of complex systems, the real messes, have refused to go away.

Hunger, poverty, environmental degradation, economic instability, unemployment, chronic disease, drug addiction, and war, for example, persist in spite of the analytical ability and technical brilliance that have been directed toward eradicating them. No one deliberately creates those problems, no one wants them to persist, but they persist nonetheless. That is because they are intrinsically systems problems—undesirable behaviors characteristic of the system structures that produce them. They will yield only as we reclaim our intuition, stop casting blame, see the system as the source of its own problems, and find the courage and wisdom to *restructure* it.

Obvious. Yet subversive. An old way of seeing. Yet somehow new. Comforting, in that the solutions are in our hands. Disturbing, because we must *do things*, or at least *see things* and *think about things*, in a different way.

This book is about that different way of seeing and thinking. It is intended for people who may be wary of the word “systems” and the field of systems analysis, even though they may have been doing systems thinking all their lives. I have kept the discussion nontechnical because I want to show what a long way you can go toward understanding systems without turning to mathematics or computers.

I have made liberal use of diagrams and time graphs in this book

because there is a problem in discussing systems only with words. Words and sentences must, by necessity, come only one at a time in linear, logical order. Systems happen all at once. They are connected not just in one direction, but in many directions simultaneously. To discuss them properly, it is necessary somehow to use a language that shares some of the same properties as the phenomena under discussion.

Pictures work for this language better than words, because you can see all the parts of a picture at once. I will build up systems pictures gradually, starting with very simple ones. I think you'll find that you can understand this graphical language easily.

I start with the basics: the definition of a system and a dissection of its parts (in a reductionist, unholistic way). Then I put the parts back together to show how they interconnect to make the basic operating unit of a system: the feedback loop.

Next I will introduce you to a systems zoo—a collection of some common and interesting types of systems. You'll see how a few of these creatures behave and why and where they can be found. You'll recognize them; they're all around you and even within you.

With a few of the zoo "animals"—a set of specific examples—as a foundation, I'll step back and talk about how and why systems work so beautifully and the reasons why they so often surprise and confound us. I'll talk about why everyone or everything in a system can act dutifully and rationally, yet all these well-meaning actions too often add up to a perfectly terrible result. And why things so often happen much faster or slower than everyone thinks they will. And why you can be doing something that has always worked and suddenly discover, to your great disappointment, that your action no longer works. And why a system might suddenly, and without warning, jump into a kind of behavior you've never seen before.

That discussion will lead to us to look at the common problems that the systems-thinking community has stumbled upon over and over again through working in corporations and governments, economies and ecosystems, physiology and psychology. "There's another case of the tragedy of the commons," we find ourselves saying as we look at an allocation system for sharing water resource among communities or financial resources among schools. Or we identify "eroding goals" as we study the business rules and incentives that help or hinder the development of new technologies. Or we see "policy resistance" as we examine decision-making power and the nature of relationships in a

family, a community, or a nation. Or we witness “addiction”—which can be caused by many more agents than caffeine, alcohol, nicotine, and narcotics.

Systems thinkers call these common structures that produce characteristic behaviors “archetypes.” When I first planned this book, I called them “system traps.” Then I added the words “and opportunities,” because these archetypes, which are responsible for some of the most intransigent and potentially dangerous problems, also can be transformed, with a little systems understanding, to produce much more desirable behaviors.

From this understanding I move into what you and I can do about restructuring the systems we live within. We can learn how to look for leverage points for change.

I conclude with the largest lessons of all, the ones derived from the wisdom shared by most systems thinkers I know. For those who want to explore systems thinking further, the Appendix provides ways to dig deeper into the subject with a glossary, a bibliography of systems thinking resources, a summary list of systems principles, and equations for the models described in Part One.

When our small research group moved from MIT to Dartmouth College years ago, one of the Dartmouth engineering professors watched us in seminars for a while, and then dropped by our offices. “You people are different,” he said. “You ask different kinds of questions. You see things I don’t see. Somehow you come at the world in a different way. How? Why?”

That’s what I hope to get across throughout this book, but especially in its conclusion. I don’t think the systems way of seeing is better than the reductionist way of thinking. I think it’s complementary, and therefore revealing. You can see some things through the lens of the human eye, other things through the lens of a microscope, others through the lens of a telescope, and still others through the lens of systems theory. Everything seen through each kind of lens is actually there. Each way of seeing allows our knowledge of the wondrous world in which we live to become a little more complete.

At a time when the world is more messy, more crowded, more interconnected, more interdependent, and more rapidly changing than ever before, the more ways of seeing, the better. The systems-thinking lens allows us to reclaim our intuition about whole systems and

- hone our abilities to understand parts,

- see interconnections,
- ask “what-if” questions about possible future behaviors, and
- be creative and courageous about system redesign.

Then we can use our insights to make a difference in ourselves and our world.

---

### **INTERLUDE • *The Blind Men and the Matter of the Elephant***

---

Beyond Ghor, there was a city. All its inhabitants were blind. A king with his entourage arrived nearby; he brought his army and camped in the desert. He had a mighty elephant, which he used to increase the people’s awe.

The populace became anxious to see the elephant, and some sightless from among this blind community ran like fools to find it.

As they did not even know the form or shape of the elephant, they groped sightlessly, gathering information by touching some part of it.

Each thought that he knew something, because he could feel a part. . . .

The man whose hand had reached an ear . . . said: “It is a large, rough thing, wide and broad, like a rug.”

And the one who had felt the trunk said: “I have the real facts about it. It is like a straight and hollow pipe, awful and destructive.”

The one who had felt its feet and legs said: “It is mighty and firm, like a pillar.”

Each had felt one part out of many. Each had perceived it wrongly. . . .<sup>2</sup>

This ancient Sufi story was told to teach a simple lesson but one that we often ignore: The behavior of a system cannot be known just by knowing the elements of which the system is made.

---





**PART ONE**

System Structure and Behavior

---



# The Basics

---

I have yet to see any problem, however complicated, which, when looked at in the right way, did not become still more complicated.

—POUL ANDERSON<sup>1</sup>

## More Than the Sum of Its Parts

A system isn't just any old collection of things. A **system**<sup>\*</sup> is an interconnected set of elements that is coherently organized in a way that achieves something. If you look at that definition closely for a minute, you can see that a system must consist of three kinds of things: *elements*, *interconnections*, and a *function* or *purpose*.

For example, the elements of your digestive system include teeth, enzymes, stomach, and intestines. They are interrelated through the physical flow of food, and through an elegant set of regulating chemical signals. The function of this system is to break down food into its basic nutrients and to transfer those nutrients into the bloodstream (another system), while discarding unusable wastes.

A football team is a system with elements such as players, coach, field, and ball. Its interconnections are the rules of the game, the coach's strategy, the players' communications, and the laws of physics that govern the motions of ball and players. The purpose of the team is to win games, or have fun, or get exercise, or make millions of dollars, or all of the above.

A school is a system. So is a city, and a factory, and a corporation, and a national economy. An animal is a system. A tree is a system, and a forest is a larger system that encompasses subsystems of trees and animals. The earth

---

<sup>1</sup> Definitions of words in bold face can be found in the Glossary.

is a system. So is the solar system; so is a galaxy. Systems can be embedded in systems, which are embedded in yet other systems.

Is there anything that is not a system? Yes—a conglomeration without any particular interconnections or function. Sand scattered on a road by happenstance is not, itself, a system. You can add sand or take away sand and you still have just sand on the road. Arbitrarily add or take away football players, or pieces of your digestive system, and you quickly no longer have the same system.

When a living creature dies, it loses its “system-ness.” The multiple interrelations that held it together no longer function, and it dissipates, although its material remains part of a larger food-web system. Some people say that an old city neighborhood where people know each other and communicate regularly is a social system, and that a new apartment block full of strangers is not—not until new relationships arise and a system forms.

---

**A system is more than the sum of its parts.** It may exhibit adaptive, dynamic, goal-seeking, self-preserving, and sometimes evolutionary behavior.

---

You can see from these examples that there is an integrity or wholeness about a system and an active set of mechanisms to maintain that integrity. Systems can change, adapt, respond to events, seek goals, mend injuries, and attend to their own survival in lifelike ways, although they may contain or consist of nonliving things. Systems can be self-organizing, and often are self-repairing over at least some range of disruptions. They are resilient, and many of them are evolutionary. Out of one system other completely new, never-before-imagined systems can arise.

### Look Beyond the Players to the Rules of the Game

You think that because you understand “one” that you must therefore understand “two” because one and one make two. But you forget that you must also understand “and.”

—Sufi teaching story

The elements of a system are often the easiest parts to notice, because many of them are visible, tangible things. The elements that make up a tree are roots, trunk, branches, and leaves. If you look more closely, you

**THINK ABOUT THIS**

How to know whether you are looking at a system or just a bunch of stuff:

- A) Can you identify parts? . . . and
- B) Do the parts affect each other? . . . and
- C) Do the parts together produce an effect that is different from the effect of each part on its own? . . . and perhaps
- D) Does the effect, the behavior over time, persist in a variety of circumstances?

see specialized cells: vessels carrying fluids up and down, chloroplasts, and so on. The system called a university is made up of buildings, students, professors, administrators, libraries, books, computers—and I could go on and say what all those things are made up of. Elements do not have to be physical things. Intangibles are also elements of a system. In a university, school pride and academic prowess are two intangibles that can be very important elements of the system. Once you start listing the elements of a system, there is almost no end to the process. You can divide elements into sub-elements and then sub-sub-elements. Pretty soon you lose sight of the system. As the saying goes, you can't see the forest for the trees.

Before going too far in that direction, it's a good idea to stop dissecting out elements and to start looking for the *interconnections*, the relationships that hold the elements together.

The interconnections in the tree system are the physical flows and chemical reactions that govern the tree's metabolic processes—the signals that allow one part to respond to what is happening in another part. For example, as the leaves lose water on a sunny day, a drop in pressure in the water-carrying vessels allows the roots to take in more water. Conversely, if the roots experience dry soil, the loss of water pressure signals the leaves to close their pores, so as not to lose even more precious water.

As the days get shorter in the temperate zones, a deciduous tree puts forth chemical messages that cause nutrients to migrate out of the leaves into the trunk and roots and that weaken the stems, allowing the leaves to

fall. There even seem to be messages that cause some trees to make repellent chemicals or harder cell walls if just one part of the plant is attacked by insects. No one understands all the relationships that allow a tree to do what it does. That lack of knowledge is not surprising. It's easier to learn about a system's elements than about its interconnections.

In the university system, interconnections include the standards for admission, the requirements for degrees, the examinations and grades, the budgets and money flows, the gossip, and most important, the communication of knowledge that is, presumably, the purpose of the whole system.

---

**Many of the interconnections in systems operate through the flow of information.** Information holds systems together and plays a great role in determining how they operate.

---

Some interconnections in systems are actual physical flows, such as the water in the tree's trunk or the students progressing through a university. Many interconnections are flows of information—signals that go to decision points or action points within a system. These kinds of interconnections are often harder to see, but the system reveals them to those who look. Students may use informal information about the probability of getting a good grade to decide what courses to take. A consumer decides what to buy using information about his or her income, savings, credit rating, stock of goods at home, prices, and availability of goods for purchase. Governments need information about kinds and quantities of water pollution before they can create sensible regulations to reduce that pollution. (Note that information about the existence of a problem may be necessary but not sufficient to trigger action—information about resources, incentives, and consequences is necessary too.)

If information-based relationships are hard to see, *functions* or *purposes* are even harder. A system's function or purpose is not necessarily spoken, written, or expressed explicitly, except through the operation of the system. The best way to deduce the system's purpose is to watch for a while to see how the system behaves.

If a frog turns right and catches a fly, and then turns left and catches a fly, and then turns around backward and catches a fly, the purpose of the frog has to do not with turning left or right or backward but with catching flies. If a government proclaims its interest in protecting the environment but allocates little money or effort toward that goal, environmental protection is not, in fact, the government's purpose. Purposes are deduced from behavior, not from rhetoric or stated goals.

### A NOTE ON LANGUAGE

The word *function* is generally used for a nonhuman system, the word *purpose* for a human one, but the distinction is not absolute, since so many systems have both human and nonhuman elements.

The function of a thermostat-furnace system is to keep a building at a given temperature. One function of a plant is to bear seeds and create more plants. One purpose of a national economy is, judging from its behavior, to keep growing larger. An important function of almost every system is to ensure its own perpetuation.

System purposes need not be human purposes and are not necessarily those intended by any single actor within the system. In fact, one of the most frustrating aspects of systems is that the purposes of subunits may add up to an overall behavior that no one wants. No one intends to produce a society with rampant drug addiction and crime, but consider the combined purposes and consequent actions of the actors involved:

- desperate people who want quick relief from psychological pain
- farmers, dealers, and bankers who want to earn money
- pushers who are less bound by civil law than are the police who oppose them
- governments that make harmful substances illegal and use police power to interdict them
- wealthy people living in close proximity to poor people
- nonaddicts who are more interested in protecting themselves than in encouraging recovery of addicts

Altogether, these make up a system from which it is extremely difficult to eradicate drug addiction and crime.

Systems can be nested within systems. Therefore, there can be purposes within purposes. The purpose of a university is to discover and preserve knowledge and pass it on to new generations. Within the university, the purpose of a student may be to get good grades, the purpose of a professor



may be to get tenure, the purpose of an administrator may be to balance the budget. Any of those sub-purposes could come into conflict with the overall purpose—the student could cheat, the professor could ignore the students in order to publish papers, the administrator could balance the budget by firing professors. Keeping sub-purposes and overall system purposes in harmony is an essential function of successful systems. I'll get back to this point later when we come to hierarchies.

You can understand the relative importance of a system's elements, interconnections, and purposes by imagining them changed one by one. Changing elements usually has the least effect on the system. If you change all the players on a football team, it is still recognizably a football team. (It may play much better or much worse—particular elements in a system can indeed be important.) A tree changes its cells constantly, its leaves

---

**The least obvious part of the system, its function or purpose, is often the most crucial determinant of the system's behavior.**

---

every year or so, but it is still essentially the same tree. Your body replaces most of its cells every few weeks, but it goes on being your body. The university has a constant flow of students and a slower flow of professors and administrators, but it is still a university. In fact it is still the same university, distinct in subtle ways from others, just as

General Motors and the U.S. Congress somehow maintain their identities even though all their members change. A system generally goes on being itself, changing only slowly if at all, even with complete substitutions of its elements—as long as its interconnections and purposes remain intact.

If the interconnections change, the system may be greatly altered. It may even become unrecognizable, even though the same players are on the team. Change the rules from those of football to those of basketball, and you've got, as they say, a whole new ball game. If you change the interconnections in the tree—say that instead of taking in carbon dioxide and emitting oxygen, it does the reverse—it would no longer be a tree. (It would be an animal.) If in a university the students graded the professors, or if arguments were won by force instead of reason, the place would need a different name. It might be an interesting organization, but it would not be a university. Changing interconnections in a system can change it dramatically.

Changes in function or purpose also can be drastic. What if you keep the players and the rules but change the purpose—from winning to losing, for example? What if the function of a tree were not to survive and repro-

duce but to capture all the nutrients in the soil and grow to unlimited size? People have imagined many purposes for a university besides disseminating knowledge—making money, indoctrinating people, winning football games. A change in purpose changes a system profoundly, even if every element and interconnection remains the same.

To ask whether elements, interconnections, or purposes are most important in a system is to ask an unsystemic question. All are essential. All interact. All have their roles. But the least obvious part of the system, its function or purpose, is often the most crucial determinant of the system's behavior. Interconnections are also critically important. Changing relationships usually changes system behavior. The elements, the parts of systems we are most likely to notice, are often (not always) least important in defining the unique characteristics of the system—*unless changing an element also results in changing relationships or purpose.*

Changing just one leader at the top—from a Brezhnev to a Gorbachev, or from a Carter to a Reagan—may or may not turn an entire nation in a new direction, though its land, factories, and hundreds of millions of people remain exactly the same. A leader can make that land and those factories and people play a different game with new rules, or can direct the play toward a new purpose.

And conversely, because land, factories, and people are long-lived, slowly changing, physical elements of the system, there is a limit to the rate at which any leader can turn the direction of a nation.

### **Bathtubs 101—Understanding System Behavior over Time**

Information contained in nature . . . allows us a partial reconstruction of the past. . . . The development of the meanders in a river, the increasing complexity of the earth's crust . . . are information-storing devices in the same manner that genetic systems are. . . . Storing information means increasing the complexity of the mechanism.

—Ramon Margalef<sup>2</sup>

A **stock** is the foundation of any system. Stocks are the elements of the system that you can see, feel, count, or measure at any given time. A system stock is just what it sounds like: a store, a quantity, an accumulation of

material or information that has built up over time. It may be the water in a bathtub, a population, the books in a bookstore, the wood in a tree, the money in a bank, your own self-confidence. A stock does not have to be

**A stock is the memory of the history of changing flows within the system.**


physical. Your reserve of good will toward others or your supply of hope that the world can be better are both stocks.

Stocks change over time through the actions of a **flow**. Flows are filling and draining, births and deaths, purchases and sales, growth and decay, deposits and withdrawals, successes and failures. A stock, then, is the present memory of the history of changing flows within the system.




**Figure 1. How to read stock-and-flow diagrams.** In this book, stocks are shown as boxes, and flows as arrow-headed “pipes” leading into or out of the stocks. The small T on each flow signifies a “faucet;” it can be turned higher or lower, on or off. The “clouds” stand for wherever the flows come from and go to—the sources and sinks that are being ignored for the purposes of the present discussion.

For example, an underground mineral deposit is a stock, out of which comes a flow of ore through mining. The inflow of ore into a mineral deposit is minute in any time period less than eons. So I have chosen to draw (Figure 2) a simplified picture of the system without any inflow. *All* system diagrams and descriptions are simplified versions of the real world.




**Figure 2.** A stock of minerals depleted by mining.

Water in a reservoir behind a dam is a stock, into which flow rain and river water, and out of which flows evaporation from the reservoir’s surface as well as the water discharged through the dam.



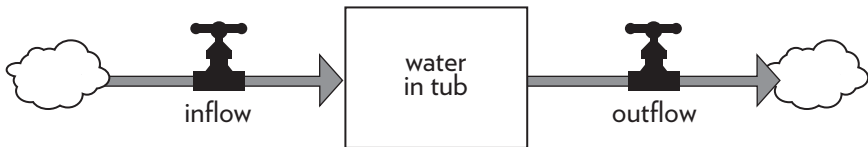
**Figure 3.** A stock of water in a reservoir with multiple inflows and outflows.

The volume of wood in the living trees in a forest is a stock. Its inflow is the growth of the trees. Its outflows are the natural deaths of trees and the harvest by loggers. The logging harvest flows into another stock, perhaps an inventory of lumber at a mill. Wood flows out of the inventory stock as lumber sold to customers.



**Figure 4.** A stock of lumber linked to a stock of trees in a forest.


If you understand the **dynamics** of stocks and flows—their behavior over time—you understand a good deal about the behavior of complex systems. And if you have had much experience with a bathtub, you understand the dynamics of stocks and flows.



**Figure 5.** The structure of a bathtub system—one stock with one inflow and one outflow.

Imagine a bathtub filled with water, with its drain plugged up and its faucets turned off—an unchanging, undynamic, boring system. Now

mentally pull the plug. The water runs out, of course. The level of water in the tub goes down until the tub is empty.



**Figure 6.** Water level in a tub when the plug is pulled.

### A NOTE ON READING GRAPHS OF BEHAVIOR OVER TIME


Systems thinkers use graphs of system behavior to understand trends over time, rather than focusing attention on individual events. We also use behavior-over-time graphs to learn whether the system is approaching a goal or a limit, and if so, how quickly.

The variable on the graph may be a stock or a flow. The pattern—the shape of the variable line—is important, as are the points at which that line changes shape or direction. The precise numbers on the axes are often less important.

The horizontal axis of time allows you to ask questions about what came before, and what might happen next. It can help you focus on the time horizon appropriate to the question or problem you are investigating.

Now imagine starting again with a full tub, and again open the drain, but this time, when the tub is about half empty, turn on the inflow faucet so the rate of water flowing in is just equal to that flowing out. What happens?

The amount of water in the tub stays constant at whatever level it had reached when the inflow became equal to the outflow. It is in a state of **dynamic equilibrium**—its level does not change, although water is continuously flowing through it.



**Figure 7.** Constant outflow, inflow turned on after 5 minutes, and the resulting changes in the stock of water in the tub.

Imagine turning the inflow on somewhat harder while keeping the outflow constant. The level of water in the tub slowly rises. If you then turn the inflow

faucet down again to match the outflow exactly, the water in the tub will stop rising. Turn it down some more, and the water level will fall slowly.

This model of a bathtub is a very simple system with just one stock, one inflow, and one outflow. Over the time period of interest (minutes), I have assumed that evaporation from the tub is insignificant, so I have not included that outflow. All models, whether mental models or mathematical models, are simplifications of the real world. You know all the dynamic possibilities of this bathtub. From it you can deduce several important principles that extend to more complicated systems:

- As long as the sum of all inflows exceeds the sum of all outflows, the level of the stock will rise.
- As long as the sum of all outflows exceeds the sum of all inflows, the level of the stock will fall.
- If the sum of all outflows equals the sum of all inflows, the stock level will not change; it will be held in dynamic equilibrium at whatever level it happened to be when the two sets of flows became equal.

The human mind seems to focus more easily on stocks than on flows. On top of that, when we do focus on flows, we tend to focus on *inflows* more easily than on outflows. Therefore, we sometimes miss seeing that we can

---

**A stock can be increased by decreasing its outflow rate as well as by increasing its inflow rate.** There's more than one way to fill a bathtub!

---

fill a bathtub not only by increasing the inflow rate, but also by decreasing the outflow rate. Everyone understands that you can prolong the life of an oil-based economy by discovering new oil deposits. It seems to be harder to understand that the same result can be achieved by burning less oil. A breakthrough in energy efficiency is equivalent, in its effect on the stock of available oil, to the discovery

of a new oil field—although different people profit from it.

Similarly, a company can build up a larger workforce by more hiring, or it can do the same thing by reducing the rates of quitting and firing. These two strategies may have very different costs. The wealth of a nation can be boosted by investment to build up a larger stock of factories and machines. It also can be boosted, often more cheaply, by decreasing the rate at which factories and machines wear out, break down, or are discarded.

You can adjust the drain or faucet of a bathtub—the flows—abruptly, but it is much more difficult to change the level of water—the stock—quickly. Water can't run out the drain instantly, even if you open the drain all the way. The tub can't fill up immediately, even with the inflow faucet on full blast. *A stock takes time to change, because flows take time to flow.* That's a vital point, a key to understanding why systems behave as they do. Stocks usually change slowly. They can act as delays, lags, buffers, ballast, and sources of momentum in a system. Stocks, especially large ones, respond to change, even sudden change, only by gradual filling or emptying.

People often underestimate the inherent momentum of a stock. It takes a long time for populations to grow or stop growing, for wood to accumulate in a forest, for a reservoir to fill up, for a mine to be depleted. An economy cannot

build up a large stock of functioning factories and highways and electric plants overnight, even if a lot of money is available. Once an economy has a lot of oil-burning furnaces and automobile engines, it cannot change quickly to furnaces and engines that burn a different fuel, even if the price of oil suddenly changes. It has taken decades to accumulate the stratospheric pollutants that destroy the earth's ozone layer; it will take decades for those pollutants to be removed.

Changes in stocks set the pace of the dynamics of systems. Industrialization cannot proceed faster than the rate at which factories and machines can be constructed and the rate at which human beings can be educated to run and maintain them. Forests can't grow overnight. Once contaminants have accumulated in groundwater, they can be washed out only at the rate of groundwater turnover, which may take decades or even centuries.

The time lags that come from slowly changing stocks can cause problems in systems, but they also can be sources of stability. Soil that has accumulated over centuries rarely erodes all at once. A population that has learned many skills doesn't forget them immediately. You can pump groundwater faster than the rate it recharges for a long time before the aquifer is drawn down far enough to be damaged. The time lags imposed by stocks allow room to maneuver, to experiment, and to revise policies that aren't working.

If you have a sense of the rates of change of stocks, you don't expect things to happen faster than they can happen. You don't give up too soon.

---

Stocks generally change slowly, even when the flows into or out of them change suddenly. Therefore, **stocks act as delays or buffers or shock absorbers in systems.**

---



You can use the opportunities presented by a system's momentum to guide it toward a good outcome—much as a judo expert uses the momentum of an opponent to achieve his or her own goals.

There is one more important principle about the role of stocks in systems, a principle that will lead us directly to the concept of feedback. The pres-

---

**Stocks allow inflows and outflows to be decoupled and to be independent and temporarily out of balance with each other.**

---

ence of stocks allows inflows and outflows to be independent of each other and temporarily out of balance with each other.

It would be hard to run an oil company if gasoline had to be produced at the refinery at exactly the rate the cars were burning it. It isn't feasible to harvest a forest at the precise rate at which the trees are growing. Gasoline in storage tanks and wood in the forest are both stocks that permit life to proceed with some certainty, continuity, and predictability, even though flows vary in the short term.

Human beings have invented hundreds of stock-maintaining mechanisms to make inflows and outflows independent and stable. Reservoirs enable residents and farmers downriver to live without constantly adjusting their lives and work to a river's varying flow, especially its droughts and floods. Banks enable you temporarily to earn money at a rate different from how you spend. Inventories of products along a chain from distributors to wholesalers to retailers allow production to proceed smoothly although customer demand varies, and allow customer demand to be filled even though production rates vary.

Most individual and institutional decisions are designed to regulate the levels in stocks. If inventories rise too high, then prices are cut or advertising budgets are increased, so that sales will go up and inventories will fall. If the stock of food in your kitchen gets low, you go to the store. As the stock of growing grain rises or fails to rise in the fields, farmers decide whether to apply water or pesticide, grain companies decide how many barges to book for the harvest, speculators bid on future values of the harvest, cattle growers build up or cut down their herds. Water levels in reservoirs cause all sorts of corrective actions if they rise too high or fall too low. The same can be said for the stock of money in your wallet, the oil reserves owned by an oil company, the pile of woodchips feeding a paper mill, and the concentration of pollutants in a lake.

People monitor stocks constantly and make decisions and take actions

designed to raise or lower stocks or to keep them within acceptable ranges. Those decisions add up to the ebbs and flows, successes and problems, of all sorts of systems. Systems thinkers see the world as a collection of stocks along with the mechanisms for regulating the levels in the stocks by manipulating flows.

That means system thinkers see the world as a collection of “feedback processes.”

### How the System Runs Itself—Feedback

Systems of information-feedback control are fundamental to all life and human endeavor, from the slow pace of biological evolution to the launching of the latest space satellite. . . . Everything we do as individuals, as an industry, or as a society is done in the context of an information-feedback system.

—Jay W. Forrester<sup>3</sup>


When a stock grows by leaps and bounds or declines swiftly or is held within a certain range no matter what else is going on around it, it is likely that there is a control mechanism at work. In other words, if you see a behavior that persists over time, there is likely a mechanism creating that consistent behavior. That mechanism operates through a **feedback loop**. It is the consistent behavior pattern over a long period of time that is the first hint of the existence of a feedback loop.

A feedback loop is formed when changes in a stock affect the flows into or out of that same stock. A feedback loop can be quite simple and direct. Think of an interest-bearing savings account in a bank. The total amount of money in the account (the stock) affects how much money comes into the account as interest. That is because the bank has a rule that the account earns a certain percent interest each year. The total dollars of interest paid into the account each year (the flow in) is not a fixed amount, but varies with the size of the total in the account.

You experience another fairly direct kind of feedback loop when you get your bank statement for your checking account each month. As your level of available cash in the checking account (a stock) goes down, you may decide to work more hours and earn more money. The money entering

your bank account is a flow that you can adjust in order to increase your stock of cash to a more desirable level. If your bank account then grows very large, you may feel free to work less (decreasing the inflow). This kind of feedback loop is keeping your level of cash available within a range that is acceptable to you. You can see that adjusting your earnings is not the only feedback loop that works on your stock of cash. You also may be able to adjust the outflow of money from your account, for example. You can imagine an outflow-adjusting feedback loop for spending.

Feedback loops can cause stocks to maintain their level within a range or grow or decline. In any case, the flows into or out of the stock are adjusted because of changes in the size of the stock itself. Whoever or whatever is monitoring the stock's level begins a corrective process, adjusting rates of inflow or outflow (or both) and so changing the stock's level. The stock level feeds back through a chain of signals and actions to control itself.



**Figure 8.** How to read a stock-and-flow diagram with feedback loops. Each diagram distinguishes the stock, the flow that changes the stock, and the information link (shown as a thin, curved arrow) that directs the action. It emphasizes that action or change always proceeds through adjusting flows.

Not all systems have feedback loops. Some systems are relatively simple open-ended chains of stocks and flows. The chain may be affected by outside factors, but the levels of the chain's stocks don't affect its flows. However, those systems that contain feedback loops are common and may be quite elegant or rather surprising, as we shall see.

---


**A feedback loop is a closed chain of causal connections from a stock, through a set of decisions or rules or physical laws or actions that are dependent on the level of the stock, and back again through a flow to change the stock.**

---

### Stabilizing Loops—Balancing Feedback

One common kind of feedback loop stabilizes the stock level, as in the checking-account example. The stock level may not remain completely fixed, but it does stay within an acceptable range. What follows are some more stabilizing feedback loops that may be familiar to you. These examples start to detail some of the steps within a feedback loop.

If you're a coffee drinker, when you feel your energy level run low, you may grab a cup of hot black stuff to perk you up again. You, as the coffee drinker, hold in your mind a desired stock level (energy for work). The purpose of this caffeine-delivery system is to keep your actual stock level near or at your desired level. (You may have other purposes for drinking coffee as well: enjoying the flavor or engaging in a social activity.) It is the



**Figure 9.** Energy level of a coffee drinker.

gap, the discrepancy, between your actual and desired levels of energy for work that drives your decisions to adjust your daily caffeine intake.


Notice that the labels in Figure 9, like all the diagram labels in this book, are direction-free. The label says “stored energy in body” not “*low* energy level,” “coffee intake” not “*more* coffee.” That’s because feedback loops often can operate in two directions. In this case, the feedback loop can correct an oversupply as well as an undersupply. If you drink too much coffee and find yourself bouncing around with extra energy, you’ll lay off the caffeine for a while. High energy creates a discrepancy that says “too much,” which then causes you to reduce your coffee intake until your energy level settles down. The diagram is intended to show that the loop works to drive the stock of energy in either direction.

I could have shown the inflow of energy coming from a cloud, but instead I made the system diagram slightly more complicated. *Remember—all system diagrams are simplifications of the real world.* We each choose how much complexity to look at. In this example, I drew another stock—the stored energy in the body that can be activated by the caffeine. I did that to indicate that there is more to the system than one simple loop. As every coffee drinker knows, caffeine is only a short-term stimulant. It lets you run your motor faster, but it doesn’t refill your personal fuel tank. Eventually the caffeine high wears off, leaving the body more energy-deficient than it was before. That drop could reactivate the feedback loop and generate another trip to the coffee pot. (See the discussion of addiction later in this book.) Or it could activate some longer-term and healthier feedback responses: Eat some food, take a walk, get some sleep.

This kind of stabilizing, goal-seeking, regulating loop is called a **balancing feedback loop**, so I put a B inside the loop in the diagram. Balancing feedback loops are *goal-seeking* or *stability-seeking*. Each tries to keep a stock at a given value or within a range of values. A balancing feedback loop opposes whatever direction of change is imposed on the system. If you push a stock too far up, a balancing loop will try to pull it back down. If you shove it too far down, a balancing loop will try to bring it back up.


Here’s another balancing feedback loop that involves coffee, but one that works through physical law rather than human decision. A hot cup of coffee will gradually cool down to room temperature. Its rate of cooling depends on the difference between the temperature of the coffee and the temperature of the room. The greater the difference, the faster the coffee

will cool. The loop works the other way too—if you make iced coffee on a hot day, it will warm up until it has the same temperature as the room. The function of this system is to bring the discrepancy between coffee’s temperature and room’s temperature to zero, no matter what the direction of the discrepancy.



**Figure 10.** A cup of coffee cooling (*left*) or warming (*right*).

Starting with coffee at different temperatures, from just below boiling to just above freezing, the graphs in Figure 11 show what will happen to the temperature over time (if you don’t drink the coffee). You can see here the “homing” behavior of a balancing feedback loop. Whatever the initial value of the system stock (coffee temperature in this case), whether it is above or below the “goal” (room temperature), the feedback loop brings it toward



**Figure 11.** Coffee temperature as it approaches a room temperature of 18°C.

the goal. The change is faster at first, and then slower, as the discrepancy between the stock and the goal decreases.

---

**Balancing feedback loops are equilibrating or goal-seeking structures in systems and are both sources of stability and sources of resistance to change.**

---

This behavior pattern—gradual approach to a system-defined goal—also can be seen when a radioactive element decays, when a missile finds its target, when an asset depreciates, when a reservoir is brought up or down to its desired level, when your body adjusts its blood-sugar concentration, when you pull your car to a stop at a stoplight. You can think of many more

examples. The world is full of goal-seeking feedback loops.

The presence of a feedback mechanism doesn't necessarily mean that the mechanism works *well*. The feedback mechanism may not be strong enough to bring the stock to the desired level. Feedbacks—the interconnections, the information part of the system—can fail for many reasons. Information can arrive too late or at the wrong place. It can be unclear or incomplete or hard to interpret. The action it triggers may be too weak or delayed or resource-constrained or simply ineffective. The goal of the feedback loop may never be reached by the actual stock. But in the simple example of a cup of coffee, the drink eventually will reach room temperature.

### Runaway Loops—Reinforcing Feedback

I'd need rest to refresh my brain, and to get rest it's necessary to travel, and to travel one must have money, and in order to get money you have to work. . . . I am in a vicious circle . . . from which it is impossible to escape.

—Honoré Balzac,<sup>4</sup> 19th century novelist and playwright

Here we meet a very important feature. It would seem as if this were circular reasoning; profits fell because investment fell, and investment fell because profits fell.

—Jan Tinbergen,<sup>5</sup> economist


The second kind of feedback loop is amplifying, reinforcing, self-multiplying, snowballing—a vicious or virtuous circle that can cause healthy growth

or runaway destruction. It is called a **reinforcing feedback loop**, and will be noted with an R in the diagrams. It generates more input to a stock the more that is already there (and less input the less that is already there). A reinforcing feedback loop enhances whatever direction of change is imposed on it.

For example:

- When we were kids, the more my brother pushed me, the more I pushed him back, so the more he pushed me back, so the more I pushed him back.
- The more prices go up, the more wages have to go up if people are to maintain their standards of living. The more wages go up, the more prices have to go up to maintain profits. This means that wages have to go up again, so prices go up again.
- The more rabbits there are, the more rabbit parents there are to make baby rabbits. The more baby rabbits there are, the more grow up to become rabbit parents, to have even more baby rabbits.
- The more soil is eroded from the land, the less plants are able to grow, so the fewer roots there are to hold the soil, so the more soil is eroded, so less plants can grow.
- The more I practice piano, the more pleasure I get from the sound, and so the more I play the piano, which gives me more practice.

Reinforcing loops are found wherever a system element has the ability to reproduce itself or to grow as a constant fraction of itself. Those elements include populations and economies. Remember the example of




**Figure 12.** Interest-bearing bank account.



the interest-bearing bank account? The more money you have in the bank, the more interest you earn, which is added to the money already in the bank, where it earns even more interest.

Figure 13 shows how this reinforcing loop multiplies money, starting with \$100 in the bank, and assuming no deposits and no withdrawals over a period of twelve years. The five lines show five different interest rates, from 2 percent to 10 percent per year.



**Figure 13.** Growth in savings with various interest rates.

This is not simple linear growth. It is not constant over time. The growth of the bank account at lower interest rates may look linear in the first few years. But, in fact, growth goes faster and faster. The more is there, the more is added. This kind of growth is called “exponential.” It’s either good news or bad news, depending on what is growing—money in the bank, people


with HIV/AIDS, pests in a cornfield, a national economy, or weapons in an arms race.

---

**Reinforcing feedback loops are self-enhancing, leading to exponential growth or to runaway collapses over time.** They are found whenever a stock has the capacity to reinforce or reproduce itself.

---

In Figure 14, the more machines and factories (collectively called “capital”) you have, the more goods and services (“output”) you can produce. The more output you can produce, the more you can invest in new machines and factories. The more you make, the more capacity you have to make even more. This reinforcing feedback loop is the central engine of growth in an economy.



**Figure 14.** Reinvestment in capital.

By now you may be seeing how basic balancing and reinforcing feedback loops are to systems. Sometimes I challenge my students to try to think of any human decision that occurs *without* a feedback loop—that is, a decision that is made without regard to any information about the level of the stock it influences. Try thinking about that yourself. The more you do, the more you’ll begin to see feedback loops everywhere.

The most common “non-feedback” decisions students suggest are falling in love and committing suicide. I’ll leave it to you to decide whether you think these are *actually* decisions made with no feedback involved.

Watch out! If you see feedback loops everywhere, you’re already in danger of becoming a systems thinker! Instead of seeing only how A causes B, you’ll begin to wonder how B may *also* influence A—and how A might reinforce or reverse itself. When you hear in the nightly news that the Federal Reserve

### HINT ON REINFORCING LOOPS AND DOUBLING TIME

Because we bump into reinforcing loops so often, it is handy to know this shortcut: The time it takes for an exponentially growing stock to double in size, the “doubling time,” equals approximately 70 divided by the growth rate (expressed as a percentage).

Example: If you put \$100 in the bank at 7% interest per year, you will double your money in 10 years ( $70 \div 7 = 10$ ). If you get only 5% interest, your money will take 14 years to double.

Bank has done something to control the economy, you'll also see that the economy must have done something to affect the Federal Reserve Bank. When someone tells you that population growth causes poverty, you'll ask yourself how poverty may cause population growth.

---

---

**THINK ABOUT THIS:**

If A causes B, is it possible that B also causes A?

You'll be thinking not in terms of a static world, but a dynamic one. You'll stop looking for who's to blame; instead you'll start asking, "What's the system?" The concept of feedback opens up the idea that a system can cause its own behavior.

So far, I have limited this discussion to one kind of feedback loop at a time. Of course, in real systems feedback loops rarely come singly. They are linked together, often in fantastically complex patterns. A single stock is likely to have several reinforcing and balancing loops of differing strengths pulling it in several directions. A single flow may be adjusted by the contents of three or five or twenty stocks. It may fill one stock while it drains another and feeds into decisions that alter yet another. The many feedback loops in a system tug against each other, trying to make stocks grow, die off, or come into balance with each other. As a result, complex systems do much more than stay steady or explode exponentially or approach goals smoothly—as we shall see.

# A Brief Visit to the Systems Zoo

---

The . . . goal of all theory is to make the . . . basic elements as simple and as few as possible without having to surrender the adequate representation of . . . experience.

—Albert Einstein,<sup>1</sup> physicist

One good way to learn something new is through specific examples rather than abstractions and generalities, so here are several common, simple but important examples of systems that are useful to understand in their own right and that will illustrate many general principles of complex systems.

This collection has some of the same strengths and weaknesses as a zoo.<sup>2</sup> It gives you an idea of the large variety of systems that exist in the world, but it is far from a complete representation of that variety. It groups the animals by family—monkeys here, bears there (single-stock systems here, two-stock systems there)—so you can observe the characteristic behaviors of monkeys, as opposed to bears. But, like a zoo, this collection is too neat. To make the animals visible and understandable, it separates them from each other and from their normal concealing environment. Just as zoo animals more naturally occur mixed together in ecosystems, so the systems animals described here normally connect and interact with each other and with others not illustrated here—all making up the buzzing, hooting, chirping, changing complexity in which we live.


Ecosystems come later. For the moment, let's look at one system animal at a time.

## One-Stock Systems

### A Stock with Two Competing Balancing Loops—a Thermostat

You already have seen the “homing in” behavior of the goal-seeking balancing feedback loop—the coffee cup cooling. What happens if there are two such loops, trying to drag a single stock toward two different goals?


One example of such a system is the thermostat mechanism that regulates the heating of your room (or cooling, if it is connected to an air conditioner instead of a furnace). Like all models, the representation of a thermostat in Figure 15 is a simplification of a real home heating system.




**Figure 15.** Room temperature regulated by a thermostat and furnace.

Whenever the room temperature falls below the thermostat setting, the thermostat detects a discrepancy and sends a signal that turns on the heat flow from the furnace, warming the room. When the room temperature rises again, the thermostat turns off the heat flow. This straightforward, stock-maintaining, balancing feedback loop is shown on the left side of Figure 15. If there were nothing else in the system, and if you start with a cold room with the thermostat set at 18°C (65°F), it would behave as shown in Figure 16. The furnace comes on, and the room warms up. When the room temperature reaches the thermostat setting, the furnace goes off, and the room stays right at the target temperature.

However, this is not the only loop in the system. Heat also leaks to the outside. The outflow of heat is governed by the second balancing feedback loop, shown on the right side of Figure 15. It is always trying to make the room temperature equal to the outside, just like a coffee cup cooling. If



**Figure 16.** A cold room warms quickly to the thermostat setting.




**Figure 17.** A warm room cools very slowly to the outside temperature of 10°C.

this were the only loop in the system (if there were no furnace), Figure 17 shows what would happen, starting with a warm room on a cold day.

The assumption is that room insulation is not perfect, and so some heat leaks out of the warm room to the cool outdoors. The better the insulation, the slower the drop in temperature would be.

Now, what happens when these two loops operate at the same time? Assuming that there is sufficient insulation and a properly sized furnace, the heating loop dominates the cooling loop. You end up with a warm room (see Figure 18), even starting with a cold room on a cold day.



**Figure 18.** The furnace warms a cool room, even as heat continues to leak from the room.

As the room heats up, the heat flowing out of it increases, because there's a larger gap between inside and outside temperatures. But the furnace keeps putting in more heat than the amount that leaks out, so the room warms nearly to the target temperature. At that point, the furnace cycles off and on as it compensates for the heat constantly flowing out of the room.

The thermostat is set at 18°C (65°F) in this simulation, but the room temperature levels off slightly below 18°C (65°F). That's because of the leak to the outside, which is draining away some heat even as the furnace is getting the signal to put it back. This is a characteristic and sometimes surprising behavior of a system with competing balancing loops. It's like trying to keep a bucket full when there's a hole in the bottom. To make things worse, water leaking out of the hole is governed by a feedback loop; the more water in the bucket, the more the water pressure at the hole increases, so the flow out increases! In this case, we are trying to keep the room warmer than the outside and the warmer the room is, the faster it loses heat to the outside. It takes time for the furnace to correct for the increased heat loss—and in that minute still more heat leaks out. In a well-insulated house, the leak will be slower and so the house more comfortable than a poorly insulated one, even a poorly insulated house with a big furnace.

With home heating systems, people have learned to set the thermostat slightly higher than the actual temperature they are aiming at. Exactly how much higher can be a tricky question, because the outflow rate is higher on cold days than on warm days. But for thermostats this control problem

isn't serious. You can muddle your way to a thermostat setting you can live with.

For other systems with this same structure of competing balancing loops, the fact that the stock goes on changing while you're trying to control it can create real problems. For example, suppose you're trying to maintain a store inventory at a certain level. You can't instantly order new stock to make up an immediately apparent shortfall. If you don't account for the goods that will be sold while you are waiting for the order to come in, your inventory will never be quite high enough. You can be fooled in the same way if you're trying to maintain a cash balance at a certain level, or the level of water in a reservoir, or the concentration of a chemical in a continuously flowing reaction system.

There's an important general principle here, and also one specific to the thermostat structure. First the general one: The information delivered by a feedback loop can only affect future behavior; it can't deliver the information, and so can't have an impact fast enough to correct behavior that drove the current feedback. A person in the system who makes a decision based on the feedback can't change the behavior of the system that drove the current feedback; the decisions he or she makes will affect only future behavior.

Why is that important? Because it means there will always be delays in responding. It says that a flow can't react instantly to a flow. It can react only to a change in a stock, and only after a slight delay to register the incoming information. In the bathtub, it takes a split second of time to assess the depth of the water and decide to adjust the flows. Many economic models make a mistake in this matter by assuming that consumption or production can respond immediately, say, to a change in price. That's one of the reasons why real economies tend not to behave exactly like many economic models.

The specific principle you can deduce from this simple system is that you must remember in thermostat-like systems to take into account whatever draining or filling processes are going on. If you don't, you won't achieve the target level of your stock. If you want your room temperature to be at 18°C (65°F), you have to set the thermostat a little above the desired

---

**The information delivered by a feedback loop—even nonphysical feedback—can only affect future behavior; it can't deliver a signal fast enough to correct behavior that drove the current feedback.** Even nonphysical information takes time to feedback into the system.

---



temperature. If you want to pay off your credit card (or the national debt), you have to raise your repayment rate high enough to cover the charges you incur while you're paying (including interest). If you're gearing up your work force to a higher level, you have to hire fast enough to correct for

---

**A stock-maintaining balancing feedback loop must have its goal set appropriately to compensate for draining or inflowing processes that affect that stock. Otherwise, the feedback process will fall short of or exceed the target for the stock.**


---

those who quit while you are hiring. In other words, your mental model of the system needs to include all the important flows, or you will be surprised by the system's behavior.

Before we leave the thermostat, we should see how it behaves with a varying outside temperature. Figure 19 shows a twenty-four-hour period of normal operation of a well-functioning thermostat system, with the outside temperature dipping well below freezing. The


inflow of heat from the furnace nicely tracks the outflow of heat to the outside. The temperature in the room varies hardly at all once the room has warmed up.

Every balancing feedback loop has its breakdown point, where other loops pull the stock away from its goal more strongly than it can pull back. That can happen in this simulated thermostat system, if I weaken the power of the heating loop (a smaller furnace that cannot put out as much heat), or if I strengthen the power of the cooling loop (colder outside tempera-



**Figure 19.** The furnace warms a cool room, even as heat leaks from the room and outside temperatures drop below freezing.

ture, less insulation, or larger leaks). Figure 20 shows what happens with the same outside temperatures as in Figure 19, but with faster heat loss from the room. At very cold temperatures, the furnace just can't keep up with the heat drain. The loop that is trying to bring the room temperature down to the outside temperature dominates the system for a while. The room gets pretty uncomfortable!



**Figure 20.** On a cold day, the furnace can't keep the room warm in this leaky house!

See if you can follow, as time unfolds, how the variables in Figure 20 relate to one another. At first, both the room and the outside temperatures are cool. The inflow of heat from the furnace exceeds the leak to the outside, and the room warms up. For an hour or two, the outside is mild enough that the furnace replaces most of the heat that's lost to the outside, and the room temperature stays near the desired temperature.


But as the outside temperature falls and the heat leak increases, the furnace cannot replace the heat fast enough. Because the furnace is generating less heat than is leaking out, the room temperature falls. Finally, the outside temperature rises again, the heat leak slows, and the furnace, still operating at full tilt, finally can pull ahead and start to warm the room again.

Just as in the rules for the bathtub, whenever the furnace is putting in more heat than is leaking out, the room temperature rises. Whenever the inflow rate falls behind the outflow rate, the temperature falls. If you

study the system changes on this graph for a while and relate them to the feedback-loop diagram of this system, you'll get a good sense of how the structural interconnections of this system—its two feedback loops and how they shift in strength relative to each other—lead to the unfolding of the system's behavior over time.

### A Stock with One Reinforcing Loop and One Balancing Loop—Population and Industrial Economy

What happens when a reinforcing and a balancing loop are both pulling on the same stock? This is one of the most common and important system structures. Among other things, it describes every living population and every economy.

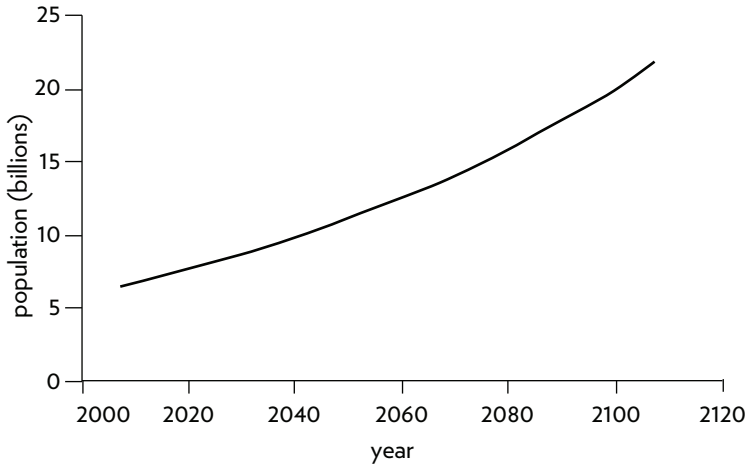


**Figure 21.** Population governed by a reinforcing loop of births and a balancing loop of deaths.

A population has a reinforcing loop causing it to grow through its birth rate, and a balancing loop causing it to die off through its death rate.

As long as fertility and mortality are constant (which in real systems they rarely are), this system has a simple behavior. It grows exponentially or dies off, depending on whether its reinforcing feedback loop determining births is stronger than its balancing feedback loop determining deaths.


For example, the 2007 world population of 6.6 billion people had a fertility rate of roughly 21 births a year for every 1,000 people in the population. Its mortality rate was 9 deaths a year out of every 1,000 people. Fertility was higher than mortality, so the reinforcing loop dominated the system. If those fertility and mortality rates continue unchanged, a child born



**Figure 22.** Population growth if fertility and mortality maintain their 2007 levels of 21 births and nine deaths per 1,000 people.

now will see the world population more than double by the time he or she reaches the age of 60, as shown in Figure 22.

If, because of a terrible disease, the mortality rate were higher, say at 30 deaths per 1,000, while the fertility rate remained at 21, then the death loop




**Figure 23.** Population decline if fertility remains at 2007 level (21 births per 1,000) but mortality is much higher, 30 deaths per 1,000.

would dominate the system. More people would die each year than would be born, and the population would gradually decrease (Figure 23).

Things get more interesting when fertility and mortality change over time. When the United Nations makes long-range population projections, it generally assumes that, as countries become more developed, average fertility will decline (approaching replacement where on average each woman has 1.85 children). Until recently, assumptions about mortality were that it would also decline, but more slowly (because it is already low in most parts of the world). However, because of the epidemic of HIV/AIDS, the UN now assumes the trend of increasing life expectancy over the next fifty years will slow in regions affected by HIV/AIDS.

Changing flows (fertility and mortality) create a change in the behavior over time of the stock (population)—the line bends. If, for example, world fertility falls steadily to equal mortality by the year 2035 and they both stay



**Figure 24.** Population stabilizes when fertility equals mortality.

constant thereafter, the population will level off, births exactly balancing deaths in dynamic equilibrium, as in Figure 24.

This behavior is an example of **shifting dominance** of feedback loops. Dominance is an important concept in systems thinking. When one loop dominates another, it has a stronger impact on behavior. Because systems often have several competing feedback loops operating simultaneously, those loops that dominate the system will determine the behavior.

At first, when fertility is higher than mortality, the reinforcing growth loop dominates the system and the resulting behavior is exponential

growth. But that loop is gradually weakened as fertility falls. Finally, it exactly equals the strength of the balancing loop of mortality. At that point neither loop dominates, and we have dynamic equilibrium.

You saw shifting dominance in the thermostat system, when the outside temperature fell and the heat leaking out of the poorly insulated house overwhelmed the ability of the furnace to put heat into the room. Dominance shifted from the heating loop to the cooling loop.

There are only a few ways a population system could behave, and these depend on what happens to the “driving” variables, fertility and mortality. These are the only ones possible for a simple system of one reinforcing and one balancing loop. A stock governed by linked reinforcing and balancing loops will grow exponentially if the reinforcing loop dominates the balancing one. It will die off if the balancing loop dominates the reinforcing one. It will level off if the two loops are of equal strength (see Figure 25). Or it will do a sequence of these things, one after another, if the relative strengths of the two loops change over time (see Figure 26).

I chose some provocative population scenarios here to illustrate a point about models and the scenarios they can generate. Whenever you are confronted with a scenario (and you are, every time you hear about an economic prediction, a corporate budget, a weather forecast, future climate change, a stockbroker saying what is going to happen to a particular holding), there are questions you need to ask that will help you decide how good a representation of reality is the underlying model.

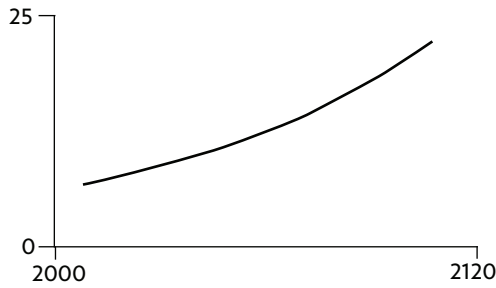


- Are the driving factors likely to unfold this way? (What are birth rate and death rate likely to do?)
- If they did, would the system react this way? (Do birth and death rates really cause the population stock to behave as we think it will?)
- What is driving the driving factors? (What affects birth rate? What affects death rate?)

The first question can't be answered factually. It's a guess about the future, and the future is inherently uncertain. Although you may have a strong

---


**Complex behaviors of systems often arise as the relative strengths of feedback loops shift, causing first one loop and then another to dominate behavior.**

---

**A: Growth****B: Decline****C: Stabilization**

**Figure 25.** Three possible behaviors of a population: growth, decline, and stabilization.

opinion about it, there's no way to prove you're right until the future actually happens. A systems analysis can test a number of scenarios to see what happens if the driving factors do different things. That's usually one purpose of a systems analysis. But you have to be the judge of which scenario, if any, should be taken seriously as a future that might really be possible.



**Figure 26.** Shifting dominance of fertility and mortality loops.

Dynamic systems studies usually are not designed to *predict* what will happen. Rather, they're designed to explore *what would happen*, if a number of driving factors unfold in a range of different ways.

The second question—whether the system really will react this way—is more scientific. It's a question about how good the model is. Does it capture the inherent dynamics of the system? Regardless of whether you think the driving factors *will* do that, *would the system behave like that* if they did?

In the population scenarios above, however likely you think they are, the answer to this second question is roughly yes, the population would behave like this, if the fertility and mortality did that. The population model I have used here is very simple. A more detailed model would distinguish age groups, for example. But basically this model responds the way a real population would, growing under the conditions when a real

---

**System dynamics models explore possible futures and ask “what if” questions.**

---

### QUESTIONS FOR TESTING THE VALUE OF A MODEL

1. Are the driving factors likely to unfold this way?
2. If they did, would the system react this way?
3. What is driving the driving factors?



**Model utility depends not on whether its driving scenarios are realistic (since no one can know that for sure), but on whether it responds with a realistic pattern of behavior.**

population would grow, declining when a real population would decline. The numbers are off, but the basic behavior pattern is realistic.


Finally, there is the third question. What is driving the driving factors? What is adjusting the inflows and outflows? This is a question about system boundaries. It requires a hard look at those driving factors to see if they are actually independent, or if they are also embedded in the

system.

Is there anything about the size of the population, for instance, that might feed back to influence fertility or mortality? Do other factors—economics, the environment, social trends—influence fertility and mortality? Does the size of the population affect those economic and environmental and social factors?

Of course, the answer to all of these questions is yes. Fertility and mortality are governed by feedback loops too. At least some of those feedback loops are themselves affected by the size of the population. This population “animal” is only one piece of a much larger system.<sup>3</sup>

One important piece of the larger system that affects population is the economy. At the heart of the economy is another reinforcing-loop-plus-balancing-loop system—the same kind of structure, with the same kinds



**Figure 27.** Like a living population, economic capital has a reinforcing loop (investment of output) governing growth and a balancing loop (depreciation) governing decline.

of behavior, as the population (see Figure 27).

The greater the stock of physical capital (machines and factories) in the economy and the efficiency of production (output per unit of capital), the more output (goods and services) can be produced each year.

The more output that is produced, the more can be invested to make new capital. This is a reinforcing loop, like the birth loop for a population. The investment fraction is equivalent to the fertility. The greater the fraction of its output a society invests, the faster its capital stock will grow.


Physical capital is drained by depreciation—obsolescence and wearing-out. The balancing loop controlling depreciation is equivalent to the death loop in a population. The “mortality” of capital is determined by the average capital lifetime. The longer the lifetime, the smaller the fraction of capital that must be retired and replaced each year.

If this system has the same structure as the population system, it must have the same repertoire of behaviors. Over recent history world capital, like world population, has been dominated by its reinforcing loop and has been growing exponentially. Whether in the future it grows or stays constant or dies off depends on whether its reinforcing growth loop remains stronger than its balancing depreciation loop. That depends on

- the investment fraction—how much output the society invests rather than consumes,
- the efficiency of capital—how much capital it takes to produce a given amount of output, and
- the average capital lifetime.

If a constant fraction of output is reinvested in the capital stock and the efficiency of that capital (its ability to produce output) is also constant, the capital stock may decline, stay constant, or grow, depending on the lifetime of the capital. The lines in Figure 28 show systems with different average capital lifetimes. With a relatively short lifetime, the capital wears out faster than it is replaced. Reinvestment does not keep up with depreciation and the economy slowly declines. When depreciation just balances investment, the economy is in dynamic equilibrium. With a long lifetime, the capital stock grows exponentially. The longer the lifetime of capital, the faster it grows.

This is another example of a principle we’ve already encountered: You can make a stock grow by decreasing its outflow rate as well as by increas-



**Figure 28.** Growth in capital stock with changes in the lifetime of the capital. In a system with output per unit capital ratio of 1:3 and an investment fraction of 20 percent, capital with a lifetime of 15 years just keeps up with depreciation. A shorter lifetime leads to a declining stock of capital.

ing its inflow rate!

Just as many factors influence the fertility and mortality of a population, so many factors influence the output ratio, investment fraction, and the lifetime of capital—interest rates, technology, tax policy, consumption habits, and prices, to name just a few. Population itself influences investment, both by contributing labor to output, and by increasing demands on consumption, thereby decreasing the investment fraction. Economic output also feeds back to influence population in many ways. A richer economy usually has better health care and a lower death rate. A richer economy also usually has a lower birth rate.

In fact, just about any long-term model of a real economy should link together the two structures of population and capital to show how they affect each other. The central question of economic development is how to keep the reinforcing loop of capital accumulation from growing more slowly than the reinforcing loop of population growth—so that people are getting richer instead of poorer.<sup>4</sup>

**Systems with similar feedback structures produce similar dynamic behaviors.**

It may seem strange to you that I call the capital system the same kind of “zoo animal” as the population system. A production system with factories and shipments and economic flows doesn’t look much like a population system with babies being born and people aging


and having more babies and dying. But from a systems point of view these systems, so dissimilar in many ways, have one important thing in common: their feedback-loop structures. Both have a stock governed by a reinforcing growth loop and a balancing death loop. Both also have an aging process. Steel mills and lathes and turbines get older and die just as people do.

One of the central insights of systems theory, as central as the observation that systems largely cause their own behavior, is that systems with similar feedback structures produce similar dynamic behaviors, even if the outward appearance of these systems is completely dissimilar.

A population is nothing like an industrial economy, except that both can reproduce themselves out of themselves and thus grow exponentially. And both age and die. A coffee cup cooling is like a warmed room cooling, and like a radioactive substance decaying, and like a population or industrial economy aging and dying. Each declines as the result of a balancing feedback loop.

#### A System with Delays—Business Inventory

Picture a stock of inventory in a store—a car dealership—with an inflow of deliveries from factories and an outflow of new car sales. By itself, this stock of cars on the dealership lot would behave like the water in a bathtub.



**Figure 29.** Inventory at a car dealership is kept steady by two competing balancing loops, one through sales and one through deliveries.


Now picture a regulatory feedback system designed to keep the inventory high enough so that it can always cover ten days' worth of sales (see Figure 29). The car dealer needs to keep some inventory because deliveries and purchases don't match perfectly every day. Customers make purchases that are unpredictable on a day-to-day basis. The car dealer also needs to provide herself with some extra inventory (a buffer) in case deliveries from suppliers are delayed occasionally.

The dealer monitors sales (perceived sales), and if, for example, they seem to be rising, she adjusts orders to the factory to bring inventory up to her new desired inventory that provides ten days' coverage at the higher sales rate. So, higher sales mean higher perceived sales, which means a higher discrepancy between inventory and desired inventory, which means higher orders, which will bring in more deliveries, which will raise inventory so it can comfortably supply the higher rate of sales.

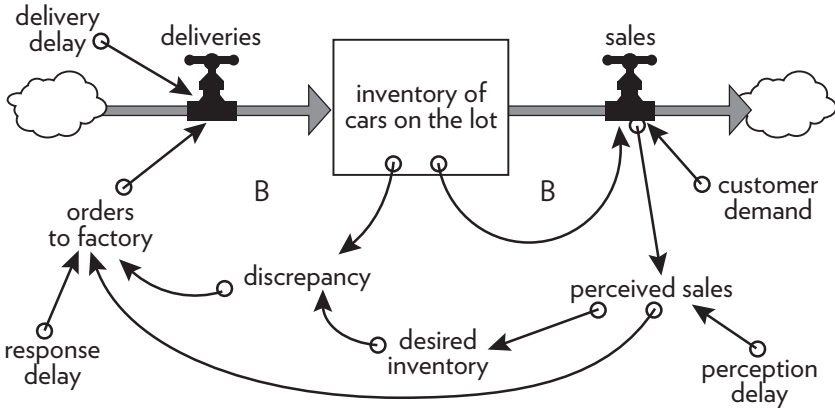
This system is a version of the thermostat system—one balancing loop of sales draining the inventory stock and a competing balancing loop maintaining the inventory by resupplying what is lost in sales. Figure 30 shows the not very surprising result of an increase in consumer demand of 10 percent.

In Figure 31, I am putting something else into this simple model—three delays that are typical of what we experience in the real world.

First, there is a perception delay, intentional in this case. The car dealer doesn't react to just any blip in sales. Before she makes ordering decisions,



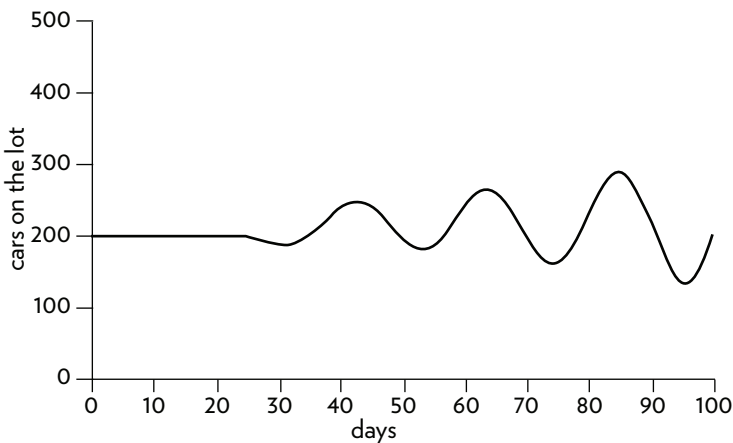
**Figure 30.** Inventory on the car dealership's lot with a permanent 10-percent increase in consumer demand starting on day 25.



**Figure 31.** Inventory at a car dealership with three common delays now included in the picture—a perception delay, a response delay, and a delivery delay.

she averages sales over the past five days to sort out real trends from temporary dips and spikes.

Second, there is a response delay. Even when it’s clear that orders need to be adjusted, she doesn’t try to make up the whole adjustment in a single order. Rather, she makes up one-third of any shortfall with each order. Another way of saying that is, she makes partial adjustments over three days to be extra sure the trend is real. Third, there is a delivery delay. It takes five days for the supplier at the factory to receive an order, process it, and



**Figure 32.** Response of inventory to a 10-percent increase in sales when there are delays in the system.

deliver it to the dealership.

Although this system still consists of just two balancing loops, like the simplified thermostat system, it doesn't behave like the thermostat system. Look at what happens, for example, as shown in Figure 32, when the business experiences the same permanent 10-percent jump in sales from an increase in customer demand.

Oscillations! A single step up in sales causes inventory to drop. The car dealer watches long enough to be sure the higher sales rate is going to last. Then she begins to order more cars to both cover the new rate of sales and bring the inventory up. But it takes time for the orders to come in. During that time inventory drops further, so orders have to go up a little more, to bring inventory back up to ten days' coverage.

Eventually, the larger volume of orders starts arriving, and inventory recovers—and more than recovers, because during the time of uncertainty about the actual trend, the owner has ordered too much. She now sees her mistake, and cuts back, but there are still high past orders coming in, so she orders even less. In fact, almost inevitably, since she still can't be sure of what is going to happen next, she orders too little. Inventory gets too low again. And so forth, through a series of oscillations around the new desired inventory level. As Figure 33 illustrates, what a difference a few delays make!

We'll see in a moment that there are ways to damp these oscillations in inventory, but first it's important to understand why they occur. It isn't because the car dealer is stupid. It's because she is struggling to operate


---

**A delay in a balancing feedback loop makes a system likely to oscillate.**

---

in a system in which she doesn't have, and can't have, timely information and in which physical delays prevent her actions from having an immediate effect on inventory. She doesn't know what her customers will do next. When they do something, she's not sure they'll keep doing it. When she issues an order, she doesn't see an immediate response. This situation of information insufficiency and physical delays is very common. Oscillations like these are frequently encountered in inventories and in many other systems. Try taking a shower sometime where there's a very long pipe between the hot- and cold-water mixer and the showerhead, and you'll experience directly the joys of hot and cold oscillations because of a long response delay.

How much of a delay causes what kind of oscillation under what circum-



**Figure 33.** The response of orders and deliveries to an increase in demand. *A* shows the small but sharp step up in sales on day 25 and the car dealer’s “perceived” sales, in which she averages the change over 3 days. *B* shows the resulting ordering pattern, tracked by the actual deliveries from the factory.

stances is not a simple matter. I can use this inventory system to show you why.

“These oscillations are intolerable,” says the car dealer (who is herself a learning system, determined now to change the behavior of the inventory system). “I’m going to shorten the delays. There’s not much I can do about the delivery delay from the factory, so I’m going to react faster myself. I’ll average sales trends over only two days instead of five before I make order adjustments.”


Figure 34 illustrates what happens when the dealer’s perception delay is




shortened from five days to two.

Not much happens when the car dealer shortens her perception delay. If anything the oscillations in the inventory of cars on the lot are a bit worse. And if, instead of shortening her perception time, the car dealer tries shortening her reaction time—making up perceived shortfalls in two days instead of three—things get very much worse, as shown in Figure 35.

Something has to change and, since this system has a learning person



**Figure 34.** The response of inventory to the same increase in demand with a shortened perception delay.



**Figure 35.** The response of inventory to the same increase in demand with a shortened reaction time. Acting faster makes the oscillations worse!

within it, something will change. “High leverage, wrong direction,” the system-thinking car dealer says to herself as she watches this failure of a policy intended to stabilize the oscillations. This perverse kind of result can be seen all the time—someone trying to fix a system is attracted intuitively to a policy lever that in fact does have a strong effect on the system. And then the well-intentioned fixer pulls the lever in the wrong direction! This is just one example of how we can be surprised by the counterintuitive behavior of systems when we start trying to change them.

Part of the problem here is that the car dealer has been reacting not too slowly, but too quickly. Given the configuration of this system, she has been overreacting. Things would go better if, instead of decreasing her response delay from three days to two, she would increase the delay from three days to six, as illustrated in Figure 36.


As Figure 36 shows, the oscillations are greatly damped with this change, and the system finds its new equilibrium fairly efficiently.

The most important delay in this system is the one that is not under the direct control of the car dealer. It’s the delay in delivery from the factory. But even without the ability to change that part of her system, the dealer can learn to manage inventory quite well.

---

Delays are pervasive in systems, and they are strong determinants of behavior. **Changing the length of a delay may** (or may not, depending on the type of delay and the relative lengths of other delays) **make a large change in the behavior of a system.**

---



**Figure 36.** The response of inventory to the same increase in demand with a slowed reaction time.

Changing the delays in a system can make it much easier or much harder to manage. You can see why system thinkers are somewhat fanatic on the subject of delays. We're always on the alert to see where delays occur in systems, how long they are, whether they are delays in information streams or in physical processes. We can't begin to understand the dynamic behavior of systems unless we know where and how long the delays are. And we are aware that some delays can be powerful policy levers. Lengthening or shortening them can produce major changes in the behavior of systems.

In the big picture, one store's inventory problem may seem trivial and fixable. But imagine that the inventory is that of all the unsold automobiles in America. Orders for more or fewer cars affect production not only at assembly plants and parts factories, but also at steel mills, rubber and glass plants, textile producers, and energy producers. Everywhere in this system are perception delays, production delays, delivery delays, and construction delays. Now consider the link between car production and jobs—increased production increases the number of jobs allowing more people to buy cars. That's a reinforcing loop, which also works in the opposite direction—less production, fewer jobs, fewer car sales, less production. Put in another reinforcing loop, as speculators buy and sell shares in the auto and auto-supply companies based on their recent performance, so that an upsurge in production produces an upsurge in stock price, and vice versa.

That very large system, with interconnected industries responding to each other through delays, entraining each other in their oscillations, and being amplified by multipliers and speculators, is the primary cause of business cycles. Those cycles don't come from presidents, although presidents can do much to ease or intensify the optimism of the upturns and the pain of the downturns. Economies are extremely complex systems; they are full of balancing feedback loops with delays, and they are inherently oscillatory.<sup>5</sup>

## Two-Stock Systems

A Renewable Stock Constrained by a Nonrenewable Stock—an Oil Economy  
The systems I've displayed so far have been free of constraints imposed by their surroundings. The capital stock of the industrial economy model didn't require raw materials to produce output. The population didn't need food. The thermostat-furnace system never ran out of oil. These simple

models of the systems have been able to operate according to their unconstrained internal dynamics, so we could see what those dynamics are.

But any real physical entity is always surrounded by and exchanging things with its environment. A corporation needs a constant supply of energy and materials and workers and managers and customers. A growing corn crop needs water and nutrients and protection from pests. A population needs food and water and living space, and if it's a human population, it needs jobs and education and health care and a multitude of other things. Any entity that is using energy and processing materials needs a place to put its wastes, or a process to carry its wastes away.

Therefore, any physical, growing system is going to run into some kind of constraint, sooner or later. That constraint will take the form of a balancing loop that in some way shifts the dominance of the reinforcing loop driving the growth behavior, either by strengthening the outflow or by weakening the inflow.

Growth in a constrained environment is very common, so common that systems thinkers call it the “limits-to-growth” archetype. (We'll explore more archetypes—frequently found system structures that produce familiar behavior patterns—in Chapter Five.) Whenever we see a growing entity, whether it be a population, a corporation, a bank account, a rumor, an epidemic, or sales of a new product, we look for the reinforcing loops that are driving it and for the balancing loops that ultimately will constrain it. We know those balancing loops are there, even if they are not yet dominating the system's behavior, because no real physical system can grow forever. Even a hot new product will saturate the market eventually. A chain reaction in a nuclear power plant or bomb will run out of fuel. A virus will run out of susceptible people to infect. An economy may be constrained by physical capital or monetary capital or labor or markets or management or resources or pollution.

---

**In physical, exponentially growing systems, there must be at least one reinforcing loop driving the growth *and* at least one balancing loop constraining the growth, because no physical system can grow forever in a finite environment.**

---

Like resources that supply the inflows to a stock, a pollution constraint can be renewable or nonrenewable. It's nonrenewable if the environment has no capacity to absorb the pollutant or make it harmless. It's renewable if the environment has a finite, usually variable, capacity for removal. Everything said here about resource-constrained systems, therefore,

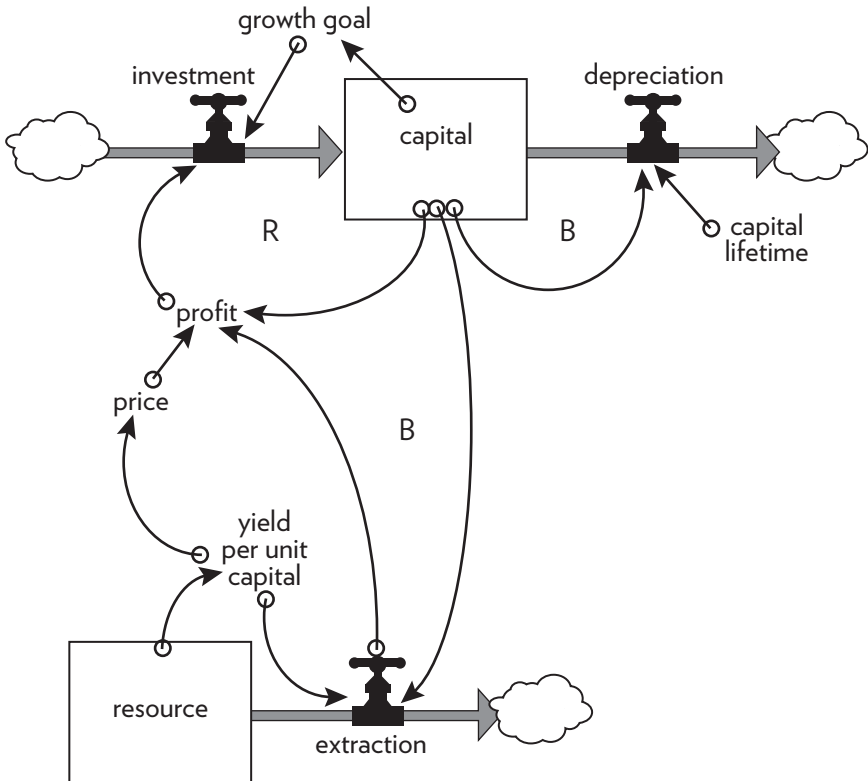
applies with the same dynamics but opposite flow directions to pollution-constrained systems.

The limits on a growing system may be temporary or permanent. The system may find ways to get around them for a short while or a long while, but eventually there must come some kind of accommodation, the system adjusting to the constraint, or the constraint to the system, or both to each other. In that accommodation come some interesting dynamics.

Whether the constraining balancing loops originate from a renewable or nonrenewable resource makes some difference, not in whether growth can continue forever, but in how growth is likely to end.

Let's look, to start, at a capital system that makes its money by extracting a nonrenewable resource—say an oil company that has just discovered a huge new oil field. See Figure 37.

The diagram in Figure 37 may look complicated, but it's no more than



**Figure 37.** Economic capital, with its reinforcing growth loop constrained by a nonrenewable resource.


a capital-growth system like the one we've already seen, using "profit" instead of "output." Driving depreciation is the now-familiar balancing loop: the more capital stock, the more machines and refineries there are that fall apart and wear out, reducing the stock of capital. In this example, the capital stock of oil drilling and refining equipment depreciates with a 20-year lifetime—meaning  $1/20$  (or 5 percent) of the stock is taken out of commission each year. It builds itself up through investment of profits from oil extraction. So we see the reinforcing loop: More capital allows more resource extraction, creating more profits that can be reinvested. I've assumed that the company has a goal of 5 percent annual growth in its business capital. If there isn't enough profit for 5 percent growth, the company invests whatever profits it can.

Profit is income minus cost. Income in this simple representation is just the price of oil times the amount of oil the company extracts. Cost is equal to capital times the operating cost (energy, labor, materials, etc.) per unit of capital. For the moment, I'll make the simplifying assumptions that both price and operating cost per unit of capital are constant.

What is not assumed to be constant is the yield of resource per unit of capital. Because this resource is not renewable, as in the case of oil, the stock feeding the extraction flow does not have an input. As the resource is extracted—as an oil well is depleted—the next barrel of oil becomes harder to get. The remaining resource is deeper down, or more dilute, or in the case of oil, under less natural pressure to force it to the surface. More and more costly and technically sophisticated measures are required to keep the resource coming.

Here is a new balancing feedback loop that ultimately will control the growth of capital: the more capital, the higher the extraction rate. The higher the extraction rate, the lower the resource stock. The lower the resource stock, the lower the yield of resource per unit of capital, so the lower the profit (with price assumed constant) and the lower the investment rate—therefore, the lower the rate of growth of capital. I could assume that resource depletion feeds back through operating cost as well as capital efficiency. In the real world it does both. In either case, the ensuing behavior pattern is the same—the classic dynamics of depletion (see Figure 38).

The system starts out with enough oil in the underground deposit to supply the initial scale of operation for 200 years. But, actual extraction peaks at about 40 years because of the surprising effect of exponential



**Figure 38.** Extraction (A) creates profits that allow for growth of capital (B) while depleting the nonrenewable resource (C). The greater the accumulation of capital, the faster the resource is depleted.

growth in extraction. At an investment rate of 10 percent per year, the capital stock and therefore the extraction rate both grow at 5 percent per year and so double in the first 14 years. After 28 years, while the capital stock has quadrupled, extraction is starting to lag because of falling yield per unit of capital. By year 50 the cost of maintaining the capital stock has overwhelmed the income from resource extraction, so profits are no longer sufficient to keep investment ahead of depreciation. The operation quickly shuts down, as the capital stock declines. The last and most expensive of the resource stays in the ground; it doesn't pay to get it out.

What happens if the original resource turns out to be twice as large as

the geologists first thought—or four times as large? Of course, that makes a huge difference in the total amount of oil that can be extracted from this field. But with the continued goal of 10 percent per year reinvestment producing 5 percent per year capital growth, each doubling of the resource makes a difference of only about 14 years in the timing of the peak extraction rate, and in the lifetime of any jobs or communities dependent on the extraction industry (see Figure 39).


---

**A quantity growing exponentially toward a constraint or limit reaches that limit in a surprisingly short time.**

---


The higher and faster you grow, the farther and faster you fall, when you're building up a capital stock dependent on a nonrenewable resource. In the face of exponential growth of extraction or use, a doubling or quadrupling of the nonrenewable resource give little added time to develop alternatives.

If your concern is to extract the resource and make money at the maximum possible rate, then the ultimate size of the resource is the most important number in this system. If, say, you're a worker at the mine or oil field, and your concern is with the lifetime of your job and stability of your community, then there are two important numbers: the size of the resource and the desired growth rate of capital. (Here is a good example of the goal of a feedback loop being crucial to the behavior of a system.) The real choice in the management of a nonrenewable resource is whether to get rich very fast or to get less rich but stay that way longer.



**Figure 39.** Extraction with two times or four times as large a resource to draw on. Each doubling of the resource makes a difference of only about fourteen years in the peak of extraction.






**Figure 40.** The peak of extraction comes much more quickly as the fraction of profits reinvested increases.

The graph in Figure 40 shows the development of the extraction rate over time, given desired growth rates above depreciation varying from 1 percent annually, to 3 percent, 5 percent, and 7 percent. With a 7 percent growth rate, extraction of this “200-year supply” peaks within 40 years. Imagine the effects of this choice not only on the profits of the company, but on the social and natural environments of the region.

Earlier I said I would make the simplifying assumption that price was constant. But what if that’s not true? Suppose that in the short term the resource is so vital to consumers that a higher price won’t decrease demand. In that case, as the resource gets scarce and price rises steeply, as shown in Figure 41.

The higher price gives the industry higher profits, so investment goes up, capital stock continues rising, and the more costly remaining resources can be extracted. If you compare Figure 41 with Figure 38, where price was held constant, you can see that the main effect of rising price is to build the capital stock higher before it collapses.

The same behavior results, by the way, if prices don’t go up but if technology brings operating costs down—as has actually happened, for example, with advanced recovery techniques from oil wells, with the beneficiation process to extract low-grade taconite from exhausted iron mines, and with the cyanide leaching process that allows profitable extraction even from the tailings of gold and silver mines.



**Figure 41.** As price goes up with increasing scarcity, there is more profit to reinvest, and the capital stock can grow larger (B) driving extraction up for longer (A). The consequence is that the resource (C) is depleted even faster at the end.

We all know that individual mines and fossil fuel deposits and groundwater aquifers can be depleted. There are abandoned mining towns and oil fields all over the world to testify to the reality of the behavior we've seen here. Resource companies understand this dynamic too. Well before depletion makes capital less efficient in one place, companies shift investment to discovery and development of another deposit somewhere else. But, if there are local limits, eventually will there be global ones?

I'll leave you to have this argument with yourself, or with someone of the

opposite persuasion. I will just point out that, according to the dynamics of depletion, the larger the stock of initial resources, the more new discoveries, the longer the growth loops elude the control loops, and the higher the capital stock and its extraction rate grow, and the earlier, faster, and farther will be the economic fall on the back side of the production peak.

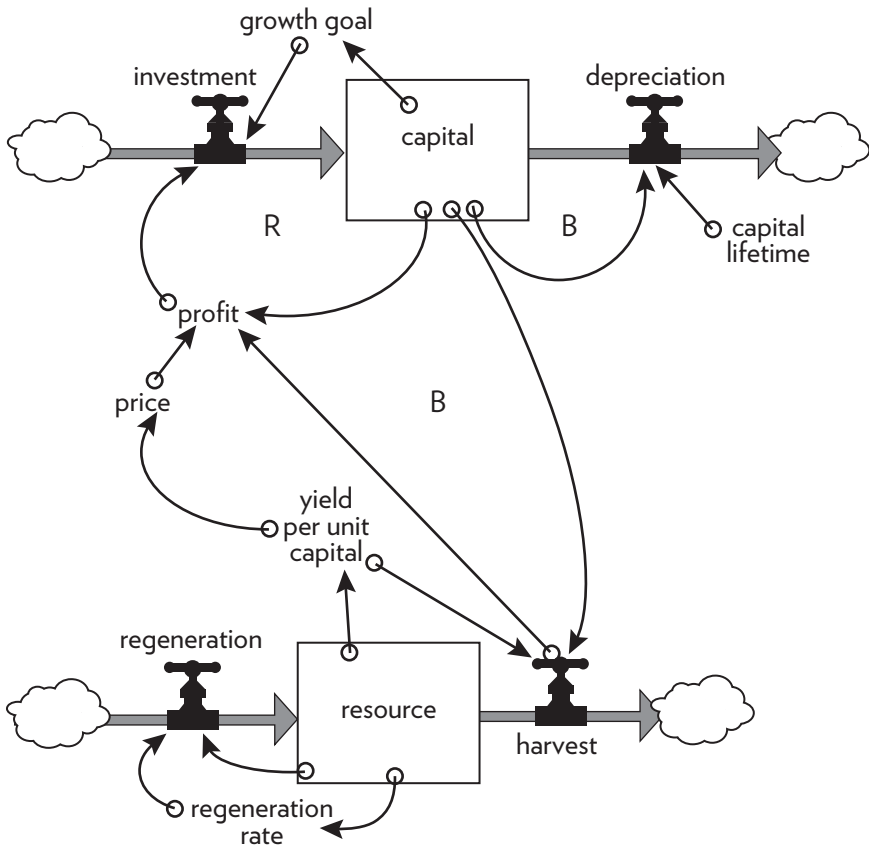
Unless, perhaps, the economy can learn to operate entirely from renewable resources.

### Renewable Stock Constrained by a Renewable Stock—a Fishing Economy

Assume the same capital system as before, except that now there is an inflow to the resource stock, making it renewable. The renewable resource in this system could be fish and the capital stock could be fishing boats. It also could be trees and sawmills, or pasture and cows. Living renewable resources such as fish or trees or grass can regenerate themselves from themselves with a reinforcing feedback loop. Nonliving renewable resources such as sunlight or wind or water in a river are regenerated not through a reinforcing loop, but through a steady input that keeps refilling the resource stock no matter what the current state of that stock might be. This same “renewable resource system” structure occurs in an epidemic of a cold virus. It spares its victims who are then able to catch another cold. Sales of a product people need to buy regularly is also a renewable resource system; the stock of potential customers is ever regenerated. Likewise an insect infestation that destroys part but not all of a plant; the plant can regenerate and the insect can eat more. In all these cases, there is an input that keeps refilling the constraining resource stock (as shown in Figure 42).

We will use the example of a fishery. Once again, assume that the lifetime of capital is 20 years and the industry will grow, if it can, at 5 percent per year. As with the nonrenewable resource, assume that as the resource gets scarce it costs more, in terms of capital, to harvest it. Bigger fishing boats that can go longer distances and are equipped with sonar are needed to find the last schools of fish. Or miles-long drift nets are needed to catch them. Or on-board refrigeration systems are needed to bring them back to port from longer distances. All this takes more capital.

The regeneration rate of the fish is not constant, but is dependent on the number of fish in the area—fish density. If the fish are very dense, their reproduction rate is near zero, limited by available food and habitat. If the fish population falls a bit, it can regenerate at a faster and faster rate,



**Figure 42.** Economic capital with its reinforcing growth loop constrained by a renewable resource.


because it can take advantage of unused nutrients or space in the ecosystem. But at some point the fish reproduction rate reaches its maximum. If the population is further depleted, it breeds not faster and faster, but slower and slower. That's because the fish can't find each other, or because another species has moved into its niche.

This simplified model of a fishery economy is affected by three nonlinear relationships: price (scarcer fish are more expensive); regeneration rate (scarcer fish don't breed much, nor do crowded fish); and yield per unit of capital (efficiency of the fishing technology and practices).

This system can produce many different sets of behaviors. Figure 43 shows one of them.

In Figure 43, we see capital and fish harvest rise exponentially at first.


The fish population (the resource stock) falls, but that stimulates the fish reproduction rate. For decades the resource can go on supplying an exponentially increasing harvest rate. Eventually, the harvest rises too far and the fish population falls low enough to reduce the profitability of the fishing fleet. The balancing feedback of falling harvest reducing profits brings





**Figure 43.** Annual harvest (A) creates profits that allow for growth of capital stock (B), but the harvest levels off, after a small overshoot in this case. The result of leveling harvest is that the resource stock (C) also stabilizes.

down the investment rate quickly enough to bring the fishing fleet into equilibrium with the fish resource. The fleet can't grow forever, but it can maintain a high and steady harvest rate forever.

Just a minor change in the strength of the controlling balancing feedback loop through yield per unit of capital, however, can make a surpris-



**Figure 44.** A slight increase in yield per unit of capital—increasingly efficient technology in this case—creates a pattern of overshoot and oscillation around a stable value in the harvest rate (A), the stock of economic capital (B), and in the resource stock.

**A: Harvest rate****B: Capital stock****C: Resource stock**

**Figure 45.** An even greater increase in yield per unit of capital creates a patterns of overshoot and collapse in the harvest (A), the economic capital (B), and the resource (C).

ing difference. Suppose that in an attempt to raise the catch in the fishery, the industry comes up with a technology to improve the efficiency of the boats (sonar, for example, to find the scarcer fish). As the fish population declines, the fleet's ability to pull in the same catch per boat is maintained just a little longer (see Figure 44).

Figure 44 shows another case of high leverage, wrong direction! This

technical change, which increases the productivity of all fishermen, throws the system into instability. Oscillations appear!

If the fishing technology gets even better, the boats can go on operating economically even at very low fish densities. The result can be a nearly complete wipeout both of the fish and of the fishing industry. The consequence is the marine equivalent of desertification. The fish have been turned, for all practical purposes, into a nonrenewable resource. Figure 45 illustrates this scenario.

In many real economies based on real renewable resources—as opposed to this simple model—the very small surviving population retains the potential to build its numbers back up again, once the capital driving the harvest is gone. The whole pattern is repeated, decades later. Very long-term renewable-resource cycles like these have been observed, for example, in the logging industry in New England, now in its third cycle of growth, overcutting, collapse, and eventual regeneration of the resource. But this is not true for all resource populations. More and more, increases in technology and harvest efficiency have the ability to drive resource populations to extinction.

Whether a real renewable resource system can survive overharvest depends on what happens to it during the time when the resource is severely depleted. A very small fish population may become especially vulnerable to pollution or storms or lack of genetic diversity. If this is a forest or grassland resource, the exposed soils may be vulnerable to erosion. Or the nearly empty ecological niche may be filled in by a competitor. Or perhaps the depleted resource can survive and rebuild itself again.

I've shown three sets of possible behaviors of this renewable resource system here:

- overshoot and adjustment to a sustainable equilibrium,

---

**Nonrenewable resources are stock-limited.** The entire stock is available at once, and can be extracted at any rate (limited mainly by extraction capital). But since the stock is not renewed, the faster the extraction rate, the shorter the lifetime of the resource.

---



---

**Renewable resources are flow-limited.** They can support extraction or harvest indefinitely, but only at a finite flow rate equal to their regeneration rate. If they are extracted faster than they regenerate, they may eventually be driven below a critical threshold and become, for all practical purposes, nonrenewable.

---



- overshoot beyond that equilibrium followed by oscillation around it, and
- overshoot followed by collapse of the resource and the industry dependent on the resource.

Which outcome actually occurs depends on two things. The first is the critical threshold beyond which the resource population's ability to regenerate itself is damaged. The second is the rapidity and effectiveness of the balancing feedback loop that slows capital growth as the resource becomes depleted. If the feedback is fast enough to stop capital growth before the critical threshold is reached, the whole system comes smoothly into equilibrium. If the balancing feedback is slower and less effective, the system oscillates. If the balancing loop is very weak, so that capital can go on growing even as the resource is reduced below its threshold ability to regenerate itself, the resource and the industry both collapse.

Neither renewable nor nonrenewable limits to growth allow a physical stock to grow forever, but the constraints they impose are dynamically quite different. The difference comes because of the difference between stocks and flows.

The trick, as with all the behavioral possibilities of complex systems, is to recognize what structures contain which latent behaviors, and what conditions release those behaviors—and, where possible, to arrange the structures and conditions to reduce the probability of destructive behaviors and to encourage the possibility of beneficial ones.