

1 Why Systems Surprise Us

The trouble . . . is that we are terrifyingly ignorant. The most learned of us are ignorant. . . . The acquisition of knowledge always involves the revelation of ignorance—almost *is* the revelation of ignorance. Our knowledge of the world instructs us first of all that the world is greater than our knowledge of it.

—Wendell Berry,¹ writer and Kentucky farmer

The simple systems in the zoo may have perplexed you with their behavior. They continue to surprise me, although I have been teaching them for years. That you and I are surprised says as much about us as it does about dynamic systems. The interactions between what I think I know about dynamic systems and my experience of the real world never fails to be humbling. They keep reminding me of three truths:

1. Everything we think we know about the world is a model. Every word and every language is a model. All maps and statistics, books and databases, equations and computer programs are models. So are the ways I picture the world in my head—my *mental* models. None of these is or ever will be the *real* world.
2. Our models usually have a strong congruence with the world. That is why we are such a successful species in the biosphere. Especially complex and sophisticated are the mental models we develop from direct, intimate experience of nature, people, and organizations immediately around us.
3. However, and conversely, our models fall far short of representing the world fully. That is why we make mistakes and why we are regularly surprised. In our heads, we can keep track of only a few variables at

one time. We often draw illogical conclusions from accurate assumptions, or logical conclusions from inaccurate assumptions. Most of us, for instance, are surprised by the amount of growth an exponential process can generate. Few of us can intuit how to damp oscillations in a complex system.

In short, this book is poised on a duality. We know a tremendous amount about how the world works, but not nearly enough. Our knowledge is amazing; our ignorance even more so. We can improve our understanding, but we can't make it perfect. I believe both sides of this duality, because I have learned much from the study of systems.

This chapter describes some of the reasons why dynamic systems are so often surprising. Alternately, it is a compilation of some of the ways our mental models fail to take into account the complications of the real world—at least those ways that one can see from a systems perspective. It is a warning list. Here is where hidden snags lie. You can't navigate well in an interconnected, feedback-dominated world unless you take your eyes off short-term events and look for long-term behavior and structure; unless you are aware of false boundaries and bounded rationality; unless you take into account limiting factors, nonlinearities and delays. You are likely to mistreat, misdesign, or misread systems if you don't respect their properties of resilience, self-organization, and hierarchy.

The bad news, or the good news, depending on your need to control the world and your willingness to be delighted by its surprises, is that even if you do understand all these system characteristics, you may be surprised less often, but you will still be surprised.

Everything we think we know about the world is a model. Our models do have a strong congruence with the world. Our models fall far short of representing the real world fully.

2

Beguiling Events

A system is a big black box
Of which we can't unlock the locks,
And all we can find out about
Is what goes in and what comes out.

Perceiving input-output pairs,
Related by parameters,
Permits us, sometimes, to relate
An input, output and a state.
If this relation's good and stable
Then to predict we may be able,
But if this fails us—heaven forbid!
We'll be compelled to force the lid!

—Kenneth Boulding,² economist

Systems fool us by presenting themselves—or we fool ourselves by seeing the world—as a series of events. The daily news tells of elections, battles, political agreements, disasters, stock market booms or busts. Much of our ordinary conversation is about specific happenings at specific times and places. A team wins. A river floods. The Dow Jones Industrial Average hits 10,000. Oil is discovered. A forest is cut. Events are the outputs, moment by moment, from the black box of the system.

Events can be spectacular: crashes, assassinations, great victories, terrible tragedies. They hook our emotions. Although we've seen many thousands of them on our TV screens or the front page of the paper, each one is different enough from the last to keep us fascinated (just as we never lose our fascination with the chaotic twists and turns of the weather). It's endlessly engrossing to take in the world as a series of events, and constantly surprising, because that way of seeing the world has almost no predictive or explanatory value. Like the tip of an iceberg rising above the water, events are the most visible aspect of a larger complex—but not always the most important.

We are less likely to be surprised if we can see how events accumulate into dynamic patterns of *behavior*. The team is on a winning streak. The variance of the river is increasing, with higher floodwaters during rains and lower flows during droughts. The Dow has been trending up for two years. Discoveries of oil are becoming less frequent. The felling of forests is happening at an ever-increasing rate.

The behavior of a system is its performance over time—its growth, stagnation, decline, oscillation, randomness, or evolution. If the news did a better job of putting events into historical context, we would have better behavior-level understanding, which is deeper than event-level under-

standing. When a systems thinker encounters a problem, the first thing he or she does is look for data, time graphs, the history of the system. That's because long-term behavior provides clues to the underlying system structure. And structure is the key to understanding not just *what* is happening, but *why*.

The structure of a system is its interlocking stocks, flows, and feedback loops. The diagrams with boxes and arrows (my students call them “spaghetti-and-meatball diagrams”) are pictures of system structure. Structure determines what behaviors are latent in the system. A goal-seeking balancing feedback loop approaches or holds a dynamic equilibrium. A reinforcing feedback loop generates exponential growth. The two of them linked together are capable of growth, decay, or equilibrium. If they also contain delays, they may produce oscillations. If they work in periodic, rapid bursts, they may produce even more surprising behaviors.

System structure is the source of system behavior. System behavior reveals itself as a series of events over time.

Systems thinking goes back and forth constantly between structure (diagrams of stocks, flows, and feedback) and behavior (time graphs). Systems thinkers strive to understand the connections between the hand releasing the Slinky (event) and the resulting oscillations (behavior) and the mechanical characteristics of the Slinky's helical coil (structure).

Simple examples like a Slinky make this event-behavior-structure distinction seem obvious. In fact, much analysis in the world goes no deeper than events. Listen to every night's explanation of why the stock market did what it did. Stocks went up (down) because the U.S. dollar fell (rose), or the prime interest rate rose (fell), or the Democrats won (lost), or one country invaded another (or didn't). Event-event analysis.

These explanations give you no ability to predict what will happen tomorrow. They give you no ability to change the behavior of the system—to make the stock market less volatile or a more reliable indicator of the health of corporations or a better vehicle to encourage investment, for instance.

Most economic analysis goes one level deeper, to behavior over time. Econometric models strive to find the statistical links among past trends in income, savings, investment, government spending, interest rates, output, or whatever, often in complicated equations.

These behavior-based models are more useful than event-based ones, but they still have fundamental problems. First, they typically overemphasize system flows and underemphasize stocks. Economists follow the behavior of flows, because that's where the interesting variations and most rapid changes in systems show up. Economic news reports on the national production (flow) of goods and services, the GNP, rather than the total physical capital (stock) of the nation's factories and farms and businesses that produce those goods and services. But without seeing how stocks affect their related flows through feedback processes, one cannot understand the dynamics of economic systems or the reasons for their behavior.

Second, and more seriously, in trying to find statistical links that relate flows to each other, econometricians are searching for something that does not exist. There's no reason to expect any flow to bear a stable relationship to any other flow. Flows go up and down, on and off, in all sorts of combinations, in response to stocks, not to other flows.

Let me use a simple example to explain what I mean. Suppose you knew nothing at all about thermostats, but you had a lot of data about past heat flows into and out of the room. You could find an equation telling you how those flows have varied together in the past, because under ordinary circumstances, being governed by the same stock (temperature of the room), they do vary together.

Your equation would hold, however, only until something changes in the system's structure—someone opens a window or improves the insulation, or tunes the furnace, or forgets to order oil. You could predict tomorrow's room temperature with your equation, as long as the system didn't change or break down. But if you were asked to make the room warmer, or if the room temperature suddenly started plummeting and you had to fix it, or if you wanted to produce the same room temperature with a lower fuel bill, your behavior-level analysis wouldn't help you. You would have to dig into the system's structure.

That's why behavior-based econometric models are pretty good at predicting the near-term performance of the economy, quite bad at predicting the longer-term performance, and terrible at telling one how to improve the performance of the economy.

And that's one reason why systems of all kinds surprise us. We are too fascinated by the events they generate. We pay too little attention to their

history. And we are insufficiently skilled at seeing in their history clues to the structures from which behavior and events flow.

3 Linear Minds in a Nonlinear World

Linear relationships are easy to think about: the more the merrier. Linear equations are solvable, which makes them suitable for textbooks. Linear systems have an important modular virtue: you can take them apart and put them together again—the pieces add up.

Nonlinear systems generally cannot be solved and cannot be added together. . . . Nonlinearity means that the act of playing the game has a way of changing the rules. . . . That twisted changeability makes nonlinearity hard to calculate, but it also creates rich kinds of behavior that never occur in linear systems.

—James Gleick, author of *Chaos: Making a New Science*³

We often are not very skilled in understanding the nature of relationships. A **linear relationship** between two elements in a system can be drawn on a graph with a straight line. It's a relationship with constant proportions. If I put 10 pounds of fertilizer on my field, my yield will go up by 2 bushels. If I put on 20 pounds, my yield will go up by 4 bushels. If I put on 30 pounds, I'll get an increase of 6 bushels.

A **nonlinear relationship** is one in which the cause does not produce a proportional effect. The relationship between cause and effect can only be drawn with curves or wiggles, not with a straight line. If I put 100 pounds of fertilizer on, my yield will go up by 10 bushels; if I put on 200, my yield will not go up at all; if I put on 300, my yield will go down. Why? I've damaged my soil with "too much of a good thing."

The world is full of nonlinearities.

So the world often surprises our linear-thinking minds. If we've learned that a small push produces a small response, we think that twice as big a push will produce twice as big a response. But in a nonlinear system, twice the push could produce one-sixth the response, or the response squared, or no response at all.

Here are some examples of nonlinearities:

- As the flow of traffic on a highway increases, car speed is affected only slightly over a large range of car density. Eventually, however, small further increases in density produce a rapid drop-off in speed. And when the number of cars on the highway builds up to a certain point, it can result in a traffic jam, and car speed drops to zero.
- Soil erosion can proceed for a long time without much affect on crop yield—until the topsoil is worn down to the depth of the root zone of the crop. Beyond that point, a little further erosion can cause yields to plummet.
- A little tasteful advertising can awaken interest in a product. A lot of blatant advertising can cause disgust for the product.

You can see why nonlinearities produce surprises. They foil the reasonable expectation that if a little of some cure did a little good, then a lot of it will do a lot of good—or alternately that if a little destructive action caused only a tolerable amount of harm, then more of that same kind of destruction will cause only a bit more harm. Reasonable expectations like these in a nonlinear world produce classic mistakes.

Nonlinearities are important not only because they confound our expectations about the relationship between action and response. They are even more important because they *change the relative strengths of feedback loops*. They can flip a system from one mode of behavior to another.

Nonlinearities are the chief cause of the shifting dominance that characterizes several of the systems in the zoo—the sudden swing between exponential growth caused by a dominant reinforcing loop, say, and then decline caused by a suddenly dominant balancing loop.

To take a dramatic example of the effects of nonlinearities, consider the destructive irruptions of the spruce budworm in North American forests.

INTERLUDE • *Spruce Budworms, Firs, and Pesticides*

Tree ring records show that the spruce budworm has been killing spruce and fir trees periodically in North America for at least 400 years. Until this century, no one much cared. The valuable tree for the lumber industry was the white pine. Spruce and fir were considered “weed species.” Eventually,

however, the stands of virgin pine were gone, and the lumber industry turned to spruce and fir. Suddenly the budworm was seen as a serious pest.

So, beginning in the 1950s, northern forests were sprayed with DDT to control the spruce budworm. In spite of the spraying, every year there was a budworm resurgence. Annual sprays were continued through the 1950s, 1960s, and 1970s, until DDT was banned. Then the sprays were changed to fenitrothion, acephate, Sevin, and methoxychlor.

Insecticides were no longer thought to be the ultimate answer to the budworm problem, but they were still seen as essential. “Insecticides buy time,” said one forester, “That’s all the forest manager wants; to preserve the trees until the mill is ready for them.”

By 1980, spraying costs were getting unmanageable—the Canadian province of New Brunswick spent \$12.5 million on budworm “control” that year. Concerned citizens were objecting to the drenching of the landscape with poisons. And, in spite of the sprays, the budworm was still killing as many as 20 million hectares (50 million acres) of trees per year.

C. S. Holling of the University of British Columbia and Gordon Baskerville of the University of New Brunswick put together a computer model to get a whole-system look at the budworm problem. They discovered that before the spraying began, the budworm had been barely detectable in most years. It was controlled by a number of predators, including birds, a spider, a parasitic wasp, and several diseases. Every few decades, however, there was a budworm outbreak, lasting from six to ten years. Then the budworm population would subside, eventually to explode again.

The budworm preferentially attacks balsam fir, secondarily spruce. Balsam fir is the most competitive tree in the northern forest. Left to its own devices, it would crowd out spruce and birch, and the forest would become a monoculture of nothing but fir. Each budworm outbreak cuts back the fir population, opening the forest for spruce and birch. Eventually fir moves back in.

As the fir population builds up, the probability of an outbreak increases—*nonlinearly*. The reproductive potential of the budworm increases more than proportionately to the availability of its favorite food supply. The final trigger is two or three warm, dry springs, perfect for the survival of budworm larvae. (If you’re doing event-level analysis, you will blame the outburst on the warm, dry springs.)

The budworm population grows too great for its natural enemies to hold in check—*nonlinearly*. Over a wide range of conditions, greater budworm populations result in more rapid multiplication of budworm predators. But beyond some point, the predators can multiply no faster. What was a reinforcing relationship—more budworms, faster predator multiplication—becomes a nonrelationship—more budworms, no faster predator multiplication—and the budworms take off, unimpeded.

Now only one thing can stop the outbreak: the insect reducing its own food supply by killing off fir trees. When that finally happens, the budworm population crashes—*nonlinearly*. The reinforcing loop of budworm reproduction yields dominance to the balancing loop of budworm starvation. Spruce and birch move into the spaces where the firs used to be, and the cycle begins again.

The budworm/spruce/fir system oscillates over decades, but it is ecologically stable within bounds. It can go on forever. The main effect of the budworm is to allow tree species other than fir to persist. But in this case what is ecologically stable is economically unstable. In eastern Canada, the

Many relationships in systems are nonlinear. Their relative strengths shift in disproportionate amounts as the stocks in the system shift. Nonlinearities in feedback systems produce shifting dominance of loops and many complexities in system behavior.

economy is almost completely dependent on the logging industry, which is dependent on a steady supply of fir and spruce.

When industry sprays insecticides, it shifts the whole system to balance uneasily on different points within its nonlinear relationships. It kills off not only the pest, but the natural enemies of the pest, thereby weakening the feedback loop that normally keeps the budworms in check. It keeps the density of fir

high, moving the budworms up their nonlinear reproduction curve to the point at which they're perpetually on the edge of population explosion.

The forest management practices have set up what Holling calls “persistent semi-outbreak conditions” over larger and larger areas. The managers have found themselves locked into a policy in which there is an incipient volcano bubbling, such that, if the policy fails, there will be an outbreak of an intensity that has never been seen before.”⁴

4 Nonexistent Boundaries

When we think in terms of systems, we see that a fundamental misconception is embedded in the popular term “side-effects.” . . . This phrase means roughly “effects which I hadn’t foreseen or don’t want to think about.” . . . Side-effects no more deserve the adjective “side” than does the “principal” effect. It is hard to think in terms of systems, and we eagerly warp our language to protect ourselves from the necessity of doing so.

—Garrett Hardin,⁵ ecologist

Remember the clouds in the structural diagrams of Chapters One and Two? Beware of clouds! They are prime sources of system surprises.

Clouds stand for the beginnings and ends of flows. They are stocks—sources and sinks—that are being ignored at the moment for the purposes of simplifying the present discussion. They mark the boundary of the system diagram. They rarely mark a real boundary, because systems rarely have real boundaries. Everything, as they say, is connected to everything else, and not neatly. There is no clearly determinable boundary between the sea and the land, between sociology and anthropology, between an automobile’s exhaust and your nose. There are only boundaries of word, thought, perception, and social agreement—artificial, mental-model boundaries.

The greatest complexities arise exactly at boundaries. There are Czechs on the German side of the border and Germans on the Czech side of the border. Forest species extend beyond the edge of the forest into the field; field species penetrate partway into the forest. Disorderly, mixed-up borders are sources of diversity and creativity.

In our system zoo, for instance, I showed the flow of cars into a car dealer’s inventory as coming from a cloud. Of course, cars don’t come from a cloud, they come from the transformation of a stock of raw materials, with the help of capital, labor, energy, technology, and management (the means of production). Similarly, the flow of cars out of the inventory goes not to a cloud, but through sales to the households or businesses of consumers.

Whether it is important to keep track of raw materials or consumers’ home stocks (whether it is legitimate to replace them in a diagram with clouds) depends on whether these stocks are likely to have a significant influence on




Figure 47. Revealing some of the stocks behind the clouds.

the behavior of the system over the time period of interest. If raw materials are guaranteed to be abundant and consumers continue to demand the products, then clouds will do. But if there could be a materials shortage or a product glut, and if we drew a mental boundary around the system that did not include these stocks, then we could be surprised by future events.

There are still clouds in Figure 47. The boundary can be expanded further. Processed raw materials come from chemical plants, smelters, or refineries, whose input comes, ultimately, from the earth. Processing creates not only products, but also employment, wages, profits, and pollution. Discarded consumers' stocks go to landfills or incinerators or recycling centers, from which they go on to have further effects on society and the environment. Landfills leach into drinking-water wells, incinerators produce smoke and ash, recycling centers move materials back into the production stream.

Whether it's important to think about the full flow from mine to dump, or as industry calls it, "from cradle to grave," depends on who wants to know, for what purpose, over how long. In the long term, the full flow is important and, as the physical economy grows and society's "ecological footprint" expands, the long term is increasingly coming to be the short term. Landfills fill up with a suddenness that has been surprising for people whose mental models picture garbage as going "away," into some sort of a cloud. Sources of raw materials—mines, wells, and oil fields—can be exhausted with surprising suddenness too.

With a long enough time horizon, even mines and dumps are not the end of the story. The great geological cycles of the earth keep moving materials around, opening and closing seas, raising up and wearing down mountains. Eons from now, everything put in a dump will end up on the top of a mountain or deep under the sea. New deposits of metals and fuels will form. On planet Earth there are no system "clouds," no ultimate boundaries. Even real clouds in the sky are part of a hydrological cycle. Everything physical comes from somewhere, everything goes somewhere, everything keeps moving.

Which is not to say that every model, mental or computer, has to follow each connection until it includes the whole planet. Clouds are a necessary part of models that describe metaphysical flows. Anger literally “comes out of a cloud,” as does love, hatred, self esteem, and so on. If we’re to understand anything, we have to simplify, which means we have to make boundaries. Often that’s a safe thing to do. It’s usually not a problem, for example, to think of populations with births and deaths coming from and going to clouds, as in Figure 48.



Figure 48. More clouds.

Figure 48 shows actual “cradle to grave” boundaries. Even these boundaries would be unserviceable, however, if the population in question experienced significant in- or out-migration, or if the problem under discussion was limited cemetery space.

The lesson of boundaries is hard even for systems thinkers to get. There is no single, legitimate boundary to draw around a system. We have to invent boundaries for clarity and sanity; and boundaries can produce problems when we forget that we’ve artificially created them.

When you draw boundaries too narrowly, the system surprises you. For example, if you try to deal with urban traffic problems without thinking about settlement patterns, you build highways, which attract housing developments along their whole length. Those households, in turn, put more cars on the highways, which then become just as clogged as before.

If you try to solve a sewage problem by throwing the waste into a river, the towns downstream make it clear that the boundary for thinking about sewage has to include the whole river. It might also have to include the soil

There are no separate systems. The world is a continuum. Where to draw a boundary around a system depends on the purpose of the discussion—the questions we want to ask.

and groundwater surrounding the river. It probably doesn't have to include the next watershed or the planetary hydrological cycle.

Planning for a national park used to stop at the physical boundary of the park. But park boundaries around the world are regularly crossed by nomadic peoples, by migrating wildlife, by waters that flow into, out of, or under the park, by the effects of economic development at the park's edges, by acid rain, and now by a climate changing from greenhouse gases in the atmosphere. Even without climate change, to manage a park you have to think about a boundary wider than the official perimeter.

Systems analysts often fall into the opposite trap: making boundaries too large. They have a habit of producing diagrams that cover several pages with small print and many arrows connecting everything with everything. *There is the system!* they say. If you have considered anything less, you are academically illegitimate.

This "my model is bigger than your model" game results in enormously complicated analyses, which produce piles of information that may only serve to obscure the answers to the questions at hand. For example, modeling the earth's climate in full detail is interesting for many reasons, but may not be necessary for figuring out how to reduce a country's CO₂ emissions to reduce climate change.

The right boundary for thinking about a problem rarely coincides with the boundary of an academic discipline, or with a political boundary. Rivers make handy borders between countries, but the worst possible borders for managing the quantity and quality of the water. Air is worse than water in its insistence on crossing political borders. National boundaries mean nothing when it comes to ozone depletion in the stratosphere, or greenhouse gases in the atmosphere, or ocean dumping.

Ideally, we would have the mental flexibility to find the appropriate boundary for thinking about each new problem. We are rarely that flexible. We get attached to the boundaries our minds happen to be accustomed to. Think how many arguments have to do with boundaries—national boundaries, trade boundaries, ethnic boundaries, boundaries between public and private responsibility, and boundaries between the rich and the poor, polluters and pollutees, people alive now and people who will come in the future. Universities can maintain disputes for years about the boundaries between economics and government, art and art history, literature and literary criticism. Too often, universities are living monuments to boundary rigidity.




Figure 49. Examples of more clouds. These are systems in which a boundary or cloud should not stop you from thinking beyond the borders of the system, but start you thinking beyond those borders. What is driving the supply of people being given new sentences? Where do the fuel rods go after replacement? What happens to an unemployed person whose registration for unemployment lapses?

It's a great art to remember that *boundaries are of our own making, and that they can and should be reconsidered for each new discussion, problem, or purpose.* It's a challenge to stay creative enough to drop the boundaries that worked for the last problem and to find the most appropriate set of boundaries for the next question. It's also a necessity, if problems are to be solved well.

Layers of Limits

Systems surprise us because our minds like to think about single causes neatly producing single effects. We like to think about one or at most a few things at a time. And we don't like, especially when our own plans and desires are involved, to think about limits.

But we live in a world in which many causes routinely come together to produce many effects. Multiple inputs produce multiple outputs, and virtually all of the inputs, and therefore outputs, are limited. For example, an industrial manufacturing process needs:

- capital
- labor
- energy
- raw materials
- land
- water
- technology
- credit
- insurance
- customers
- good management
- public-funded infrastructure and government services (such as police and fire protection and education for managers and workers)
- functioning families to bring up and care for both producers and consumers
- a healthy ecosystem to supply or support all these inputs and to absorb or carry away their wastes

A patch of growing grain needs:

- sunlight
- air
- water
- nitrogen
- phosphorus

- potassium
- dozens of minor nutrients
- a friable soil and the services of a microbial soil community
- some system to control weeds and pests
- protection from the wastes of the industrial manufacturer

It was with regard to grain that Justus von Liebig came up with his famous “law of the minimum.” It doesn’t matter how much nitrogen is available to the grain, he said, if what’s short is phosphorus. It does no good to pour on more phosphorus, if the problem is low potassium.

Bread will not rise without yeast, no matter how much flour it has. Children will not thrive without protein, no matter how many carbohydrates they eat. Companies can’t keep going without energy, no matter how many customers they have—or without customers, no matter how much energy they have.

This concept of a **limiting factor** is simple and widely misunderstood. Agronomists assume, for example, that they know what to put in artificial fertilizer, because they have identified many of the major and minor nutrients in good soil. Are there any essential nutrients they have not identified? How do artificial fertilizers affect soil microbe communities? Do they interfere with, and therefore limit, any other functions of good soil? And what limits the production of artificial fertilizers?

At any given time, the input that is most important to a system is the one that is most limiting.

Rich countries transfer capital or technology to poor ones and wonder why the economies of the receiving countries still don’t develop, never thinking that capital or technology may not be the most limiting factors.

Economics evolved in a time when labor and capital were the most common limiting factors to production. Therefore, most economic production functions keep track only of these two factors (and sometimes technology). As the economy grows relative to the ecosystem, however, and the limiting factors shift to clean water, clean air, dump space, and acceptable forms of energy and raw materials, the traditional focus on only capital and labor becomes increasingly unhelpful.

One of the classic models taught to systems students at MIT is Jay Forrester’s corporate-growth model. It starts with a successful young company, growing rapidly. The problem for this company is to recognize

and deal with its shifting limits—limits that change in response to the company's own growth.

The company may hire salespeople, for example, who are so good that they generate orders faster than the factory can produce. Delivery delays increase and customers are lost, because production capacity is the most limiting factor. So the managers expand the capital stock of production plants. New people are hired in a hurry and trained too little. Quality suffers and customers are lost because labor skill is the most limiting factor. So management invests in worker training. Quality improves, new orders pour in, and the order-fulfillment and record-keeping system clogs. And so forth.

There are layers of limits around every growing plant, child, epidemic, new product, technological advance, company, city, economy, and population. Insight comes not only from recognizing which factor is limiting, but from seeing that *growth itself depletes or enhances limits* and therefore changes what is limiting. The interplay between a growing plant and the soil, a growing company and its market, a growing economy and its resource base, is dynamic. Whenever one factor ceases to be limiting, growth occurs, and the growth itself changes the relative scarcity of factors until another becomes limiting. To shift attention from the abundant factors to the next potential limiting factor is to gain real understanding of, and control over, the growth process.

Any physical entity with multiple inputs and outputs—a population, a production process, an economy—is surrounded by layers of limits. As the system develops, it interacts with and affects its own limits. The growing entity and its limited environment together form a coevolving dynamic system.

Understanding layers of limits and keeping an eye on the next upcoming limiting factor is not a recipe for perpetual growth, however. For any physical entity in a finite environment, perpetual growth is impossible. Ultimately, the choice is not to grow forever but to decide what limits to live within. If a company produces a perfect product or service at an affordable price, it

Any physical entity with multiple inputs and outputs is surrounded by layers of limits.

will be swamped with orders until it grows to the point at which some limit decreases the perfection of the product or raises its price. If a city meets the needs of all its inhabitants better than any other city, people will flock there until some limit brings down the city's ability to satisfy peoples' needs.⁶

~~There always will be limits to growth. They can be self-imposed. If they aren't, they will be system-imposed. No physical entity can grow forever. If company managers, city governments, the human population do not choose and enforce their own limits to keep growth within the capacity of the supporting environment, then the environment will choose and enforce limits.~~

There always will be limits to growth. They can be self-imposed. If they aren't, they will be system-imposed.

5 Ubiquitous Delays

I realize with fright that my impatience for the re-establishment of democracy had something almost communist in it; or, more generally, something rationalist. I had wanted to make history move ahead in the same way that a child pulls on a plant to make it grow more quickly.

I believe we must learn to wait as we learn to create. We have to patiently sow the seeds, assiduously water the earth where they are sown and give the plants the time that is their own. One cannot fool a plant any more than one can fool history.

—Václav Havel,⁷ playwright, last President of Czechoslovakia
and first president of the Czech Republic

It takes time for a plant or a forest or a democracy to grow; time for letters put into a mailbox to reach their destinations; time for consumers to absorb information about changing prices and alter their buying behavior, or for a nuclear power plant to be built, or a machine to wear out, or a new technology to penetrate an economy.

We are surprised over and over again at how much time things take. Jay Forrester used to tell us, when we were modeling a construction or processing delay, to ask everyone in the system how long they thought the delay was, make our best guess, and then multiply by three. (That correction factor also works perfectly, I have found, for estimating how long it will take to write a book!)

Delays are ubiquitous in systems. Every stock is a delay. Most flows have delays—shipping delays, perception delays, processing delays, maturation

delays. Here are just a few of the delays we have found important to include in various models we have made:

- The delay between catching an infectious disease and getting sick enough to be diagnosed—days to years, depending on the disease.
- The delay between pollution emission and the diffusion or percolation or concentration of the pollutant in the ecosystem to the point at which it does harm.
- The gestation and maturation delay in building up breeding populations of animals or plants, causing the characteristic oscillations of commodity prices: 4-year cycles for pigs, 7 years for cows, 11 years for cocoa trees.⁸
- The delay in changing the social norms for desirable family size—at least one generation.
- The delay in retooling a production stream and the delay in turning over a capital stock. It takes 3 to 8 years to design a new car and bring it to the market. That model may have 5 years of life on the new-car market. Cars stay on the road an average of 10 to 15 years.

Just as the appropriate boundaries to draw around one's picture of a system depend on the purpose of the discussion, so do the important delays. If you're worrying about oscillations that take weeks, you probably don't have to think about delays that take minutes, or years. If you're concerned about the decades-long development of a population and economy, you usually can ignore oscillations that take weeks. The world peeps, squawks, bangs, and thunders at many frequencies all at once. What is a significant delay depends—usually—on which set of frequencies you're trying to understand.

The systems zoo has already demonstrated how important delays in feedback are to the behavior of systems. Changing the length of a delay may utterly change behavior. Delays are often sensitive leverage points for policy, if they can be made shorter or longer. You can see why that is. If a decision point in a system (or a person working in that part of the system) is responding to delayed information, or responding with a delay, the decisions will be off target. Actions will be too much or too little to achieve the

decision maker's goals. On the other hand, if action is taken too fast, it may nervously amplify short-term variation and create unnecessary instability. Delays determine how fast systems can react, how accurately they hit their targets, and how timely is the information passed around a system. Overshoots, oscillations, and collapses are always caused by delays.

Understanding delays helps one understand why Mikhail Gorbachev could transform the information system of the Soviet Union virtually overnight, but not the physical economy. (That takes decades.) It helps one see why the absorption of East Germany by West Germany produced more hardship over a longer time than the politicians foresaw. Because of long delays in building new power plants, the electricity industry is plagued with cycles of overcapacity and then undercapacity leading to brownouts. Because of decades-long delays as the earth's oceans respond to warmer temperatures, human fossil-fuel emissions have already induced changes in climate that will not be fully revealed for a generation or two.

When there are long delays in feedback loops, some sort of foresight is essential. To act only when a problem becomes obvious is to miss an important opportunity to solve the problem.

6 Bounded Rationality

As every individual, therefore, endeavours as much as he can both to employ his capital in the support of domestic industry, and so to direct that industry that its produce may be of greatest value. . . he generally, indeed, neither intends to promote the public interest, nor knows how much he is promoting it. . . . He intends his own security; . . . he intends only his own gain and he is in this . . . led by an invisible hand to promote an end which was no part of his intention. By pursuing his own interest he frequently promotes that of society more effectually than when he really intends to promote it.

—Adam Smith,⁹ 18th century political economist

It would be so nice if the “invisible hand” of the market really did lead individuals to make decisions that add up to the good of the whole. Then not only would material selfishness be a social virtue, but mathematical

models of the economy would be much easier to make. There would be no need to think about the good of other people or about the operations of complex feedback systems. No wonder Adam Smith's model has had such strong appeal for two hundred years!

Unfortunately, the world presents us with multiple examples of people acting rationally in their short-term best interests and producing aggregate results that no one likes. Tourists flock to places like Waikiki or Zermatt and then complain that those places have been ruined by all the tourists. Farmers produce surpluses of wheat, butter, or cheese, and prices plummet. Fishermen overfish and destroy their own livelihood. Corporations collectively make investment decisions that cause business-cycle downturns. Poor people have more babies than they can support.

Why?

Because of what World Bank economist Herman Daly calls the "invisible foot" or what Nobel Prize-winning economist Herbert Simon calls **bounded rationality**.¹⁰

Bounded rationality means that people make quite reasonable decisions based on the information they have. But they don't have perfect information, especially about more distant parts of the system. Fishermen don't know how many fish there are, much less how many fish will be caught by other fishermen that same day.

Businessmen don't know for sure what other businessmen are planning to invest, or what consumers will be willing to buy, or how their products will compete. They don't know their current market share, and they don't know the size of the market. Their information about these things is incomplete and delayed, and their own responses are delayed. So they systematically under- and overinvest.

We are not omniscient, rational optimizers, says Simon. Rather, we are blundering "satisficers," attempting to meet (*satisfy*) our needs well enough (*sufficiently*) before moving on to the next decision.¹¹ We do our best to further our own nearby interests in a rational way, but we can take into account only what we know. We don't know what others are planning to do, until they do it. We rarely see the full range of possibilities before us. We often don't foresee (or choose to ignore) the impacts of our actions on the whole system. So instead of finding a long-term optimum, we discover within our limited purview a choice we can live with for now, and we stick to it, changing our behavior only when forced to.

We don't even interpret perfectly the imperfect information that we do have, say behavioral scientists. We misperceive risk, assuming that some things are much more dangerous than they really are and others much less. We live in an exaggerated present—we pay too much attention to recent experience and too little attention to the past, focusing on current events rather than long-term behavior. We discount the future at rates that make no economic or ecological sense. We don't give all incoming signals their appropriate weights. We don't let in at all news we don't like, or information that doesn't fit our mental models. Which is to say, we don't even make decisions that optimize our own individual good, much less the good of the system as a whole.

When the theory of bounded rationality challenged two hundred years of economics based on the teachings of political economist Adam Smith, you can imagine the controversy that resulted—one that is far from over. Economic theory as derived from Adam Smith assumes first that *homo economicus* acts with perfect optimality on complete information, and second that when many of the species *homo economicus* do that, their actions add up to the best possible outcome for everybody.

Neither of these assumptions stands up long against the evidence. In the next chapter on system traps and opportunities, I will describe some of the most commonly encountered structures that can cause bounded rationality to lead to disaster. They include such familiar phenomena as addiction, policy resistance, arms races, drift to low performance, and the tragedy of the commons. For now, I want to make just one point about the biggest surprise that comes from not understanding bounded rationality.

Suppose you are for some reason lifted out of your accustomed place in society and put in the place of someone whose behavior you have never understood. Having been a staunch critic of government, you suddenly become part of government. Or having been a laborer in opposition to management, you become management (or vice versa). Perhaps having been an environmental critic of big business, you find yourself making environmental decisions for big business. Would that such transitions could happen much more often, in all directions, to broaden everyone's horizons!

In your new position, you experience the information flows, the incentives and disincentives, the goals and discrepancies, the pressures—the bounded rationality—that goes with that position. It's possible that you

could retain your memory of how things look from another angle, and that you burst forth with innovations that transform the system, but it's distinctly unlikely. If you become a manager, you probably will stop seeing labor as a deserving partner in production, and start seeing it as a cost to be minimized. If you become a financier, you probably will overinvest during booms and underinvest during busts, along with all the other financiers. If you become very poor, you will see the short-term rationality, the hope, the opportunity, the necessity of having many children. If you are now a fisherman with a mortgage on your boat, a family to support, and imperfect knowledge of the state of the fish population, you will overfish.

We teach this point by playing games in which students are put into situations in which they experience the realistic, partial information streams seen by various actors in real systems. As simulated fishermen, they overfish. As ministers of simulated developing nations, they favor the needs of their industries over the needs of their people. As the upper class, they feather their own nests; as the lower class, they become apathetic or rebellious. So would you. In the famous Stanford prison experiment by psychologist Philip Zimbardo, players even took on, in an amazingly short time, the attitudes and behaviors of prison guards and prisoners.¹²

Seeing how individual decisions are rational within the bounds of the information available does not provide an excuse for narrow-minded behavior. It provides an understanding of why that behavior arises. Within the bounds of what a person in that part of the system can see and know, the behavior is reasonable. Taking out one individual from a position of bounded rationality and putting in another person is not likely to make much difference. Blaming the individual rarely helps create a more desirable outcome.

Change comes first from stepping outside the limited information that can be seen from any single place in the system and getting an overview. From a wider perspective, information flows, goals, incentives, and disincentives can be restructured so that separate, bounded, rational actions do add up to results that everyone desires.

It's amazing how quickly and easily behavior changes can come, with even slight enlargement of bounded rationality, by providing better, more complete, timelier information.

INTERLUDE • *Electric Meters in Dutch Houses*

Near Amsterdam, there is a suburb of single-family houses all built at the same time, all alike. Well, nearly alike. For unknown reasons it happened that some of the houses were built with the electric meter down in the basement. In other houses, the electric meter was installed in the front hall.

These were the sort of electric meters that have a glass bubble with a small horizontal metal wheel inside. As the household uses more electricity, the wheel turns faster and a dial adds up the accumulated kilowatt-hours.

During the oil embargo and energy crisis of the early 1970s, the Dutch began to pay close attention to their energy use. It was discovered that some of the houses in this subdivision used one-third less electricity than the other houses. No one could explain this. All houses were charged the same price for electricity, all contained similar families.

The difference, it turned out, was in the position of the electric meter. The families with high electricity use were the ones with the meter in the basement, where people rarely saw it. The ones with low use had the meter in the front hall where people passed, the little wheel turning around, adding up the monthly electricity bill many times a day.¹³

Some systems are structured to function well despite bounded rationality. The right feedback gets to the right place at the right time. Under ordinary circumstances, your liver gets just the information it needs to do its job. In undisturbed ecosystems and traditional cultures, the average individual, species, or population, left to its own devices, behaves in ways that serve and stabilize the whole. These systems and others are self-regulatory. They do not cause problems. We don't have government agencies and dozens of failed policies about them.

Since Adam Smith, it has been widely believed that the free, competitive market is one of these properly structured self-regulating systems. In some ways, it is. In other ways, obvious to anyone who is willing to look, it isn't. A free market does allow producers and consumers, who have the best information about production opportunities and consumption choices, to make fairly uninhibited and locally rational decisions. But those decisions can't, by themselves, correct the overall system's tendency to create monopolies and undesirable side effects (externalities), to discriminate against the poor, or to overshoot its sustainable carrying capacity.

To paraphrase a common prayer: God grant us the serenity to exercise our bounded rationality freely in the systems that are structured appropriately, the courage to restructure the systems that aren't, and the wisdom to know the difference!

The bounded rationality of each actor in a system—determined by the information, incentives, disincentives, goals, stresses, and constraints impinging on that actor—may or may not lead to decisions that further the welfare of the system as a whole. If they do not, putting new actors into the same system will not improve the system's performance. What makes a difference is redesigning the system to improve the information, incentives, disincentives, goals, stresses, and constraints that have an effect on specific actors.

The bounded rationality of each actor in a system may not lead to decisions that further the welfare of the system as a whole.
