— THREE —

Why Systems Work So Well

> If the land mechanism as a whole is good, then every part is good, whether we understand it or not. If the biota, in the course of aeons, has built something we like but do not understand, then who but a fool would discard seemingly useless parts? To keep every cog and wheel is the first precaution of intelligent tinkering.
>
> —Aldo Leopold,[1] forester

Chapter Two introduced simple systems that create their own behavior based on their structures. Some are quite elegant—surviving the buffeting of the world—and, within limits, regaining their composure and proceeding on about their business of maintaining a room's temperature, depleting an oil field, or bringing into balance the size of a fishing fleet with the productivity of a fishery resource.

If pushed too far, systems may well fall apart or exhibit heretofore unobserved behavior. But, by and large, they manage quite well. And that is the beauty of systems: They can work so well. When systems work well, we see a kind of harmony in their functioning. Think of a community kicking in to high gear to respond to a storm. People work long hours to help victims, talents and skills emerge; once the emergency is over, life goes back to "normal."

Why do systems work so well? Consider the properties of highly functional systems—machines or human communities or ecosystems—which are familiar to you. Chances are good that you may have observed one of three characteristics: resilience, self-organization, or hierarchy.

Resilience

> Placing a system in a straitjacket of constancy can cause fragility to evolve.
>
> —C. S. Holling,[2] ecologist

Resilience has many definitions, depending on the branch of engineering, ecology, or system science doing the defining. For our purposes, the normal dictionary meaning will do: "the ability to bounce or spring back into shape, position, etc., after being pressed or stretched. Elasticity. The ability to recover strength, spirits, good humor, or any other aspect quickly." Resilience is a measure of a system's ability to survive and persist within a variable environment. The opposite of resilience is brittleness or rigidity.

Resilience arises from a rich structure of many feedback loops that can work in different ways to restore a system even after a large perturbation. A single balancing loop brings a system stock back to its desired state. Resilience is provided by several such loops, operating through different mechanisms, at different time scales, and with redundancy—one kicking in if another one fails.

A set of feedback loops that can *restore or rebuild feedback loops* is resilience at a still higher level—meta-resilience, if you will. Even higher meta-meta-resilience comes from feedback loops that can *learn*, *create*, *design*, *and evolve* ever more complex restorative structures. Systems that can do this are self-organizing, which will be the next surprising system characteristic I come to.

The human body is an astonishing example of a resilient system. It can fend off thousands of different kinds of invaders, it can tolerate wide ranges of temperature and wide variations in food supply, it can reallocate blood supply, repair rips, gear up or slow down metabolism, and compensate to some extent for missing or defective parts. Add to it a self-organizing intelligence that can learn, socialize, design technologies, and even transplant body parts, and you have a formidably resilient system—although not infinitely so, because, so far at least, no human body-plus-intelligence has been resilient enough to keep itself or any other body from eventually dying.

There are always limits to resilience.

Ecosystems are also remarkably resilient, with multiple species holding each other in check, moving around in space, multiplying or declining over time in response to weather and the availability of nutrients and the impacts of human activities. Populations and ecosystems also have the ability to "learn" and evolve through their incredibly rich genetic variability. They can, given enough time, come up with whole new systems to take advantage of changing opportunities for life support.

Resilience is not the same thing as being static or constant over time. Resilient systems can be very dynamic. Short-term oscillations, or periodic outbreaks, or long cycles of succession, climax, and collapse may in fact be the normal condition, which resilience acts to restore!

And, conversely, systems that are constant over time can be unresilient. This distinction between static stability and resilience is important. Static stability is something you can see; it's measured by variation in the condition of a system week by week or year by year. Resilience is something that may be very hard to see, unless you exceed its limits, overwhelm and damage the balancing loops, and the system structure breaks down. Because resilience may not be obvious without a whole-system view, people often sacrifice resilience for stability, or for productivity, or for some other more immediately recognizable system property.

- Injections of genetically engineered bovine growth hormone increase the milk production of a cow without proportionately increasing the cow's food intake. The hormone diverts some of the cow's metabolic energy from other bodily functions to milk production. (Cattle breeding over centuries has done much the same thing but not to the same degree.) The cost of increased production is lowered resilience. The cow is less healthy, less long-lived, more dependent on human management.
- Just-in-time deliveries of products to retailers or parts to manufacturers have reduced inventory instabilities and brought down costs in many industries. The just-in-time model also has made the production system more vulnerable, however, to perturbations in fuel supply, traffic flow, computer breakdown, labor availability, and other possible glitches.
- Hundreds of years of intensive management of the forests of Europe gradually have replaced native ecosystems with single-age, single-species plantations, often of nonnative trees. These

> forests are designed to yield wood and pulp at a high rate indefinitely. However, without multiple species interacting with each other and drawing and returning varying combinations of nutrients from the soil, these forests have lost their resilience. They seem to be especially vulnerable to a new form of insult: industrial air pollution.

Many chronic diseases, such as cancer and heart disease, come from breakdown of resilience mechanisms that repair DNA, keep blood vessels flexible, or control cell division. Ecological disasters in many places come from loss of resilience, as species are removed from ecosystems, soil chemistry and biology are disturbed, or toxins build up. Large organizations of all kinds, from corporations to governments, lose their resilience simply because the feedback mechanisms by which they sense and respond to their environment have to travel through too many layers of delay and distortion. (More on that in a minute, when we come to hierarchies.)

I think of resilience as a plateau upon which the system can play, performing its normal functions in safety. A resilient system has a big plateau, a lot of space over which it can wander, with gentle, elastic walls that will bounce it back, if it comes near a dangerous edge. As a system loses its resilience, its plateau shrinks, and its protective walls become lower and more rigid, until the system is operating on a knife-edge, likely to fall off in one direction or another whenever it makes a move. Loss of resilience can come as a surprise, because the system usually is paying much more attention to its play than to its playing space. One day it does something it has done a hundred times before and crashes.

Systems need to be managed not only for productivity or stability, they also need to be managed for resilience— the ability to recover from perturbation, the ability to restore or repair themselves.

Awareness of resilience enables one to see many ways to preserve or enhance a system's own restorative powers. That awareness is behind the encouragement of natural ecosystems on farms, so that predators can take on more of the job of controlling pests. It is behind "holistic" health care that tries not only to cure disease but also to build up a body's internal resistance. It is behind aid programs that do more than give food or money—that try to change the circumstances that obstruct peoples' ability to provide their own food or money.

Self-Organization

> [Evolution] appears to be not a series of accidents the course of which is determined only by the change of environments during earth history and the resulting struggle for existence, . . . but is governed by definite laws. . . . The discovery of these laws constitutes one of the most important tasks of the future.
>
> —Ludwig von Bertalanffy,[3] biologist

The most marvelous characteristic of some complex systems is their ability to learn, diversify, complexify, evolve. It is the ability of a single fertilized ovum to generate, out of itself, the incredible complexity of a mature frog, or chicken, or person. It is the ability of nature to have diversified millions of fantastic species out of a puddle of organic chemicals. It is the ability of a society to take the ideas of burning coal, making steam, pumping water, and specializing labor, and develop them eventually into an automobile assembly plant, a city of skyscrapers, a worldwide network of communications.

This capacity of a system to make its own structure more complex is called **self-organization**. You see self-organization in a small, mechanistic way whenever you see a snowflake, or ice feathers on a poorly insulated window, or a supersaturated solution suddenly forming a garden of crystals. You see self-organization in a more profound way whenever a seed sprouts, or a baby learns to speak, or a neighborhood decides to come together to oppose a toxic waste dump.

Self-organization is such a common property, particularly of living systems, that we take it for granted. If we didn't, we would be dazzled by the unfolding systems of our world. And if we weren't nearly blind to the property of self-organization, we would do better at encouraging, rather than destroying, the self-organizing capacities of the systems of which we are a part.

Like resilience, self-organization is often sacrificed for purposes of short-term productivity and stability. Productivity and stability are the usual excuses for turning creative human beings into mechanical adjuncts to production processes. Or for narrowing the genetic variability of crop plants. Or for establishing bureaucracies and theories of knowledge that treat people as if they were only numbers.

Self-organization produces heterogeneity and unpredictability. It is likely to come up with whole new structures, whole new ways of doing things. It requires freedom and experimentation, and a certain amount of disorder. These conditions that encourage self-organization often can be scary for individuals and threatening to power structures. As a consequence, education systems may restrict the creative powers of children instead of stimulating those powers. Economic policies may lean toward supporting established, powerful enterprises rather than upstart, new ones. And many governments prefer their people not to be too self-organizing.

Fortunately, self-organization is such a basic property of living systems that even the most overbearing power structure can never fully kill it, although in the name of law and order, self-organization can be suppressed for long, barren, cruel, boring periods.

Systems theorists used to think that self-organization was such a complex property of systems that it could never be understood. Computers were used to model mechanistic, "deterministic" systems, not evolutionary ones, because it was suspected, without much thought, that evolutionary systems were simply not understandable.

New discoveries, however, suggest that just a few simple organizing principles can lead to wildly diverse self-organizing structures. Imagine a triangle with three equal sides. Add to the middle of each side another equilateral triangle, one-third the size of the first one. Add to each of the new sides another triangle, one-third smaller. And so on. The result is called a Koch snowflake. (See Figure 46.) Its edge has tremendous length—but it can be contained within a circle. This structure is one simple example of fractal geometry—a realm of mathematics and art populated by elaborate shapes formed by relatively simple rules.

Similarly, the delicate, beautiful, intricate structure of a stylized fern can be generated by a computer with just a few simple fractal rules. The

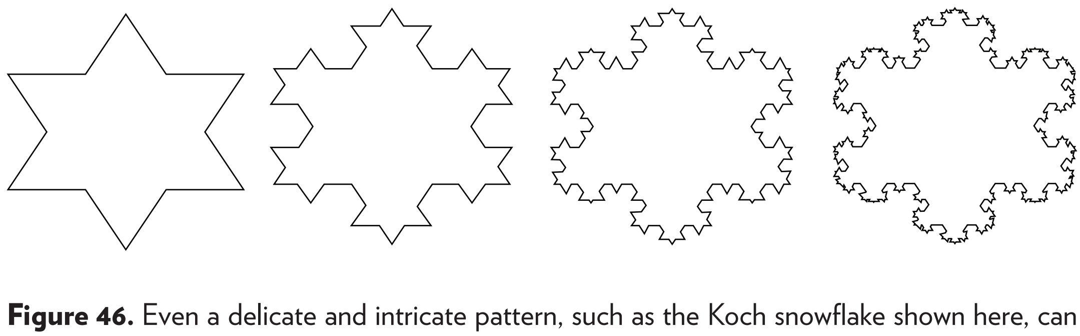

Figure 46. Even a delicate and intricate pattern, such as the Koch snowflake shown here, can evolve from a simple set of organizing principles or decision rules.

differentiation of a single cell into a human being probably proceeds by some similar set of geometric rules, basically simple, but generating utter complexity. (It is because of fractal geometry that the average human lung has enough surface area to cover a tennis court.)

Here are some other examples of simple organizing rules that have led to self-organizing systems of great complexity:

- All of life, from viruses to redwood trees, from amoebas to elephants, is based on the basic organizing rules encapsulated in the chemistry of DNA, RNA, and protein molecules.
- The agricultural revolution and all that followed started with the simple, shocking ideas that people could stay settled in one place, own land, select and cultivate crops.
- "God created the universe with the earth at its center, the land with the castle at its center, and humanity with the Church at its center"—the organizing principle for the elaborate social and physical structures of Europe in the Middle Ages.
- "God and morality are outmoded ideas; people should be objective and scientific, should own and multiply the means of production, and should treat people and nature as instrumental inputs to production"—the organizing principles of the Industrial Revolution.

Out of simple rules of self-organization can grow enormous, diversifying crystals of technology, physical structures, organizations, and cultures.

Science knows now that self-organizing systems can arise from simple rules. Science, itself a self-organizing system, likes to think that all the complexity of the world must arise, ultimately, from simple rules. Whether that actually happens is something that science does not yet know.

Systems often have the property of self-organization—the ability to structure themselves, to create new structure, to learn, diversify, and complexify. Even complex forms of self-organization may arise from relatively simple organizing rules—or may not.

Hierarchy

> So, naturalists observe, a flea
> Has smaller Fleas that on him prey;
> And these have smaller still to bite 'em,
> And so proceed *ad infinitum*.
>
> —Jonathan Swift,[4] 18th century poet

In the process of creating new structures and increasing complexity, one thing that a self-organizing system often generates is **hierarchy**.

The world, or at least the parts of it humans think they understand, is organized in subsystems aggregated into larger subsystems, aggregated into still larger subsystems. A cell in your liver is a subsystem of an organ, which is a subsystem of you as an organism, and you are a subsystem of a family, an athletic team, a musical group, and so forth. These groups are subsystems of a town or city, and then a nation, and then the whole global socioeconomic system that dwells within the biosphere system. This arrangement of systems and subsystems is called a hierarchy.

Corporate systems, military systems, ecological systems, economic systems, living organisms, are arranged in hierarchies. It is no accident that that is so. If subsystems can largely take care of themselves, regulate themselves, maintain themselves, and yet serve the needs of the larger system, while the larger system coordinates and enhances the functioning of the subsystems, a stable, resilient, and efficient structure results. It is hard to imagine how any other kind of arrangement could have come to be.

INTERLUDE • *Why the Universe Is Organized into Hierarchies—a Fable*

There once were two watchmakers, named Hora and Tempus. Both of them made fine watches, and they both had many customers. People dropped into their stores, and their phones rang constantly with new orders. Over the years, however, Hora prospered, while Tempus became poorer and poorer. That's because Hora discovered the principle of hierarchy. . . .

The watches made by both Hora and Tempus consisted of about one thousand parts each. Tempus put his together in such a way that if he had one partly assembled and had to put it down—to answer the phone, say—it fell to pieces. When he came back to it, Tempus would have to start all over again. The more his customers phoned him, the harder it became for him to find enough uninterrupted time to finish a watch.

Hora's watches were no less complex than those of Tempus, but he put together stable subassemblies of about ten elements each. Then he put ten of these subassemblies together into a larger assembly; and ten of those assemblies constituted the whole watch. Whenever Hora had to put down a partly completed watch to answer the phone, he lost only a small part of his work. So he made his watches much faster and more efficiently than did Tempus.

Complex systems can evolve from simple systems only if there are stable intermediate forms. The resulting complex forms will naturally be hierarchic. That may explain why hierarchies are so common in the systems nature presents to us. Among all possible complex forms, hierarchies are the only ones that have had the time to evolve.[5]

Hierarchies are brilliant systems inventions, not only because they give a system stability and resilience, but also because they reduce the amount of information that any part of the system has to keep track of.

In hierarchical systems relationships *within* each subsystem are denser and stronger than relationships *between* subsystems. Everything is still connected to everything else, but not equally strongly. People in the same university department talk to each other more than they talk to people in other departments. The cells that constitute the liver are in closer communication with each other than they are with the cells of the heart. If these differential information links within and between each level of the hierarchy are designed right, feedback delays are minimized. No level is overwhelmed with information. The system works with efficiency and resilience.

Hierarchical systems are partially decomposable. They can be taken apart and the subsystems with their especially dense information links can function, at least partially, as systems in their own right. When hierarchies break down, they usually split along their subsystem boundaries. Much can be learned by taking apart systems at different hierarchical levels—cells or organs, for example—and studying them separately. Hence, systems thinkers would say, the reductionist dissection of regular science teaches us a lot. However, one should not lose sight of the important relationships that bind each subsystem to the others and to the higher levels of the hierarchy, or one will be in for surprises.

If you have a liver disease, for example, a doctor usually can treat it without paying much attention to your heart or your tonsils (to stay on the same hierarchical level) or your personality (to move up a level or two) or the DNA in the nuclei of the liver cells (to move down several levels). There are just enough exceptions to that rule, however, to reinforce the necessity of stepping back to consider the whole hierarchy. Maybe your job exposes you to a chemical that is damaging your liver. Maybe the disease originates in a malfunction of the DNA.

What you need to think about may change over time, as self-organizing systems evolve new degrees of hierarchy and integration. The energy systems of nations were once almost completely decomposable one from another. That is no longer true. People whose thinking has not evolved as fast as the energy economy has may be shocked to discover how dependent they have become on resources and decisions halfway around the world.

You can watch self-organizing systems form hierarchies. A self-employed person gets too much work and hires some helpers. A small, informal nonprofit organization attracts many members and a bigger budget and one day the members decide, "Hey, we need someone to organize all this." A cluster of dividing cells differentiates into special functions and generates a branching circulatory system to feed all cells, and a branching nervous system to coordinate them.

Hierarchies evolve from the lowest level up—from the pieces to the whole, from cell to organ to organism, from individual to team, from actual production to management of production. Early farmers decided to come together and form cities for self-protection and for making trade more efficient. Life started with single-cell bacteria, not with elephants. The original purpose of a hierarchy is always to help its originating subsystems do their jobs better. This is something, unfortunately, that both the higher and the lower levels of a greatly articulated hierarchy easily can forget. Therefore, many systems are not meeting our goals because of malfunctioning hierarchies.

If a team member is more interested in personal glory than in the team winning, he or she can cause the team to lose. If a body cell breaks free from its hierarchical function and starts multiplying wildly, we call it a cancer. If students think their purpose is to maximize personal grades instead of seeking knowledge, cheating and other counterproductive behaviors break out. If a single corporation bribes the government to favor that corporation, the advantages of the competitive market and the good of the whole society are eroded.

When a subsystem's goals dominate at the expense of the total system's goals, the resulting behavior is called **suboptimization**.

Just as damaging as suboptimization, of course, is the problem of too much central control. If the brain controlled each cell so tightly that the cell could not perform its self-maintenance functions, the whole organism could die. If central rules and regulations prevent students or faculty from exploring fields of knowledge freely, the purpose of the university is not served. The coach of a team might interfere with the on-the-spot perceptions of a good player, to the detriment of the team. Economic examples of overcontrol from the top, from companies to nations, are the causes of some of the great catastrophes of history, all of which are by no means behind us.

To be a highly functional system, hierarchy must balance the welfare, freedoms, and responsibilities of the subsystems and total system—there must be enough central control to achieve coordination toward the large-system goal, and enough autonomy to keep all subsystems flourishing, functioning, and self-organizing.

Resilience, self-organization, and hierarchy are three of the reasons dynamic systems can work so well. Promoting or managing for these properties of a system can improve its ability to function well over the long term—to be sustainable. But watching how systems behave also can be full of surprises.

Hierarchical systems evolve from the bottom up. The purpose of the upper layers of the hierarchy is to serve the purposes of the lower layers.

— FOUR —

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.

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.

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.

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 understanding. 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.

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: