The Cultural Dynamics Project:

Mapping and Simulating the American Cultural Ecosystem

October 4 & 5, 2003

Madison, Wisconsin

Project Home Page: http://www.bolzcenter.org/culturaldynamics.php

---

Readings

The following readings are intended to provide a brief overview of systems thinking, while framing its potential, challenges, and implications. These readings are clearly not intended to be comprehensive, but merely to provide a flavor for our focus in Madison. If you have other articles, references, books, or resources to suggest, please send them along. We would love to develop a resource listing to distribute after the event.

Project Brief: The Cultural Dynamics Project

> You received a version of this brief in your original invitation. It provides an overview of the project and its goals.

Richmond, Barry, "A Pressing Need: Improving Performance," Chapter 1 from *An Introduction To Systems Thinking*, High Performance Systems, Inc., 2001, available on-line: http://www.hps-inc.com/hps_resources.htm

> While this excerpt has an initial focus on corporate process improvement and reengineering, it offers a quick overview of the needs for systems thinking, and some of the basic principals.

Meadows, Donella, "Dancing with Systems," *Whole Earth*, Winter 2001, available on-line: http://www.wholeearthmag.com/ArticleBin/447.html

"Places to Intervene in a System," *Whole Earth*, Winter 1997, available on-line: http://www.wholeearthmag.com/ArticleBin/109.html

> An early student of systems thinking at MIT, Donella Meadows provided a broad perspective of the potential and pitfalls of the discipline for addressing complex social and economic systems.

The Cultural Dynamics Project:
Mapping and Simulating the American Cultural Ecosystem
October 4 & 5, 2003
University of Wisconsin-Madison

A collaborative project of
Bolz Center for Arts Administration,
Cultural Initiatives Silicon Valley, and
National Arts Strategies

Funded by
The William and Flora Hewlett Foundation and
The David and Lucile Packard Foundation

In all its forms and functions, "policy" comprises three elements: an understanding of the way things work (systemic knowledge), a vision for how they might work better, and specific interventions aimed at achieving the desired vision. Modern medicine, in effect, applies these three elements of policy in the course of diagnosing and treating an individual patient. In medical practice, effective "policy" for a patient consists of a thorough diagnosis of a patient's physical/mental condition (understanding of the patient's system, what is working and what is not), comparison of that condition to a vision of good health, and the provision of specific interventions intended to advance the patient's condition toward that vision of health. While these three components of a systematically formulated "policy" for an individual's health may seem self-evident, they are often missing or incomplete in public policies and organizational planning. In arts and culture, such an approach seems entirely absent.

Too often, what passes for cultural policy or organizational planning is simply a reaction to a crisis: the threatened reduction of funding for the National Endowment for Arts, the loss of artist live/work space in the face of urban development, the sudden shift in purchase patterns among traditional audiences. In these moments of perceived crisis, resources are mobilized to push back on the threat and restore the status quo. Only rarely are there thoughtful attempts to understand the broad condition of the arts and culture, to formulate integrated visions of how they might be fully healthy, and to enact well-conceived interventions at the macro and micro level.

The case can be made that the arts and culture in America constitute a unique ecosystem whose features have been shaped, and continue to evolve, through the confluence of a wide range of influences. Governmental tax policies, subventions and regulations are obvious influences, but many other factors have been even more significant. Over more than a century, trends in immigration, technology, education, labor, philanthropy, capital infrastructure, urbanization, consumer fashion, consumer spending, and the economy have affected the content, quantity, organizational structures, and creativity of the American cultural sector.

The purpose of the Cultural Dynamics Project is to map the key factors that are driving the American cultural ecosystem using the tools of systems thinking. In the course of the mapping, the principal forces and the directions of their causality will be described graphically: what affects what, and to what extent? The map will be formulated by a veteran group of about 15 authorities who have observed or researched the evolution of broad facets of the ecosystem over the last four decades. The two-day session will provide an overview of systems thinking, and the basic elements of systems mapping, but move quickly to the active application of that discipline to the arts.

Following the meeting and the creation of the map, a leading expert on systems dynamics will develop a working computer simulation. The simulation will include a user-friendly interface that will enable a variety of users to visualize the workings of the cultural ecosystem and to project policy scenarios into the future. Once the map and simulation have been created, they will be circulated into several policy, academic, and research circles to ascertain whether they are worthwhile tools. If the response is positive, wider circulation will be undertaken, and more ambitious macro and micro mapping and simulation projects will ensue.

Systems thinking has provided a useful toolset for disciplines as diverse as meteorology, biology, social science, and human systems, areas that all share a complex structure of dynamically interconnected elements. The Cultural Dynamics Project seeks to turn that same toolset and conceptual framework toward the reasoned development of cultural policy, the responsive management of cultural institutions, and the dynamic training of future leaders in the field.

ABOUT THE CONVENERS

The project is being convened by three organizations that have already begun working toward a more systemic view of arts and culture: Cultural Initiatives Silicon Valley (www.ci-sv.org) has used systems simulations/games as a device for analysis and teaching in the domain of regional cultural policy; National Arts Strategies (www.artstrategies.org) has been redirecting its national consulting and support services toward adaptability and responsive management; and the Bolz Center for Arts Administration (www.bolzcenter.org) has been reshaping its graduate business degree curriculum drawing on the concepts and skills of systems thinking. All three organizations are working to expand and explore these early steps through the insights and input of the October convening, and to build a national network of experts to carry the conversation forward.

ABOUT THE FUNDERS

The project has been funded by The William and Flora Hewlett Foundation and The David and Lucile Packard Foundation, whose representatives will also participate in the Madison event.

PROJECT SCHEDULE

All sessions will be hosted in the Fluno Center for Executive Education on the University of Wisconsin-Madison campus (www.fluno.com). Participants should plan to arrive on the afternoon or evening of Friday, October 3, and depart in the late afternoon or evening of Sunday, October 5 (we will plan to complete our meeting by Noon on Sunday).

Chapter 1

A Pressing Need:

Improving Performance

It is estimated that more than 75% of reengineering efforts do not produce targeted performance improvements. The collapse of the dot.com boom bears vivid testimony to the fact that growth strategies often fail to yield real growth. The great majority of large-scale projects overrun both schedule and budget by very wide margins. Among the avalanche of mergers and acquisitions that has unfolded over the last decade, those that have realized anticipated synergies, number in the small handfuls. Stories abound of costly organizational change efforts that either have fizzled, or worse, exacerbated the situations they aimed at improving. The number of organizations with Balanced Scorecards—replete with metrics that no one understands how to use to improve performance—is approaching epidemic proportions.

How come? Why do so many well-intentioned performance-improvement efforts, conceived by so many smart people, so often miss the mark? And, perhaps more importantly, what can we do about it? What will it take to significantly increase the likelihood that the initiatives we design can achieve the results we intend? These are the questions we'll explore in this Chapter.

Getting to Root Cause

The first step in "fixing" *anything* is to understand why it's broken. If, in general, our performance improvement initiatives too often fall short, a good place to start looking for "why" is at the process by which these initiatives come into being. So how *do* our performance initiatives come into being? The simple answer is: *We think 'em up*! That is, they arise out of the process of thinking. So, let's take a closer look at *that* process.

The first thing to note about thinking is that when we ponder something, we do not actually have that "something" in our head. Think about it… You're trying to figure out whether you should let your kid drive to the party. You're struggling to decide whether to quit your steady, but relatively unchallenging day job, to pursue wild and wooly challenges at a start-up. You're wondering about the best way to reduce cycle-time in your customer support process. Whatever it is you are thinking about, you do not have it in your head. Then, what do you have? What are you working with when you're "thinking?"

You're working with a "mental model"—which is to say, a "selective abstraction" of the reality about which you are thinking. You've constructed that model using certain assumptions about how reality, in general, works, and also certain specific assumptions about the particular piece of reality you're thinking about. Let's go through a simple example to make these ideas more concrete.

You're at a nice restaurant. You are thinking about what to have for dinner. The mental model you are "working with" probably includes certain *general* assumptions about the reality of eating, such as: eating makes my hunger go away; when I eat too fast I get indigestion; if I eat dinner with my hands, people will think I'm a slob; and so forth. I'll refer to such general assumptions as "meta assumptions," because they transcend the specifics of any given eating situation. As you'll see, the "meta assumptions" we use when constructing our mental models will play an important role in explaining why our performance-improvement initiatives often don't fare so well. Your dinner-related model also will include some assumptions *specific* to the particular eating situation: the beef here is superb; I'll have a dry, red with dinner; and so forth.

Once you've assembled a preliminary set of assumptions into a mental model, you then "think" with them. I'll use a more operational term to describe what you are doing with them. I'll call it "mental simulation." You are *simulating* your mental model; you're "running what if's"…"Yah, the beef is good here, but what about my cholesterol? I can already taste the wine, but the roads are icy and I don't want to chance it." And so on. You run these simulations in an effort to predict what outcomes in reality are implied by the set of assumptions that constitute your mental model of that reality.

Does this square with your experience of what goes on when *you* "think?" The description seems to work pretty well for the people with whom I've talked. And so…if, when we create any sort of performance-improvement initiative, we think…and, when we think, we construct, then simulate, a mental model…then, if our performance-improvement initiatives consistently come up short of the mark, it is reasonable to suspect that something is awry in the processes by which we construct and simulate our mental models.

"What's up" with our Mental Model Construction & Simulation Processes?

Each of us has been constructing and simulating mental models for virtually our entire lifetime. And, since practice makes perfect, we ought to be pretty good at doing so! Let's test this plausible conjecture...

What follows is a passage that describes a very simple supply chain. Use it to construct a mental model. Then, simulate the model in order to predict how the system will perform in response to the "disturbance" to which it will be exposed.

> A retailer maintains an inventory of product that is shipped to customers on demand. Upon shipping, the retailer orders more product (to re-stock inventory) from the firm that supplies it. The retailer always emails an order to the supplier for an amount of product exactly equal to what was shipped in a given day. If ten units go out in a day, the retailer emails an order for ten units at the end of the day. The retailer never delays in placing the order, and always orders *exactly* the amount of product that was shipped in a given day.
>
> The supplier also is very regular. The supplier always processes the retailer's order immediately upon receipt, then ships the requested amount of product to the retailer. Product always arrives six days after the retailer places the order. The supplier has never been out-of-stock (and never will be!), and has always (and will always) be able to get product to the retailer exactly six days after the retailer's order is placed. Furthermore, no product shipped by the supplier is ever, or will ever be, defective, damaged or lost in transit.
>
> This simple supply chain has been in steady-state for some time. This means that the volume of product being demanded at retail by customers has been constant at some level for a long time, as has therefore the volume of product the retailer has been ordering from the supplier, as well as the amount the supplier has been shipping to the retailer. Everything is in perfect, constant balance. Now suppose, all of a sudden, the volume of demand from customers coming into the retailer steps up to a new higher level, and then remains there (i.e., a one-time, step-increase occurs). On the axes provided in Figure 1-1, sketch the *pattern* you think will be traced by the level of the retailer's inventory, over time, following the one-time step-increase in customer demand.

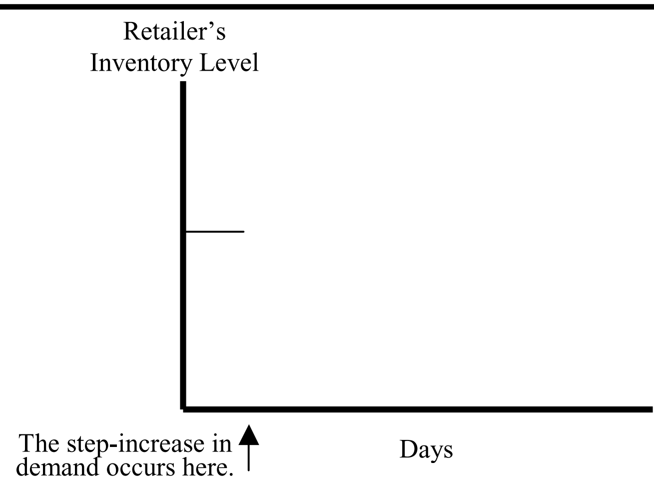

Figure 1-1.
A Sketch of Your Prediction.

Usually, upwards of *80%* of any group who is asked to conduct this type of thought experiment traces an incorrect pattern! The correct pattern is that: *following the step-increase in demand, the Retailer's inventory will decline in a straight-line manner for six days; it then will level off and remain at the new, lower level*. (You'll develop an understanding of why in the next chapter). The relatively small percentage of people who *do* trace the correct pattern has proven to be independent of culture, education level, or experience with supply chains. These results strongly suggest that human beings, in general, either are not very good at constructing mental models (of even very simple systems!), performing mental simulations of these models, or both!

Why Are We Not So Good at Constructing & Simulating Mental Models?

So how come we're not better at constructing and/or simulating mental models—especially given all the experience we've had doing it? I will argue that it's due to a difference in the speed with which biological and socio-cultural systems evolve. The differential speed of evolution has produced a human species whose cognitive machinery is pretty much what it always was, and an operating reality that has become vastly more complex and interdependent. It's this mismatch that's the root of the problem.

Our Simulation Machinery

Simply stated, when our ancestors got thumbs and began to stand up, they unfortunately didn't simultaneously get a huge boost in their cognitive capacities. And, they really didn't need one…at that time. Back when we still lived in caves, our mental simulations served us well. The rules were simple. See bear, whack bear, eat bear…maybe even share. Bear were abundant. Clubs and rocks were "local" weapons. Bear meat wasn't laced with additives, heavy metals, and/or pesticides. We didn't have to trade off time spent hunting, with our day jobs and the kids' soccer practice. Lawyers weren't yet invented. Life was straightforward. Our mental models were very simple. The associated simulations were slam-dunks.

Then came "progress." We created tools, used them to decimate most of the bear, started wearing bear coats and growing our own food, someone invented MTV…and the rest is, as they say, history! Life got complex. It became difficult to do anything without inadvertently causing a bunch of other things to happen—most of which we remained oblivious to. Everything became a "competition." We began competing for resources, people, time, mind-share. Free lunches were all eaten.

The problem was simply that socio-cultural evolution happened too fast for cognitive evolution to keep pace. To this day, we still can't juggle more than a few variables in our head at a time. And, as far as reliably tracing out the consequences of an action over more than a very limited time horizon…fugeddaboudit! As the little mental simulation exercise you just completed demonstrates, our cognitive machinery limits our ability to conduct reliable mental simulations of even the most elementary sets of relationships.

And, while inadequate mental simulation capability is bad enough, unfortunately, there's *more* bad news! Growing evidence, not the least of which is our record with performance-improvement initiatives, suggests that the mental models we construct do not capture enough of the essence of how reality actually works! There are three reasons why these models don't pass muster: (1) what's in them, (2) how what's in them is represented, and (3) the process for honing both content and representation. We'll examine each…

Reason 1 for Poor Quality Mental Models: *Content*

Problems with the quality of our mental models begin with what we choose to put in them…and what we choose to leave out—that is, how we choose to "filter" reality for purposes of selecting material for inclusion in our mental models.

The "contents" problem again harkens back to our ancestral past as individual actors in a perilous natural environment. Our neurobiology was honed to respond to what was right in front of us—both in space and time. And for good reason: what was right in front of us could *kill* us—a fact which, unfortunately, remains too true even today! Content-wise, our ancestors' mental models contained lots of detail about what was *immediate,* in both space and time. We knew a lot…about a little. The fact that our weed-level perspective afforded only a limited view of the overall garden was OK, because cause and effect connections were short and direct. Our actions had immediate-in-time, local-in-space, impacts. "Overall garden" impacts just weren't an issue. Our neurobiology was well-adapted for surviving in the primeval garden.

And survive, we did. In fact, we thrived! Our "garden" is now pretty much fully populated—we now number in the billions. And instead of operating as individual actors, we're now members of communities and organizations who operate within a highly-interdependent web. Actions taken by individuals now regularly have "whole garden" impacts. Yet our neurobiological machinery remains essentially the same as when all we had to focus on was immediate! To make matters worse, the structure of many of today's organizations plays to the tendencies toward "localness" inherent in our neurobiology. Manufacturing, Sales, R&D, Finance, IT, HR, and Marketing "silos"—each with its own dialect and culture, each with its well-defined spatial boundaries—encourage the development of highly "local" mental models. Like our ancestors, we continue to know a lot about a little. And, Wall Street does its part to make sure we don't forget about Bears—keeping us locally-focused in time, by making everything ride on *this* quarter's earnings.

So, while almost any action to improve performance taken today has extensive ramifications, both spatial and temporal, the *contents* of our mental models (i.e., the associated boundaries) do not allow us to "think through" these ramifications! As a result, we get "surprised" a lot—and usually the surprises are not pleasant. In addition, because we don't capture the ramifications, it's not possible to *learn* from them! Hence, we are destined to re-live past mistakes. Figure 1-2 depicts the situation…

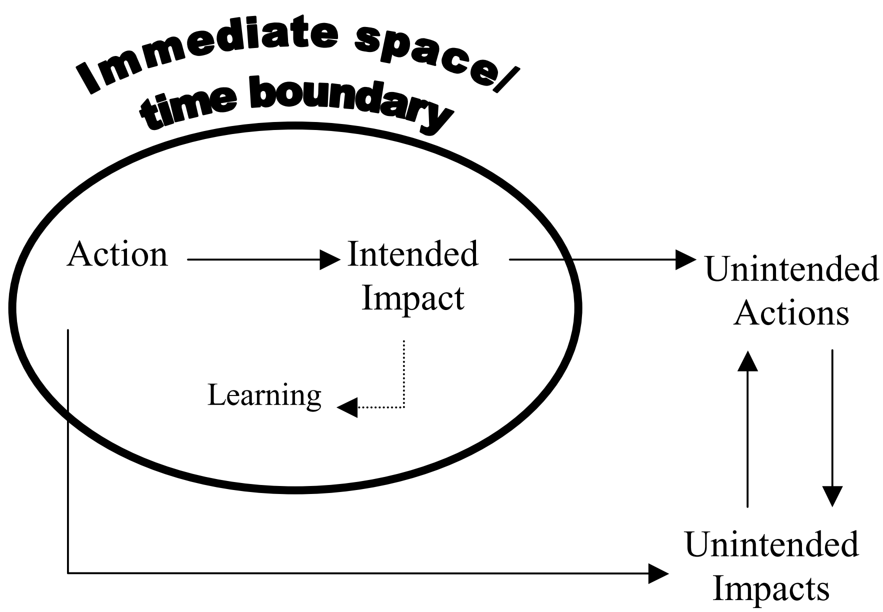

Figure 1-2.
Deep, Narrow Content Undermines the Reliability of Mental Simulation and Limits Learning.

Thus, the first step in improving the quality of our mental models is to improve their content. To do that, we need a better "filter." We need a perspective that allows us to capture content that will enable us to "see" beyond the immediate in space and time, and that will prevent us from getting so bogged down in the weeds that we can't appreciate the "whole garden." As we'll see in Chapter 2, Systems Thinking offers one such perspective.

Reason 2 for Poor Quality Mental Models: Representation of Content

Even if we were able to improve the filter we use for selecting content for our mental models, we'd still need to improve the way we *represent* that content. Simply stated, the "meta assumptions" we use to structure our mental models are not sufficiently congruent with reality. As a result, the "structure" of our mental models does not mirror reality closely enough to yield reliable inferences when simulating them.

Because we make such extensive use of "meta assumptions," they submerge…outright disappear from consciousness! They become so "obviously true," they're no longer subject to scrutiny or question. But if we are to have any hope of improving upon these assumptions, we must first bring them back into view. One way to surface them is to identify conceptual frameworks and analytical tools that are in widespread use in diverse arenas. The fact that they are widely used suggests they mask a set of commonly embraced "meta assumptions."

A popular candidate on the conceptual framework front is what we might label "Critical Success Factors Thinking." Most organizations have identified a set of *critical success factors*. The set most often manifests as a list of "drivers of the business." You see them tacked up on cubicle partitions, taped to conference room walls, and on little laminated cards that people carry around in their wallets. From service delivery to heavy manufacturing to educational institutions, all sorts of organizations have them. And, individuals also have embraced the critical success factors framework. One of best-selling popular books of all time is Steven Covey's *The Seven Habits of Highly Effective People*—critical success factors for individuals seeking to live the "right life." Numerous other best-sellers offer similar success factor recipes for "prevailing" in our complex, fast-paced times.

If we were to diagram the generic structure that underlies a "critical success factor" (CSF) model, it would look like what you see in Figure 1-3.

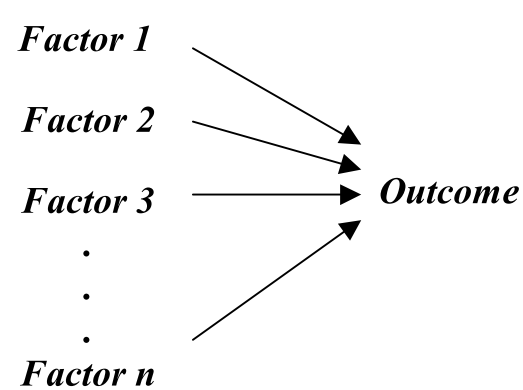

Figure 1-3.
The Generic Structure of a "Critical Success Factor" Model.

OK, so what "meta assumptions" does this structure reveal? Two obvious ones suggest themselves. The first is that the "Factors" operate *independently*. Each "impacts" the outcome, but it does so, independently. The second is that the "Outcome" does not *cycle back* to influence any of the Factors. That is, *causality is assumed to run one-way*—from Factor to Outcome, but not back again.

Both "meta assumptions" are highly suspect! In today's highly-interdependent world, it's difficult to find *any* "Factor" that doesn't influence, and isn't influenced by, multiple other Factors. Consider an example from an organizational context. A firm might list, say, technology, good people, and learning as three "drivers" of success. But is it not the case that, top-quality people create good technology, and that good technology is part of what enables people to remain "top-quality?" And further, isn't it "learning" that drives technological advance, and technological advance that, in turn, drives learning? Don't top-quality people learn faster and better than lower-quality people? And isn't it the opportunity to learn that's key to attracting and retaining top-quality people? So much for the *independence* of "Factors" assumption!

The other "meta assumption"—that causality runs one-way, from driver to outcome (and not back again)—is equally easy to dispatch. Certainly it's true that top-quality people help to create successful organizations. But does it end there? Is that it? You get top-quality people, they help create a successful organization, and everyone lives happily ever after, end of story? Isn't the following storyline more congruent with reality as you know it? An organization is spawned by some top-quality people who, if everything comes together, begin to have some success. The success, in turn, attracts the attention of other high-quality people who are offered opportunities in the expanding organization. More success is created, and more top-quality people are attracted…and we're off to the *reciprocal causality* races. At some point, the organization will encounter some type of "limits to growth" (nothing can spiral forever!). How the organization addresses these limits will determine whether the spiral continues upward, reverses direction producing a nosedive, or settles into some sort of steady-state.

And so, isn't there really a reciprocal, or closed-loop, causal relationship between top-quality people (or any of the other "Factors") and organizational success? Success is not just an outcome, something that is "driven" by a set of Factors. Success is, itself, a driver! Causality runs *both* ways, not *one-way*! Commonly-employed "meta assumption" number two lands with a thud!

If we look a little more closely at "Critical Success Factors" models, we can infer the existence of other "meta assumptions." The assumptions also are clearly evident in some of the highly popular analytical tools in use today. So, let's use them for our examples.

One of these tools is the spreadsheet. Another is *The Balanced Scorecard* bubble diagram. A third is "root cause" or "fishbone" diagrams. In the artifacts created by each tool, like the CSF framework, we find "logic trees" with associated causality running only one way. We also often find more independent than interdependent factors. But these popular tools also generally reflect two other "meta assumptions," as well. The first of these is that impacts are felt *instantaneously* (i.e., delays are largely ignored). The second is that impacts are *linear* and *constant* (i.e., an x% change in input always results in a y% change in output).

Looking at "instantaneous impacts," virtually every system/process known to humankind has some inertia in it. Almost nothing responds instantly—at least not the *total* response! There may be some immediate reactions to things, but these usually set in motion other reactions that take time to play out. Delays are a ubiquitous fact of life! They're an important attribute of both organizational and individual reality. Similarly, looking at the second assumption (impacts are linear), what makes life interesting, and impacts so difficult to predict, is that sometimes you can push "a ton" and get an ounce, while other times the tickle of a feather brings down the house! Like delays, non-linear relationships are an essential characteristic of operating reality. The validity of two more popular "meta assumptions" are thus called into question.

If we are to improve the quality of the representations of content within our mental models, we need a better set of "meta assumptions!" In place of the assumptions of independence, one-way causality, and impacts that are instantaneous and linear, we need assumptions that celebrate interdependence, closed-loop causality, delays and non-linearities! Only when the representations in our mental models commonly bear these characteristics, will we increase the likelihood that the initiatives we design will create the outcomes we intend.

Reason 3 for Poor Quality Mental Models: The Honing Process

So, fine…our biology and modern-day organizational structures encourage us to form narrow "filters" that restrict the content of our mental models. And, the "meta assumptions" we employ destine us to represent that content in ways that do not mirror how reality actually works. But, as a result, after "getting it wrong" so many times, why haven't we figured it out and improved our mental models? We continue to lack a process for systematically improving the quality of the content, the representation of content, and the simulation of our mental models. In short, neither our individual, nor organizational, learning processes are very effective. We're pretty good at Knowledge Management (collecting, storing and retrieving knowledge), but we're *very* poor at Understanding Management (collecting, storing and retrieving understanding). Why? First, we don't have a sharable language for integrating our "piece understanding" into a coherent picture of "the whole." And second, we don't have tools for then testing the validity of that understanding. We'll take them one at a time…

On the sharable language score, as already noted, most organizations are collections of functional, divisional, and/or geographic fiefdoms. People who understand "the whole" are rare. Those who understand a "piece" are abundant. If it were possible to somehow *knit together* the "piece understanding" into a *manageable* picture of the whole, we'd all be working with a fundamentally better mental model of the reality within which we are operating. So, what stands in our way? Two things. The first is the absence of an *Esperanto*, a universal language that offers a common set of symbols for accomplishing the "knitting together." The second is a framework that provides a "filter" that passes just what's essential about the way the whole works, without admitting all of the piece detail. This gives us the "manageable" part. Systems Thinking, as you'll discover in Chapter 2, can provide both!

On the tools front, assuming we succeed in knitting together piece understanding into a manageable picture of the "whole," we'd then need a way to rigorously test the assumptions that constitute this understanding against reality. We need to test our assumptions both *before* implementing our initiatives, and we also need to be able to double-back to re-visit them *after* reality has performed *its* simulation! Pre-implementation tests give us the opportunity to ferret out internal inconsistencies and to surface "blind spots" (places where we need further information and understanding). Tools, here, are serving as "practice fields"—no risk, rapid-turnaround opportunities to learn before having to do it for real. Post-implementation tests provide opportunities to discover how and why model-projected outcomes differed from what reality actually served up. When discrepancies arise, model assumptions can be modified to better reflect how reality actually works. As a result, over time, the organization's collective understanding can be continuously and systematically improved.

As you'll see in Chapter 2, the *ithink* software is a tool that has been designed to play the aforedescribed role. Used in conjunction with Systems Thinking, it can serve as a powerful resource for meeting the challenge of creating effective performance-improvement initiatives.

What's Next

In this Chapter, I've teed up the challenge: *improving our ability to create effective performance-improvement initiatives*. I've argued that the reason the record of success is not very distinguished is that the quality of the mental models underlying our performance-improvement initiatives is poor, and that the simulation of these models is unreliable. I've also asserted that Systems Thinking and the *ithink* software constitute a powerful tandem for supporting your efforts to improve this situation. Chapter 2 takes on the task of supporting this assertion.

Dancing with Systems
What to do when systems resist change;
an excerpt from Donella Meadows's unfinished last book.
http://www.wholeearthmag.com/ArticleBin/447.html

By Donella Meadows
(Whole Earth Winter 2001)

People who are raised in the industrial world and who get enthused about systems thinking are likely to make a terrible mistake. They are likely to assume that here, in systems analysis, in interconnection and complication, in the power of the computer, here at last, is the key to prediction and control. This mistake is likely because the mindset of the industrial world assumes that there is a key to prediction and control.

I assumed that at first too. We all assumed it, as eager systems students at the great institution called MIT. More or less innocently, enchanted by what we could see through our new lens, we did what many discoverers do. We exaggerated our own ability to change the world. We did so not with any intent to deceive others, but in the expression of our own expectations and hopes. Systems thinking for us was more than subtle, complicated mindplay. It was going to Make Systems Work.

But self-organizing, nonlinear, feedback systems are inherently unpredictable. They are not controllable. They are understandable only in the most general way. The goal of foreseeing the future exactly and preparing for it perfectly is unrealizable. The idea of making a complex system do just what you want it to do can be achieved only temporarily, at best. We can never fully understand our world, not in the way our reductionistic science has led us to expect. Our science itself, from quantum theory to the mathematics of chaos, leads us into irreducible uncertainty. For any objective other than the most trivial, we can't optimize; we don't even know what to optimize. We can't keep track of everything. We can't find a proper, sustainable relationship to nature, each other, or the institutions we create, if we try to do it from the role of omniscient conqueror.

For those who stake their identity on the role of omniscient conqueror, the uncertainty exposed by systems thinking is hard to take. If you can't understand, predict, and control, what is there to do?

Systems thinking leads to another conclusion, however—waiting, shining, obvious as soon as we stop being blinded by the illusion of control. It says that there is plenty to do, of a different sort of "doing." The future can't be predicted, but it can be envisioned and brought lovingly into being. Systems can't be controlled, but they can be designed and redesigned. We can't surge forward with certainty into a world of no surprises, but we can expect surprises and learn from them and even profit from them. We can't impose our will upon a system. We can listen to what the system tells us, and discover how its properties and our values can work together to bring forth something much better than could ever be produced by our will alone.

We can't control systems or figure them out. But we can dance with them! I already knew that, in a way before I began to study systems. I had learned about dancing with great powers from whitewater kayaking, from gardening, from playing music, from skiing. All those endeavors require one to stay wide awake, pay close attention, participate flat out, and respond to feedback. It had never occurred to me that those same requirements might apply to intellectual work, to management, to government, to getting along with people.

But there it was, the message emerging from every computer model we made. Living successfully in a world of systems requires more of us than our ability to calculate. It requires our full humanity—our rationality, our ability to sort out truth from falsehood, our intuition, our compassion, our vision, and our morality.

I will summarize the most general "systems wisdoms" I have absorbed from modeling complex systems and hanging out with modelers. These are the take-home lessons, the concepts and practices that penetrate the discipline of systems so deeply that one begins, however imperfectly, to practice them not just in one's profession, but in all of life.

The list probably isn't complete, because I am still a student in the school of systems. And it isn't unique to systems thinking. There are many ways to learn to dance. But here, as a start-off dancing lesson, are the practices I see my colleagues adopting, consciously or unconsciously, as they encounter systems.

Get the beat.
Before you disturb the system in any way, watch how it behaves. If it's a piece of music or a whitewater rapid or a fluctuation in a commodity price, study its beat. If it's a social system, watch it work. Learn its history. Ask people who've been around a long time to tell you what has happened. If possible, find or make a time graph of actual data from the system. Peoples' memories are not always reliable when it comes to timing.

Starting with the behavior of the system forces you to focus on facts, not theories. It keeps you from falling too quickly into your own beliefs or misconceptions, or those of others. It's amazing how many misconceptions there can be. People will swear that rainfall is decreasing, say, but when you look at the data, you find that what is really happening is that variability is increasing—the droughts are deeper, but the floods are greater too. I have been told with great authority that milk price was going up when it was going down, that real interest rates were falling when they were rising, that the deficit was a higher fraction of the GNP than ever before when it wasn't.

Starting with the behavior of the system directs one's thoughts to dynamic, not static analysis—not only to "what's wrong?" but also to "how did we get there?" and "what behavior modes are possible?" and "if we don't change direction, where are we going to end up?"

And finally, starting with history discourages the common and distracting tendency we all have to define a problem not by the system's actual behavior, but by the lack of our favorite solution. (The problem is, we need to find more oil. The problem is, we need to ban abortion. The problem is, how can we attract more growth to this town?)

Listen to the wisdom of the system.
Aid and encourage the forces and structures that help the system run itself. Don't be an unthinking intervener and destroy the system's own self-maintenance capacities. Before you charge in to make things better, pay attention to the value of what's already there.

A friend of mine, Nathan Gray, was once an aid worker in Guatemala. He told me of his frustration with agencies that would arrive with the intention of "creating jobs" and "increasing entrepreneurial abilities" and "attracting outside investors." They would walk right past the thriving local market, where small-scale business people of all kinds, from basket-makers to vegetable growers to butchers to candy sellers, were displaying their entrepreneurial abilities in jobs they had created for themselves. Nathan spent his time talking to the people in the market, asking about their lives and businesses, learning what was in the way of those businesses expanding and incomes rising. He concluded that what was needed was not outside investors, but inside ones. Small loans available at reasonable interest rates, and classes in literacy and accounting, would produce much more long-term good for the community than bringing in a factory or assembly plant from outside.

Expose your mental models to the open air.
Remember, always, that everything you know, and everything everyone knows, is only a model. Get your model out there where it can be shot at. Invite others to challenge your assumptions and add their own. Instead of becoming a champion for one possible explanation or hypothesis or model, collect as many as possible. Consider all of them plausible until you find some evidence that causes you to rule one out. That way you will be emotionally able to see the evidence that rules out an assumption with which you might have confused your own identity.

You don't have to put forth your mental model with diagrams and equations, though that's a good discipline. You can do it with words or lists or pictures or arrows showing what you think is connected to what. The more you do that, in any form, the clearer your thinking will become, the faster you will admit your uncertainties and correct your mistakes, and the more flexible you will learn to be. Mental flexibility—the willingness to redraw boundaries, to notice that a system has shifted into a new mode, to see how to redesign structure—is a necessity when you live in a world of flexible systems.

Stay humble. Stay a learner.
Systems thinking has taught me to trust my intuition more and my figuring-out rationality less, to lean on both as much as I can, but still to be prepared for surprises. Working with systems, on the computer, in nature, among people, in organizations, constantly reminds me of how incomplete my mental models are, how complex the world is, and how much I don't know.

The thing to do, when you don't know, is not to bluff and not to freeze, but to learn. The way you learn is by experiment—or, as Buckminster Fuller put it, by trial and error, error, error. In a world of complex systems it is not appropriate to charge forward with rigid, undeviating directives. "Stay the course" is only a good idea if you're sure you're on course. Pretending you're in control even when you aren't is a recipe not only for mistakes, but for not learning from mistakes. What's appropriate when you're learning is small steps, constant monitoring, and a willingness to change course as you find out more about where it's leading.

That's hard. It means making mistakes and, worse, admitting them. It means what psychologist Don Michael calls "error-embracing." It takes a lot of courage to embrace your errors.

Honor and protect information.
A decision-maker can't respond to information he or she doesn't have, can't respond accurately to information that is inaccurate, can't respond in a timely way to information that is late. I would guess that 99 percent of what goes wrong in systems goes wrong because of faulty or missing information.

If I could, I would add an Eleventh Commandment: Thou shalt not distort, delay, or sequester information. You can drive a system crazy by muddying its information streams. You can make a system work better with surprising ease if you can give it more timely, accurate, and complete information.

For example, in 1986 new federal legislation required US companies to report all chemical emissions from each of their plants. Through the Freedom of Information Act (from a systems point of view one of the most important laws in the nation) that information became a matter of public record. In July 1988 the first data on chemical emissions became available. The reported emissions were not illegal, but they didn't look very good when they were published in local papers by enterprising reporters, who had a tendency to make lists of "the top ten local polluters." That's all that happened. There were no lawsuits, no required reductions, no fines, no penalties. But within two years chemical emissions nationwide (as least as reported, and presumably also in fact) had decreased by 40 percent. Some companies were launching policies to bring their emissions down by 90 percent, just because of the release of previously sequestered information.

Locate responsibility in the system.
Look for the ways the system creates its own behavior. Do pay attention to the triggering events, the outside influences that bring forth one kind of behavior from the system rather than another. Sometimes those outside events can be controlled (as in reducing the pathogens in drinking water to keep down incidences of infectious disease). But sometimes they can't. And sometimes blaming or trying to control the outside influence blinds one to the easier task of increasing responsibility within the system.

"Intrinsic responsibility" means that the system is designed to send feedback about the consequences of decision-making directly and quickly and compellingly to the decision-makers.

Dartmouth College reduced intrinsic responsibility when it took thermostats out of individual offices and classrooms and put temperature-control decisions under the guidance of a central computer. That was done as an energy-saving measure. My observation from a low level in the hierarchy is that the main consequence was greater oscillations in room temperature. When my office gets overheated now, instead of turning down the thermostat, I have to call an office across campus, which gets around to making corrections over a period of hours or days, and which often overcorrects, setting up the need for another phone call. One way of making that system more, rather than less, responsible, might have been to let professors keep control of their own thermostats and charge them directly for the amount of energy they use. (Thereby privatizing a commons!)

Designing a system for intrinsic responsibility could mean, for example, requiring all towns or companies that emit wastewater into a stream to place their intake pipe downstream from their outflow pipe. It could mean that neither insurance companies nor public funds should pay for medical costs resulting from smoking or from accidents in which a motorcycle rider didn't wear a helmet or a car rider didn't fasten the seat belt. It could mean Congress would no longer be allowed to legislate rules from which it exempts itself.

Make feedback policies for feedback systems.
President Jimmy Carter had an unusual ability to think in feedback terms and to make feedback policies. Unfortunately he had a hard time explaining them to a press and public that didn't understand feedback.

He suggested, at a time when oil imports were soaring, that there be a tax on gasoline proportional to the fraction of US oil consumption that had to be imported. If imports continued to rise the tax would rise, until it suppressed demand and brought forth substitutes and reduced imports. If imports fell to zero, the tax would fall to zero.

The tax never got passed.

Carter was also trying to deal with a flood of illegal immigrants from Mexico. He suggested that nothing could be done about that immigration as long as there was a great gap in opportunity and living standards between the US and Mexico. Rather than spending money on border guards and barriers, he said, we should spend money helping to build the Mexican economy, and we should continue to do so until the immigration stopped.

That never happened either.

You can imagine why a dynamic, self-adjusting system cannot be governed by a static, unbending policy. It's easier, more effective, and usually much cheaper to design policies that change depending on the state of the system. Especially where there are great uncertainties, the best policies not only contain feedback loops, but meta-feedback loops—loops that alter, correct, and expand loops. These are policies that design learning into the management process.

Pay attention to what is important, not just what is quantifiable.
Our culture, obsessed with numbers, has given us the idea that what we can measure is more important than what we can't measure. You can look around and make up your own mind about whether quantity or quality is the outstanding characteristic of the world in which you live.

If something is ugly, say so. If it is tacky, inappropriate, out of proportion, unsustainable, morally degrading, ecologically impoverishing, or humanly demeaning, don't let it pass. Don't be stopped by the "if you can't define it and measure it, I don't have to pay attention to it" ploy. No one can [precisely] define or measure justice, democracy, security, freedom, truth, or love. No one can [precisely] define or measure any value. But if no one speaks up for them, if systems aren't designed to produce them, if we don't speak about them and point toward their presence or absence, they will cease to exist.

Go for the good of the whole.
Don't maximize parts of systems or subsystems while ignoring the whole. As Kenneth Boulding once said, don't go to great trouble to optimize something that never should be

done at all. Aim to enhance total systems properties, such as [creativity], stability, diversity, resilience, and sustainability—whether they are easily measured or not.

As you think about a system, spend part of your time from a vantage point that lets you see the whole system, not just the problem that may have drawn you to focus on the system to begin with. And realize that, especially in the short term, changes for the good of the whole may sometimes seem to be counter to the interests of a part of the system. It helps to remember that the parts of a system cannot survive without the whole. The long-term interests of your liver require the long-term health of your body, and the long-term interests of sawmills require the long-term health of forests.

Expand time horizons.
The official time horizon of industrial society doesn't extend beyond what will happen after the next election or beyond the payback period of current investments. The time horizon of most families still extends farther than that—through the lifetimes of children or grandchildren. Many Native American cultures actively spoke of and considered in their decisions the effects upon the seventh generation to come. The longer the operant time horizon, the better the chances for survival.

In the strict systems sense there is no long-term/short-term distinction. Phenomena at different timescales are nested within each other. Actions taken now have some immediate effects and some that radiate out for decades to come. We experience now the consequences of actions set in motion yesterday and decades ago and centuries ago.

When you're walking along a tricky, curving, unknown, surprising, obstacle-strewn path, you'd be a fool to keep your head down and look just at the next step in front of you. You'd be equally a fool just to peer far ahead and never notice what's immediately under your feet. You need to be watching both the short and long terms—the whole system.

Expand thought horizons.
Defy the disciplines. In spite of what you majored in, or what the textbooks say, or what you think you're an expert at, follow a system wherever it leads. It will be sure to lead across traditional disciplinary lines. To understand that system, you will have to be able to learn from—while not being limited by—economists and chemists and psychologists and theologians. You will have to penetrate their jargons, integrate what they tell you, recognize what they can honestly see through their particular lenses, and discard the distortions that come from the narrowness and incompleteness of their lenses. They won't make it easy for you.

Seeing systems whole requires more than being "interdisciplinary," if that word means, as it usually does, putting together people from different disciplines and letting them talk past each other. Interdisciplinary communication works only if there is a real problem to be solved, and if the representatives from the various disciplines are more committed to solving the problem than to being academically correct. They will have to go into learning mode, to admit ignorance and be willing to be taught, by each other and by the system.

It can be done. It's very exciting when it happens.

Expand the boundary of caring.
Living successfully in a world of complex systems means expanding not only time horizons and thought horizons; above all it means expanding the horizons of caring. There are moral reasons for doing that, of course. And if moral arguments are not sufficient, systems thinking provides the practical reasons to back up the moral ones. The real system is interconnected. No part of the human race is separate either from other human beings or from the global ecosystem. It will not be possible in this integrated world for your heart to succeed if your lungs fail, or for your company to succeed if your workers fail, or for the rich in Los Angeles to succeed if the poor in Los Angeles fail, or for Europe to succeed if Africa fails, or for the global economy to succeed if the global environment fails.

As with everything else about systems, most people already know the interconnections that make moral and practical rules turn out to be the same rules. They just have to bring themselves to believe what they know.

Celebrate complexity.
Let's face it, the universe is messy. It is nonlinear, turbulent, and chaotic. It is dynamic. It spends its time in transient behavior on its way to somewhere else, not in mathematically neat equilibria. It self-organizes and evolves. It creates diversity, not uniformity. That's what makes the world interesting, that's what makes it beautiful, and that's what makes it work.

There's something within the human mind that is attracted to straight lines and not curves, to whole numbers and not fractions, to uniformity and not diversity, and to certainties and not mystery. But there is something else within us that has the opposite set of tendencies, since we ourselves evolved out of and are shaped by and structured as complex feedback systems. Only a part of us, a part that has emerged recently, designs buildings as boxes with uncompromising straight lines and flat surfaces. Another part of us recognizes instinctively that nature designs in fractals, with intriguing detail on every scale from the microscopic to the macroscopic. That part of us makes Gothic cathedrals and Persian carpets, symphonies and novels, Mardi Gras costumes and artificial intelligence programs, all with embellishments almost as complex as the ones we find in the world around us.

Hold fast to the goal of goodness.
Examples of bad human behavior are held up, magnified by the media, affirmed by the culture, as typical. Just what you would expect. After all, we're only human. The far more numerous examples of human goodness are barely noticed. They are Not News. They are exceptions. Must have been a saint. Can't expect everyone to behave like that.

And so expectations are lowered. The gap between desired behavior and actual behavior narrows. Fewer actions are taken to affirm and instill ideals. The public discourse is full of cynicism. Public leaders are visibly, unrepentantly, amoral or immoral and are not held to account. Idealism is ridiculed. Statements of moral belief are suspect. It is much easier to talk about hate in public than to talk about love.

We know what to do about eroding goals. Don't weigh the bad news more heavily than the good. And keep standards absolute.

This is quite a list. Systems thinking can only tell us to do these things. It can't do them for us. And so we are brought to the gap between understanding and implementation. Systems thinking by itself cannot bridge that gap. But it can lead us to the edge of what analysis can do and then point beyond—to what can and must be done by the human spirit.

Donella Meadows died in the spring of 2001. This article was excerpted from the manuscript of her unfinished last book.

Places to Intervene in a System
By Donella H. Meadows
(Whole Earth Winter 1997)
http://www.wholeearthmag.com/ArticleBin/109.html

Folks who do systems analysis have a great belief in "leverage points." These are places within a complex system (a corporation, an economy, a living body, a city, an ecosystem) where a small shift in one thing can produce big changes in everything.

The systems community has a lot of lore about leverage points. Those of us who were trained by the great Jay Forrester at MIT have absorbed one of his favorite stories. "People know intuitively where leverage points are. Time after time I've done an analysis of a company, and I've figured out a leverage point. Then I've gone to the company and discovered that everyone is pushing it in the wrong direction!"

The classic example of that backward intuition was Forrester's first world model. Asked by the Club of Rome to show how major global problems—poverty and hunger, environmental destruction, resource depletion, urban deterioration, unemployment—are related and how they might be solved, Forrester came out with a clear leverage point: Growth. Both population and economic growth. Growth has costs—among which are poverty and hunger, environmental destruction—the whole list of problems we are trying to solve with growth!

The world's leaders are correctly fixated on economic growth as the answer to virtually all problems, but they're pushing with all their might in the wrong direction.

Counterintuitive. That's Forrester's word to describe complex systems. The systems analysts I know have come up with no quick or easy formulas for finding leverage points. Our counterintuitions aren't that well developed. Give us a few months or years and we'll model the system and figure it out. We know from bitter experience that when we do discover the system's leverage points, hardly anybody will believe us.

Very frustrating. So one day I was sitting in a meeting about the new global trade regime, NAFTA and GATT and the World Trade Organization. The more I listened, the more I began to simmer inside. "This is a HUGE NEW SYSTEM people are inventing!" I said to myself. "They haven't the slightest idea how it will behave," myself said back to me. "It's cranking the system in the wrong direction—growth, growth at any price!! And the control measures these nice folks are talking about—small parameter adjustments, weak negative feedback loops—are PUNY!"

Suddenly, without quite knowing what was happening, I got up, marched to the flip chart, tossed over a clean page, and wrote: " Places to Intervene in a System ," followed by nine items:

9. Numbers (subsidies, taxes, standards).
8. Material stocks and flows.
7. Regulating negative feedback loops.
6. Driving positive feedback loops.
5. Information flows.
4. The rules of the system (incentives, punishment, constraints).
3. The power of self-organization.

2. The goals of the system.
1. The mindset or paradigm out of which the goals, rules, feedback structure arise.

Everyone in the meeting blinked in surprise, including me. "That's brilliant!" someone breathed. "Huh?" said someone else.

I realized that I had a lot of explaining to do.

In a minute I'll go through the list, translate the jargon, give examples and exceptions. First I want to place the list in a context of humility. What bubbled up in me that day was distilled from decades of rigorous analysis of many different kinds of systems done by many smart people. But complex systems are, well, complex. It's dangerous to generalize about them. What you are about to read is not a recipe for finding leverage points. Rather it's an invitation to think more broadly about system change.

That's why leverage points are not intuitive.

9. Numbers.
Numbers ("parameters" in systems jargon) determine how much of a discrepancy turns which faucet how fast. Maybe the faucet turns hard, so it takes a while to get the water flowing. Maybe the drain is blocked and can allow only a small flow, no matter how open it is. Maybe the faucet can deliver with the force of a fire hose. These considerations are a matter of numbers, some of which are physically locked in, but most of which are popular intervention points.

Consider the national debt. It's a negative bathtub, a money hole. The rate at which it sinks is the annual deficit. Tax income makes it rise, government expenditures make it fall. Congress and the president argue endlessly about the many parameters that open and close tax faucets and spending drains. Since those faucets and drains are connected to the voters, these are politically charged parameters. But, despite all the fireworks, and no matter which party is in charge, the money hole goes on sinking, just at different rates.

The amount of land we set aside for conservation. The minimum wage. How much we spend on AIDS research or Stealth bombers. The service charge the bank extracts from your account. All these are numbers, adjustments to faucets. So, by the way, is firing people and getting new ones. Putting different hands on the faucets may change the rate at which they turn, but if they're the same old faucets, plumbed into the same system, turned according to the same information and rules and goals, the system isn't going to change much. Bill Clinton is different from George Bush, but not all that different.

Numbers are last on my list of leverage points. Diddling with details, arranging the deck chairs on the Titanic. Probably ninety-five percent of our attention goes to numbers, but there's not a lot of power in them.

Not that parameters aren't important—they can be, especially in the short term and to the individual who's standing directly in the flow. But they RARELY CHANGE BEHAVIOR. If the system is chronically stagnant, parameter changes rarely kick-start it. If it's wildly variable, they don't usually stabilize it. If it's growing out of control, they don't brake it.

Whatever cap we put on campaign contributions, it doesn't clean up politics. The Feds fiddling with the interest rate haven't made business cycles go away. (We always forget that during upturns, and are shocked, shocked by the downturns.) Spending more on police doesn't make crime go away.

However, there are critical exceptions. Numbers become leverage points when they go into ranges that kick off one of the items higher on this list. Interest rates or birth rates control the gains around positive feedback loops. System goals are parameters that can make big differences. Sometimes a system gets onto a chaotic edge, where the tiniest change in a number can drive it from order to what appears to be wild disorder.

Probably the most common kind of critical number is the length of delay in a feedback loop. Remember that bathtub on the fourth floor I mentioned, with the water heater in the basement? I actually experienced one of those once, in an old hotel in London. It wasn't even a bathtub with buffering capacity; it was a shower. The water temperature took at least a minute to respond to my faucet twists. Guess what my shower was like. Right, oscillations from hot to cold and back to hot, punctuated with expletives. Delays in negative feedback loops cause oscillations. If you're trying to adjust a system state to your goal, but you only receive delayed information about what the system state is, you will overshoot and undershoot.

Same if your information is timely, but your response isn't. For example, it takes several years to build an electric power plant, and then that plant lasts, say, thirty years. Those delays make it impossible to build exactly the right number of plants to supply a rapidly changing demand. Even with immense effort at forecasting, almost every electricity industry in the world experiences long oscillations between overcapacity and undercapacity. A system just can't respond to short-term changes when it has long-term delays. That's why a massive central-planning system, such as the Soviet Union or General Motors, necessarily functions poorly.

A delay in a feedback process is critical RELATIVE TO RATES OF CHANGE (growth, fluctuation, decay) IN THE SYSTEM STATE THAT THE FEEDBACK LOOP IS TRYING TO CONTROL. Delays that are too short cause overreaction, oscillations amplified by the jumpiness of the response. Delays that are too long cause damped, sustained, or exploding oscillations, depending on how much too long. At the extreme they cause chaos. Delays in a system with a threshold, a danger point, a range past which irreversible damage can occur, cause overshoot and collapse.

Delay length would be a high leverage point, except for the fact that delays are not often easily changeable. Things take as long as they take. You can't do a lot about the construction time of a major piece of capital, or the maturation time of a child, or the growth rate of a forest. It's usually easier to slow down the change rate (positive feedback loops, higher on this list), so feedback delays won't cause so much trouble. Critical numbers are not nearly as common as people seem to think they are. Most systems have evolved or are designed to stay out of sensitive parameter ranges. Mostly, the numbers are not worth the sweat put into them.

8. Material stocks and flows.

The plumbing structure, the stocks and flows and their physical arrangement, can have an enormous effect on how a system operates.

When the Hungarian road system was laid out so all traffic from one side of the nation to the other had to pass through central Budapest, that determined a lot about air pollution and commuting delays that are not easily fixed by pollution control devices, traffic lights, or speed limits. The only way to fix a system that is laid out wrong is to rebuild it, if you can.

Often you can't, because physical building is a slow and expensive kind of change. Some stock-and-flow structures are just plain unchangeable.

The baby-boom swell in the US population first caused pressure on the elementary school system, then high schools and colleges, then jobs and housing, and now we're looking forward to supporting its retirement. Not much to do about it, because five-year-olds become six-year-olds, and sixty-four-year-olds become sixty-five-year-olds predictably and unstoppably. The same can be said for the lifetime of destructive CFC molecules in the ozone layer, for the rate at which contaminants get washed out of aquifers, for the fact that an inefficient car fleet takes ten to twenty years to turn over.

The possible exceptional leverage point here is in the size of stocks, or buffers. Consider a huge bathtub with slow in and outflows. Now think about a small one with fast flows. That's the difference between a lake and a river. You hear about catastrophic river floods much more often than catastrophic lake floods, because stocks that are big, relative to their flows, are more stable than small ones. A big, stabilizing stock is a buffer.

The stabilizing power of buffers is why you keep money in the bank rather than living from the flow of change through your pocket. It's why stores hold inventory instead of calling for new stock just as customers carry the old stock out the door. It's why we need to maintain more than the minimum breeding population of an endangered species. Soils in the eastern US are more sensitive to acid rain than soils in the west, because they haven't got big buffers of calcium to neutralize acid. You can often stabilize a system by increasing the capacity of a buffer. But if a buffer is too big, the system gets inflexible. It reacts too slowly. Businesses invented just-in-time inventories, because occasional vulnerability to fluctuations or screw-ups is cheaper than certain, constant inventory costs—and because small-to-vanishing inventories allow more flexible response to shifting demand.

There's leverage, sometimes magical, in changing the size of buffers. But buffers are usually physical entities, not easy to change.

The acid absorption capacity of eastern soils is not a leverage point for alleviating acid rain damage. The storage capacity of a dam is literally cast in concrete. Physical structure is crucial in a system, but the leverage point is in proper design in the first place. After the structure is built, the leverage is in understanding its limitations and bottlenecks and refraining from fluctutions or expansions that strain its capacity.

7. Regulating negative feedback loops.
Now we're beginning to move from the physical part of the system to the information and control parts, where more leverage can be found. Nature evolves negative feedback loops and humans invent them to keep system states within safe bounds.

A thermostat loop is the classic example. Its purpose is to keep the system state called "room temperature" fairly constant at a desired level. Any negative feedback loop needs a goal (the thermostat setting), a monitoring and signaling device to detect excursions from the goal (the thermostat), and a response mechanism (the furnace and/or air conditioner, fans, heat pipes, fuel, etc.).

A complex system usually has numerous negative feedback loops it can bring into play, so it can self-correct under different conditions and impacts. Some of those loops may be inactive much of the time—like the emergency cooling system in a nuclear power plant, or your ability to sweat or shiver to maintain your body temperature. One of the big mistakes we make is to strip away these emergency response mechanisms because they aren't often used and they appear to be costly. In the short term we see no effect from doing this. In the long term, we narrow the range of conditions over which the system can survive.

One of the most heartbreaking ways we do this is in encroaching on the habitats of endangered species. Another is in encroaching on our own time for rest, recreation, socialization, and meditation.

The "strength" of a negative loop—its ability to keep its appointed stock at or near its goal—depends on the combination of all its parameters and links—the accuracy and rapidity of monitoring, the quickness and power of response, the directness and size of corrective flows.

There can be leverage points here. Take markets, for example, the negative feedback systems that are all but worshiped by economists—and they can indeed be marvels of self-correction, as prices vary to keep supply and demand in balance. The more the price—the central signal to both producers and consumers—is kept clear, unambiguous, timely, and truthful, the more smoothly markets will operate. Prices that reflect full costs will tell consumers how much they can actually afford and will reward efficient producers. Companies and governments are fatally attracted to the price leverage point, of course, all of them pushing in the wrong direction with subsidies, fixes, externalities, taxes, and other forms of confusion. The REAL leverage here is to keep them from doing it. Hence anti-trust laws, truth-in-advertising laws, attempts to internalize costs (such as pollution taxes), the removal of perverse subsidies, and other ways of leveling market playing fields.

The strength of a negative feedback loop is important RELATIVE TO THE IMPACT IT IS DESIGNED TO CORRECT. If the impact increases in strength, the feedbacks have to be strengthened too.

A thermostat system may work fine on a cold winter day—but open all the windows and its corrective power will fail. Democracy worked better before the advent of the brainwashing power of centralized mass communications. Traditional controls on fishing were sufficient until radar spotting and drift nets and other technologies made it possible for a few actors to wipe out the fish. The power of big industry calls for the power of big government to hold it in check; a global economy makes necessary a global government.

Here are some other examples of strengthening negative feedback controls to improve a system's self-correcting abilities: preventive medicine, exercise, and good nutrition to bolster the body's ability to fight disease, integrated pest management to encourage natural predators of crop pests, the Freedom of Information Act to reduce government secrecy, protection for whistle blowers, impact fees, pollution taxes, and performance bonds to recapture the externalized public costs of private benefits.

6. Driving positive feedback loops.
A positive feedback loop is self-reinforcing. The more it works, the more it gains power to work some more.

The more people catch the flu, the more they infect other people. The more babies are born, the more people grow up to have babies. The more money you have in the bank, the more interest you earn, the more money you have in the bank. The more the soil erodes, the less vegetation it can support, the fewer roots and leaves to soften rain and runoff, the more soil erodes. The more high-energy neutrons in the critical mass, the more they knock into nuclei and generate more.

Positive feedback loops drive growth, explosion, erosion, and collapse in systems. A system with an unchecked positive loop ultimately will destroy itself. That's why there are so few of them.

Usually a negative loop kicks in sooner or later. The epidemic runs out of infectable people—or people take increasingly strong steps to avoid being infected. The death rate rises to equal the birth rate—or people see the consequences of unchecked population growth and have fewer babies. The soil erodes away to bedrock, and after a million years the bedrock crumbles into new soil—or people put up check dams and plant trees.

In those examples, the first outcome is what happens if the positive loop runs its course, the second is what happens if there's an intervention to reduce its power.

Reducing the gain around a positive loop—slowing the growth—is usually a more powerful leverage point in systems than strengthening negative loops, and much preferable to letting the positive loop run.

Population and economic growth rates in the world model are leverage points, because slowing them gives the many negative loops, through technology and markets and other forms of adaptation, time to function. It's the same as slowing the car when you're driving too fast, rather than calling for more responsive brakes or technical advances in steering.

The most interesting behavior that rapidly turning positive loops can trigger is chaos. This wild, unpredictable, unreplicable, and yet bounded behavior happens when a system starts changing much, much faster than its negative loops can react to it.

For example, if you keep raising the capital growth rate in the world model, eventually you get to a point where one tiny increase more will shift the economy from exponential growth to oscillation. Another nudge upward gives the oscillation a double beat. And just the tiniest further nudge sends it into chaos.

I don't expect the world economy to turn chaotic any time soon (not for that reason, anyway). That behavior occurs only in unrealistic parameter ranges, equivalent to doubling the size of the economy within a year. Real-world systems do turn chaotic, however, if something in them can grow or decline very fast. Fast-replicating bacteria or insect populations, very infectious epidemics, wild speculative bubbles in money systems, neutron fluxes in the guts of nuclear power plants. These systems are hard to control, and control must involve slowing down the positive feedbacks.

In more ordinary systems, look for leverage points around birth rates, interest rates, erosion rates, "success to the successful" loops, any place where the more you have of something, the more you have the possibility of having more.

5. Information flows.

There was this subdivision of identical houses, the story goes, except that the electric meter in some of the houses was installed in the basement and in others it was installed in the front hall, where the residents could see it constantly, going round faster or slower as they used more or less electricity. Electricity consumption was 30 percent lower in the houses where the meter was in the front hall.

Systems-heads love that story because it's an example of a high leverage point in the information structure of the system. It's not a parameter adjustment, not a strengthening or weakening of an existing loop. It's a NEW LOOP, delivering feedback to a place where it wasn't going before.

In 1986 the US government required that every factory releasing hazardous air pollutants report those emissions publicly. Suddenly everyone could find out precisely what was coming out of the smokestacks in town. There was no law against those emissions, no fines, no determination of "safe" levels, just information. But by 1990 emissions dropped 40 percent. One chemical company that found itself on the Top Ten Polluters list reduced its emissions by 90 percent, just to "get off that list."

Missing feedback is a common cause of system malfunction. Adding or rerouting information can be a powerful intervention, usually easier and cheaper than rebuilding physical structure.

The tragedy of the commons that is exhausting the world's commercial fisheries occurs because there is no feedback from the state of the fish population to the decision to invest in fishing vessels. (Contrary to economic opinion, the price of fish doesn't provide that feedback. As the fish get more scarce and hence more expensive, it becomes all the more profitable to go out and catch them. That's a perverse feedback, a positive loop that leads to collapse.)

It's important that the missing feedback be restored to the right place and in compelling form. It's not enough to inform all the users of an aquifer that the groundwater level is dropping. That could trigger a race to the bottom. It would be more effective to set a water price that rises steeply as the pumping rate exceeds the recharge rate.

Suppose taxpayers got to specify on their return forms what government services their tax payments must be spent on. (Radical democracy!) Suppose any town or company that puts a water intake pipe in a river had to put it immediately DOWNSTREAM from its own outflow pipe. Suppose any public or private official who made the decision to invest in a nuclear power plant got the waste from that plant stored on his/her lawn.

There is a systematic tendency on the part of human beings to avoid accountability for their own decisions. That's why there are so many missing feedback loops—and why this kind of leverage point is so often popular with the masses, unpopular with the powers that be, and effective, if you can get the powers that be to permit it to happen or go around them and make it happen anyway.

4. The rules of the system (incentives, punishments, constraints).
The rules of the system define its scope, boundaries, degrees of freedom. Thou shalt not kill. Everyone has the right of free speech. Contracts are to be honored. The president serves four-year terms and cannot serve more than two of them. Nine people on a team, you have to touch every base, three strikes and you're out. If you get caught robbing a bank, you go to jail.

Mikhail Gorbachev came to power in the USSR and opened information flows (glasnost) and changed the economic rules (perestroika), and look what happened.

Constitutions are strong social rules. Physical laws such as the second law of thermodynamics are absolute rules, if we understand them correctly. Laws, punishments, incentives, and informal social agreements are progressively weaker rules.

To demonstrate the power of rules, I ask my students to imagine different ones for a college. Suppose the students graded the teachers. Suppose you come to college when you want to learn something, and you leave when you've learned it. Suppose professors were hired according to their ability to solve real-world problems, rather than to publish academic papers. Suppose a class got graded as a group, instead of as individuals.

Rules change behavior. Power over rules is real power.

That's why lobbyists congregate when Congress writes laws, and why the Supreme Court, which interprets and delineates the Constitution—the rules for writing the rules—has even more power than Congress.

If you want to understand the deepest malfunctions of systems, pay attention to the rules, and to who has power over them.

That's why my systems intuition was sending off alarm bells as the new world trade system was explained to me. It is a system with rules designed by corporations, run by corporations, for the benefit of corporations. Its rules exclude almost any feedback from other sectors of society. Most of its meetings are closed to the press (no information, no feedback). It forces nations into positive loops, competing with each other to weaken environmental and social safeguards in order to attract corporate investment. It's a recipe for unleashing "success to the succesful" loops.

3. The power of self-organization.
The most stunning thing living systems can do is to change themselves utterly by creating whole new structures and behaviors. In biological systems that power is called evolution. In human economies it's called technical advance or social revolution. In systems lingo it's called self-organization.

Self-organization means changing any aspect of a system lower on this list—adding or deleting new physical structure, adding or deleting negative or positive loops or information flows or rules. The ability to self-organize is the strongest form of system resilience, the ability to survive change by changing.

The human immune system can develop responses to (some kinds of) insults it has never before encountered. The human brain can take in new information and pop out completely new thoughts.

Self-organization seems so wondrous that we tend to regard it as mysterious, miraculous. Economists often model technology as literal manna from heaven—coming from nowhere, costing nothing, increasing the productivity of an economy by some steady percent each year. For centuries people have regarded the spectacular variety of nature with the same awe. Only a divine creator could bring forth such a creation.

In fact the divine creator does not have to produce miracles. He, she, or it just has to write clever RULES FOR SELF-ORGANIZATION. These rules govern how, where, and what the system can add onto or subtract from itself under what conditions.

Self-organizing computer models demonstrate that delightful, mind-boggling patterns can evolve from simple evolutionary algorithms. (That need not mean that real-world algorithms are simple, only that they can be.) The genetic code that is the basis of all biological evolution contains just four letters, combined into words of three letters each. That code, and the rules for replicating and rearranging it, has spewed out an unimaginable variety of creatures.

Self-organization is basically a matter of evolutionary raw material—a stock of information from which to select possible patterns—and a means for testing them. For biological evolution the raw material is DNA, one source of variety is spontaneous mutation, and the testing mechanism is something like punctuated Darwinian selection. For technology the raw material is the body of understanding science has accumulated. The source of variety is human creativity (whatever THAT is) and the selection mechanism is whatever the market will reward or whatever governments and foundations will fund or whatever tickles the fancy of crazy inventors.

When you understand the power of self-organization, you begin to understand why biologists worship biodiversity even more than economists worship technology. The wildly varied stock of DNA, evolved and accumulated over billions of years, is the source of evolutionary potential, just as science libraries and labs and scientists are the source of technological potential. Allowing species to go extinct is a systems crime, just as randomly eliminating all copies of particular science journals, or particular kinds of scientists, would be.

The same could be said of human cultures, which are the store of behavioral repertoires accumulated over not billions, but hundreds of thousands of years. They are a stock out of which social evolution can arise. Unfortunately, people appreciate the evolutionary potential of cultures even less than they understand the potential of every genetic variation in ground squirrels. I guess that's because one aspect of almost every culture is a belief in the utter superiority of that culture.

Any system, biological, economic, or social, that scorns experimentation and wipes out the raw material of innovation is doomed over the long term on this highly variable planet.

The intervention point here is obvious but unpopular. Encouraging diversity means losing control. Let a thousand flowers bloom and ANYTHING could happen!

Who wants that?

2. The goals of the system.
Right there, the push for control, is an example of why the goal of a system is even more of a leverage point than the self-organizing ability of a system.

If the goal is to bring more and more of the world under the control of one central planning system (the empire of Genghis Khan, the world of Islam, the People's Republic of China, Wal-Mart, Disney), then everything further down the list, even self-organizing behavior, will be pressured or weakened to conform to that goal.

That's why I can't get into arguments about whether genetic engineering is a good or a bad thing. Like all technologies, it depends upon who is wielding it, with what goal. The only thing one can say is that if corporations wield it for the purpose of generating marketable products, that is a very different goal, a different direction for evolution than anything the planet has seen so far.

There is a hierarchy of goals in systems. Most negative feedback loops have their own goals—to keep the bath water at the right level, to keep the room temperature comfortable, to keep inventories stocked at sufficient levels. They are small leverage points. The big leverage points are the goals of entire systems.

People within systems don't often recognize what whole-system goal they are serving. To make profits, most corporations would say, but that's just a rule, a necessary condition to stay in the game. What is the point of the game? To grow, to increase market share, to bring the world (customers, suppliers, regulators) more under the control of the corporation, so that its operations become ever more shielded from uncertainty. That's the goal of a cancer cell too and of every living population. It's only a bad one when it isn't countered by higher-level negative feedback loops with goals of keeping the system in balance. The goal of keeping the market competitive has to trump the goal of each corporation to eliminate its competitors. The goal of keeping populations in balance and evolving has to trump the goal of each population to commandeer all resources into its own metabolism.

I said a while back that changing the players in a system is a low-level intervention, as long as the players fit into the same old system. The exception to that rule is at the top, if a single player can change the system's goal.

I have watched in wonder as—only very occasionally—a new leader in an organization, from Dartmouth College to Nazi Germany, comes in, enunciates a new goal, and single-handedly changes the behavior of hundreds or thousands or millions of perfectly rational people.

That's what Ronald Reagan did. Not long before he came to office, a president could say, "Ask not what government can do for you, ask what you can do for the government," and no one even laughed. Reagan said the goal is not to get the people to help the government and not to get government to help the people, but to get the government off our backs. One can argue, and I would, that larger system changes let him get away with that. But the thoroughness with which behavior in the US and even the world has been changed since Reagan is testimony to the high leverage of articulating, repeating, standing for, insisting upon new system goals.

1. The mindset or paradigm out of which the system arises.
Another of Jay Forrester's systems sayings goes: It doesn't matter how the tax law of a country is written. There is a shared idea in the minds of the society about what a "fair" distribution of the tax load is. Whatever the rules say, by fair means or foul, by complications, cheating, exemptions or deductions, by constant sniping at the rules, the actual distribution of taxes will push right up against the accepted idea of "fairness."

The shared idea in the minds of society, the great unstated assumptions—unstated because unnecessary to state; everyone knows them—constitute that society's deepest set of beliefs about how the world works. There is a difference between nouns and verbs. People who are paid less are worth less. Growth is good. Nature is a stock of resources to be converted to human purposes. Evolution stopped with the emergence of Homo sapiens . One can "own" land. Those are just a few of the paradigmatic assumptions of our culture, all of which utterly dumbfound people of other cultures.

Paradigms are the sources of systems. From them come goals, information flows, feedbacks, stocks, flows.

The ancient Egyptians built pyramids because they believed in an afterlife. We build skyscrapers, because we believe that space in downtown cities is enormously valuable. (Except for blighted spaces, often near the skyscrapers, which we believe are worthless.) Whether it was Copernicus and Kepler showing that the earth is not the center of the universe, or Einstein hypothesizing that matter and energy are interchangeable, or Adam Smith postulating that the selfish actions of individual players in markets wonderfully accumulate to the common good.

People who manage to intervene in systems at the level of paradigm hit a leverage point that totally transforms systems.

You could say paradigms are harder to change than anything else about a system, and therefore this item should be lowest on the list, not the highest. But there's nothing physical or expensive or even slow about paradigm change. In a single individual it can happen in a millisecond. All it takes is a click in the mind, a new way of seeing. Of course individuals and societies do resist challenges to their paradigm harder than they resist any other kind of change.

So how do you change paradigms? Thomas Kuhn, who wrote the seminal book about the great paradigm shifts of science, has a lot to say about that. In a nutshell, you keep pointing at the anomalies and failures in the old paradigm, you come yourself, loudly, with assurance, from the new one, you insert people with the new paradigm in places of public visibility and power. You don't waste time with reactionaries; rather you work with active change agents and with the vast middle ground of people who are open-minded.

Systems folks would say one way to change a paradigm is to model a system, which takes you outside the system and forces you to see it whole. We say that because our own paradigms have been changed that way.

0. The power to transcend paradigms.
Sorry, but to be truthful and complete, I have to add this kicker.

The highest leverage of all is to keep oneself unattached in the arena of paradigms, to realize that NO paradigm is "true," that even the one that sweetly shapes one's comfortable worldview is a tremendously limited understanding of an immense and amazing universe.

It is to "get" at a gut level the paradigm that there are paradigms, and to see that that itself is a paradigm, and to regard that whole realization as devastatingly funny. It is to let go into Not Knowing.

People who cling to paradigms (just about all of us) take one look at the spacious possibility that everything we think is guaranteed to be nonsense and pedal rapidly in the opposite direction. Surely there is no power, no control, not even a reason for being, much less acting, in the experience that there is no certainty in any worldview. But everyone who has managed to entertain that idea, for a moment or for a lifetime, has found it a basis for radical empowerment. If no paradigm is right, you can choose one that will help achieve your purpose. If you have no idea where to get a purpose, you can listen to the universe (or put in the name of your favorite deity here) and do his, her, its will, which is a lot better informed than your will.

It is in the space of mastery over paradigms that people throw off addictions, live in constant joy, bring down empires, get locked up or burned at the stake or crucified or shot, and have impacts that last for millennia.

Back from the sublime to the ridiculous, from enlightenment to caveats. There is so much that has to be said to qualify this list. It is tentative and its order is slithery. There are exceptions to every item on it. Having the list percolating in my subconscious for years has not transformed me into a Superwoman. I seem to spend my time running up and down the list, trying out leverage points wherever I can find them. The higher the leverage point, the more the system resists changing it-that's why societies rub out truly enlightened beings.

I don't think there are cheap tickets to system change. You have to work at it, whether that means rigorously analyzing a system or rigorously casting off paradigms. In the end, it seems that leverage has less to do with pushing levers than it does with disciplined thinking combined with strategically, profoundly, madly letting go.