

T O O L B O X R E P R I N T S E R I E S



SYSTEMS THINKING TOOLS

A User's Reference Guide ←



→ BY DANIEL H. KIM

THE TOOLBOX REPRINT SERIES

Systems Archetypes I: Diagnosing Systemic Issues and Designing High-Leverage Interventions

Systems Archetypes II: Using Systems Archetypes to Take Effective Action

Systems Archetypes III: Understanding Patterns of Behavior and Delay

Systems Thinking Tools: A User's Reference Guide

The "Thinking" in Systems Thinking: Seven Essential Skills

Systems Thinking Tools: A User's Reference Guide

by **Daniel H. Kim**

© 1994, 2000 by Pegasus Communications, Inc.

First printing January 1994.

All rights reserved. No part of this book may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording, or by any information storage or retrieval system, without written permission from the publisher.

Acquiring editor: Kellie Wardman O'Reilly

Production: Nancy Daugherty

ISBN 1-883823-02-1



PEGASUS COMMUNICATIONS, INC.

One Moody Street

Waltham, MA 02453-5339 USA

Phone 800-272-0945 / 781-398-9700

Fax 781-894-7175

customerservice@pegasuscom.com / info@pegasuscom.com

www.pegasuscom.com

T A B L E O F C O N T E N T S

Introduction	3
PART I: AN OVERVIEW	
Systems Thinking as a Language <i>by Michael R. Goodman</i>	6
Levels of Understanding: “Fire-Fighting” at Multiple Levels	8
A Palette of Systems Thinking Tools	10
PART II: DYNAMIC THINKING TOOLS	
Reinforcing and Balancing Loops: Building Blocks of Dynamic Systems	14
Balancing Loops with Delays: Teeter-Tottering on Seesaws	16
Guidelines for Drawing Causal Loop Diagrams	18
Systems Archetypes at a Glance	20
PART III: STRUCTURAL THINKING TOOLS	
From Causal Loops to Graphical Functions: Articulating Chaos	24
Graphical Functions: “Seeing” the Full Story	26
Structural Thinking: The World According to Accumulators and Flows	28
Accumulators: Bathtubs, Bathtubs, Everywhere	30
Accumulation Management: Avoiding the “Pack Rat” Syndrome	32
Delays: Accumulators in Disguise	34
S-Shaped Growth and the Law of Diminishing Returns	36
PART IV: COMPUTER-BASED TOOLS	
Modeling for What Purpose? <i>by Jay W. Forrester</i>	40
Management Flight Simulators: Flight Training for Managers—Part I	42
Management Flight Simulators: Flight Training for Managers—Part II	44
Learning Laboratories: Practicing Between Performances	46
PART V: REFERENCE GUIDE	
The Vocabulary of Systems Thinking: A Pocket Guide	50
The Do’s and Don’t’s of Systems Thinking on the Job <i>by Michael R. Goodman</i>	52
Further Reading on the Ten Tools of Systems Thinking	54
Index to <i>THE SYSTEMS THINKER</i>	55



We are in the midst of a changing of an age—from the age of machines to the systems age. Our past was defined by a view of the world as a machine that could be understood by breaking it into smaller and smaller parts. In the machine age view of the world, the parts are what is most important—by understanding each of the parts, we build up our understanding of the larger wholes. In the systems view, it is the whole that is most important—parts in isolation have no meaning in and of themselves. Systems thinking embodies the idea that the interrelationships among parts relative to a common purpose of a system are what is important.

There is nothing more powerful than an idea whose time has come. But ideas without practical tools can take us only so far in making any meaningful changes that will have an impact on the world. Systems thinking provides the ideas that can help us see the world in new ways, as well as the tools that can help us take new actions that are systemic and more effective. This booklet provides a basic introduction to the various tools of systems thinking that have been developed and used over the last 50 years.


ACKNOWLEDGMENTS

Much of this work has been developed over the years through the efforts of many system dynamicists. *Systems Thinking Tools: A User's Reference Guide*, part of the Toolbox Reprint Series, was created and compiled by Kellie Wardman O'Reilly.

THE LANGUAGE OF LINKS AND LOOPS


- ^s A causal link between two variables, where a change in X causes a change in Y in the same direction, or where X adds to Y.
- ^o A causal link between two variables, where a change in X causes a change in Y in the opposite direction, or where X subtracts from Y.
- R A "reinforcing" feedback loop that amplifies change.
- B A "balancing" feedback loop that seeks equilibrium.

BALANCING LOOP EXAMPLE



If there is a gap between the desired level and the actual level, adjustments are made until the actual equals the desired level. The starting variable is grey.

THE LANGUAGE OF ACCUMULATORS



PART I

AN OVERVIEW

Systems thinking can be thought of as a language. As a language, it is a specific way of viewing the world; it affects thought, and thought in turn affects how we look at the world. “Systems Thinking as a Language” (p. 6) offers insight into how systems thinking can be a useful framework for communicating about complex issues.

By “conversing” in the language of feedback loops, we can learn to better articulate the complex interconnections of circular causality in which we live. Learning the language of systems thinking requires us to understand our world on at least four levels—events, pattern of events, systemic structure, and shared vision. “Levels of Understanding: ‘Fire-Fighting’ at Multiple Levels” (p. 8) describes these levels and the specific action mode associated with each one.

This section closes with “A Palette of Systems Thinking Tools” (p. 10), which outlines ten tools of systems thinking. Seven of these tools are covered in the subsequent sections of this booklet.



SYSTEMS THINKING AS A LANGUAGE

BY MICHAEL R. GOODMAN

Language has a subtle, yet powerful effect on the way we view the world. English, like most other Western languages, is linear—its basic sentence construction, noun-verb-noun, translates into a world-view of “x causes y.” This linearity predisposes us to focus on one-way relationships rather than circular or mutually causative ones, where x influences y, and y in turn influences x. Unfortunately, many of the most vexing problems confronting managers and corporations today are caused by a web of tightly interconnected circular relationships. To enhance our understanding and communication of such problems, we need a language more naturally suited to the task.

ELEMENTS OF THE LANGUAGE

Systems thinking can be thought of as a language for communicating about complexities and interdependencies. In particular, the following qualities make systems thinking a useful framework for discussing and analyzing complex issues:

- **Focuses on “closed interdependencies.”**

The language of systems thinking is circular rather than linear. It focuses on closed interdependencies, where x influences y, y influences z, and z influences x.

- **Offers a “visual” language.** Many of the systems thinking tools—causal loop diagrams, behavior over time diagrams, systems archetypes, and structural diagrams—have a strong visual component. They help clarify complex issues by summing up, concisely and clearly, the key elements involved.

Diagrams also facilitate learning. Studies have shown that many people learn best through representational images, such as pictures or stories. A systems diagram is a powerful means of communication because it distills the essence of a problem into a format that can be easily remembered, yet is rich in implications and insights.

A systems diagram is a powerful means of communication because it distills the essence of a problem into a format that can be easily remembered, yet is rich in implications and insights.

- **Adds precision.** The specific set of “syntactical” rules that govern systems diagrams greatly reduce the ambiguities and miscommunications that can occur when we tackle complex issues.

Example: In drawing out the relationships between key aspects of a problem, causal links are not only indicated by arrows, but are labeled “s” (same) or “o” (opposite) to specify how one variable affects another. Such labeling makes the nature of the relationship more precise, ensuring only one possible interpretation.

- **Forces an “explicitness” of mental models.** The systems thinking language translates “war stories” and individual

perceptions of a problem into black-and-white pictures that can reveal subtle differences in viewpoint.

Example: In one systems thinking course, a team of managers was working on an issue they had been wrestling with for months. One manager was explaining his position, tracing through the loops he had drawn, when a team member stopped him. “Does that model represent your thinking about this problem?” he asked.

The presenter hesitated a bit, reviewed his diagram, and finally answered, “Yes.”

The first man, evidently relieved, responded, “After all of these months, I finally *really understand* your thoughts on this issue. I disagree with it, but at least now that we are clear on our different viewpoints, we can work together to clarify the problem.”

- **Allows examination and inquiry.**

Systems diagrams can be powerful means for fostering a collective understanding of a problem. Once individuals have stated their understanding of the problem, they can collaborate on addressing the challenges it poses. And by focusing the discussion on the diagrams, systems thinking defuses much of the defensiveness that can arise in a high-level debate.

Example: When carrying on a systems discussion, differing opinions are no longer viewed as “human resources’ view of our productivity problem” or “marketing’s description of decreasing customer satisfaction,” but simply different structural representations of the system. This shifts the focus of the discussion from whether human resources or marketing is right, to



constructing a diagram that best captures the behavior of the system.

• **Embodies a worldview that looks at wholes, rather than parts**, and that recognizes the importance of understanding how the different segments of a system are interconnected. An inherent assumption of the systems thinking worldview is that problems are internally generated—that we often create our own “worst nightmares.”

Example: At systems thinking courses at Innovation Associates, participants play a board game known as the Beer Game, where they assume the position of retailer, wholesaler, distributor, or producer. Each player tries to achieve a careful balance between carrying too much inventory or being backlogged. When things go wrong, many people blame their supplier (“I kept ordering more, but he didn’t respond”) or the buyers (“fickle consumers—one day they’re buying it by the truckload, the next day they won’t even touch the stuff”). In reality, neither the buyers nor the suppliers are responsible for the wide fluctuations in inventory—they are a natural consequence of the structure of the system in which the players are functioning.

The systems thinking worldview dispels the “us versus them” mentality by expanding the boundary of our thinking. Within the framework of systems thinking, “us” and “them” are part of the same system and thus responsible for both the problems and their solutions.

LEARNING THE LANGUAGE

Learning systems thinking can be likened to mastering a foreign language. In school, we studied a foreign language by first memorizing the essential vocabulary words and verb conjugations. Then we began putting together the pieces into simple sentences. In the language of systems thinking, systems diagrams such as causal loops can be thought of as sentences con-

structed by linking together key variables and indicating the causal relationships between them. By stringing together several loops, we can create a “paragraph” that tells a coherent story about a particular problem under study.

If there were a Berlitz guide to systems thinking, archetypes such as “Fixes That Fail” or “Shifting the Burden” would be listed as “commonly used phrases.” They provide a ready-made library of common structures and behaviors that can apply to many situations. Memorizing them can help you recognize a business situation or problem that is exhibiting common symptoms of a systemic breakdown.

An inherent assumption of the systems thinking worldview is that problems are internally generated—we often create our own “worst nightmares.”

Of course, the key to becoming more proficient in any language is to practice—and practice often. When reading a newspaper article, for example, try to “translate” it into a systems perspective:

- Take events reported in the newspaper and try to trace out an underlying pattern that is at work.
- Check whether the story fits one of the systems archetypes, or whether it is perhaps a combination of several archetypes.
- Try to sketch out a causal loop or two that captures the structure producing that pattern.

Don’t expect to be fluent in systems thinking right away. Remember, after your first few Latin classes, you still couldn’t read *The Odyssey*. For that matter, you probably knew only a few key phrases and vocabulary words, but you improved your

skills by using the language as often as possible. The same holds true for systems thinking.

When sitting in a meeting, see if you can inform your understanding of a problem by applying a systems perspective. Look for key words that suggest linear thinking is occurring—statements such as “we need more of the same” or “that solution worked for us the last time this happened, why not use it again?” You can also create low-key practice sessions by working with a small team of colleagues to diagram a particular problem or issue.

BECOMING FLUENT

We say someone is fluent when they begin to think in a particular language and no longer have to translate. But fluency means more than just an ability to communicate in a language; it means understanding the surrounding culture of the language—the worldview. As with any foreign language, mastering systems thinking will allow us to fully engage in and absorb the worldview that pervades it. By learning the language of systems thinking, we will hopefully change not only the way we discuss complex issues, but the way we think about them as well. ■

Michael Goodman is an associate director of Innovation Associates, Inc. (Cambridge, MA). The material in this article was drawn from his 20 years of experience in the field, as well as from business courses developed by Innovation Associates.



LEVELS OF UNDERSTANDING: “FIRE-FIGHTING” AT MULTIPLE LEVELS

It's another busy night in the hospital emergency room. Several car accident victims have been rushed into surgery, one little boy is having a broken arm set, a drug overdose victim is being treated, and numerous other people fill the chairs in the waiting room. Each night is different, and yet each one is also the same. The doctors and nurses must act fast to treat the most seriously injured, while the others wait their turn. Like an assembly line of defective parts, patients are diagnosed, treated, and released. Each injury is a crisis that demands immediate attention.

So what's wrong with this picture? After all, isn't this what emergency rooms are meant to do? The answer depends on the level of understanding at which we are looking at the situation.

LEVELS OF UNDERSTANDING

There are multiple levels from which we can view and understand the world. From a systemic perspective, we are interested in

four distinct levels—events, patterns of events, systemic structure, and shared vision (see “Levels of Understanding”). Events are the things we encounter on a day-to-day basis: a machine breaks, it rains, we eat dinner, see a movie, or write a report. Patterns of events are the accumulated memories of events—when strung together in a series over time, they reveal recurring patterns. Systemic structure can be viewed as “event generators” because they are responsible for producing the events. Similarly, shared vision can be viewed as “systemic structure generators” because they are the guiding force behind the creation or change of all kinds of structures.

We live in an event-oriented world, and our language is rooted at the level of events. At work, we encounter a series of events, which often appear in the form of problems that we must “solve.” Our solutions, however, may be short-lived, and the symptoms can eventually return as seemingly new problems (see “Using ‘Fixes That Fail’ to Get off the Problem-Solving Treadmill,”

THE SYSTEMS THINKER, V3N7). This is consistent with our evolutionary history, which was geared toward responding to those things that posed an immediate danger to our well-being.

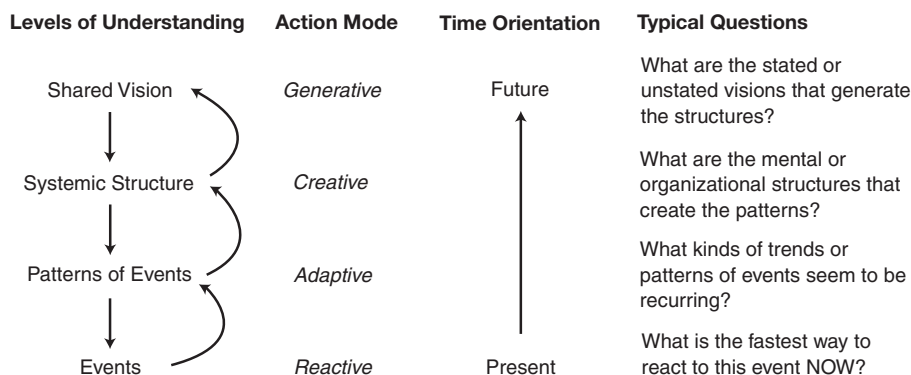
Events require an immediate response. If a house is burning, we react by taking action to put out the fire. Putting out the fire is appropriate, but if it is the only action that is ever taken, it is inadequate from a systemic perspective. Why? Although it solved the immediate problem (the burning house), it has done nothing to alter the fundamental structure that caused that event (e.g., inadequate building codes, lack of fire detectors, fire prevention education). The “Levels of Understanding” diagram and framework can help us go beyond typical event-orientation responses and begin to look for higher leverage actions.

FROM FIRE-FIGHTING TO FIRE PREVENTION

At the event level, if a house is on fire, all we can do is react as quickly as possible to put the fire out. The only mode of action that is appropriate and available is to be reactive. If we reacted to fires only at the events level, we would put all of our energy into fighting fires—and we would probably have a lot more fire stations than we do today.

If we look at the problem of fires at the pattern of events level, we can begin to anticipate where they are more likely to occur. We may notice that certain neighborhoods seem to have more fires than others. We are able to be adaptive by locating more fire stations in those areas, and

LEVELS OF UNDERSTANDING





staffing them accordingly (based on past patterns of usage). Since the stations are a lot closer, we can be more effective at putting out fires by getting to them sooner. Yet while being adaptive allows us to be more effective fire-fighters, it does nothing to reduce the actual occurrence of fires.

At the systemic structure level we begin asking questions: “Are smoke detectors being used? What kinds of building materials are less flammable? What safety features reduce fatalities?” Actions taken at this level can actually reduce the number and severity of fires. Establishing fire codes with requirements such as automatic sprinkler systems, fireproof materials, fire walls, and fire alarm systems saves lives by preventing or containing fires. Actions taken at this level are creative because they help create a different future.

Systemic structure includes not only the organizational structures and physical buildings, but people’s mental models and habits as well. Where do the systemic structures come from? They are usually a reflection of a shared vision of what is valued or desired. In the case of fire-fighting, the new structures (e.g., fire codes) are born out of a shared value of the importance of protecting human lives, combined with the desire to live and work in safe buildings. At the level of shared vision, our actions can be generative, bringing something into being that did not exist before. We begin asking questions like “What’s the role of the fire-fighting function in this community? What are the trade-offs we are willing to make as a community between the amount of resources devoted to fire-fighting compared to other things?”

It is important to remember that the process of gaining deeper understanding is not a linear one. Our understanding of a situation at one level can feed back and inform our awareness at another level. Events and patterns of events, for example, can cause us to change systemic structures

and can also challenge our shared vision. To be most effective, the full range of levels must be considered simultaneously. The danger lies in operating at any one level to the exclusion of the others.

Our ability to influence the future does increase, however, as we move from the level of events to shared vision. Does this mean that high-leverage actions can only be found at higher levels? No, because leverage is a relative concept, not an absolute. When someone is bleeding, the highest leverage action at that moment is to stop

The process of gaining deeper understanding is not a linear one. Our understanding of a situation at one level can feed back and inform our awareness at another level.

the bleeding—any other action would be inappropriate. As we move up the levels from events to shared vision, the focus moves from being present-oriented to being future-oriented. Consequently, the actions we take at the higher levels have more impact on future outcomes, not present events.

BACK AT THE EMERGENCY ROOM

While the emergency room (ER) offers a graphic example of a situation in which people must be focused on the present, it also reveals the limitations of the events-oriented response. ER treatment offers maximal leverage to affect the present situation with each patient, but it provides very little leverage for changing the future. If we go up one level and examine ER use from a patterns of events level, we may discover that certain areas of a city seem to have higher emergency room needs. We

may try an adaptive response and increase ER capacity in those regions. If diversion rates are high, we can also find out where the ambulances are being diverted from and try to enhance capacity there.

At the systemic structure level, we can begin to explore why certain regions have an increased need for ERs. We may discover, for example, that 40 percent of the ER admissions are children who are poisoned, because a large percentage of the community cannot read English and all warning labels are printed in English. By redrawing the boundary of the ER issue to include the community, we can take actions that will change the inflow of patients. Electrical utilities have been doing this for some time. Instead of building another expensive power plant to supply more power, they are working with customers to reduce the demand for power.

At a community-wide level, we may want to explore the question, “What is our shared vision of the role our healthcare system plays in our lives?” Perhaps the resources that are going into ERs could be better utilized elsewhere, such as community education and prevention programs. The highest leverage lies in clarifying the quality of life we envision for ourselves, and then using that as a guide for creating the systemic structures that will help us achieve that vision.

The basic message of the “Levels of Understanding” diagram is the importance of recognizing the level at which you are operating, and evaluating whether or not it provides the highest leverage for that situation. Each level offers different opportunities for high-leverage action, but they also have their limits. The challenge is to choose the appropriate response for the immediate situation and find ways to alter the future occurrence of those events. ■



A PALETTE OF SYSTEMS THINKING TOOLS

There is a full array of systems thinking tools that you can think of in the same way as a painter views colors—many shades can be created out of three primary colors, but having a full range of ready-made colors makes painting much easier.

There are at least 10 distinct types of systems thinking tools (an abbreviated summary diagram appears on the facing page). They fall under four broad categories: brainstorming tools, dynamic thinking tools, structural thinking tools, and computer-based tools. Although each of the tools is designed to stand alone, they also build upon one another and can be used in combination to achieve deeper insights into dynamic behavior.

BRAINSTORMING TOOLS

The Double-Q (QQ) Diagram is based on what is commonly known as a fishbone or cause-and-effect diagram. The Qs stand for *qualitative* and *quantitative*, and the technique is designed to help participants begin to see the whole system. During a structured brainstorming session with the QQ diagram, both sides of an issue remain equally visible and properly balanced, avoiding a “top-heavy” perspective. The diagram also provides a visual map of the key factors involved. Once those factors are pinpointed, Behavior Over Time Diagrams and/or Causal Loop Diagrams can be used to explore how they interact.

A QQ diagram begins with a heavy horizontal arrow that points to the issue being addressed. Major “hard” (quantitative) factors branch off along the top and “soft” (qualitative) factors run along the bottom. Arrows leading off of the major factors represent sub-factors, which can in

turn have *sub* sub-factors. Many layers of nesting, however, may be a sign that one of the sub-factors should be turned into a major factor.

DYNAMIC THINKING TOOLS

Behavior Over Time (BOT) Diagrams are more than simple line projections—they capture the dynamic relationships among variables. For example, say we were trying to project the relationship between sales, inventory, and production. If sales jump 20 percent, production cannot jump instantaneously to the new sales number. In addition, inventory must drop below its previous level while production catches up with sales. By sketching out the behavior of different variables on the same graph, we can gain a more explicit understanding of how these variables interrelate.

Causal Loop Diagrams (CLDs) provide a useful way to represent dynamic interrelationships. CLDs make explicit one’s understanding of a system’s structure, provide a visual representation to help communicate that understanding, and capture complex systems in a succinct form. CLDs can be combined with BOTs to form structure-behavior pairs, which provide a rich framework for describing complex dynamic phenomena. CLDs are the systems thinker’s equivalent of the painter’s primary colors.

Systems Archetypes is the name given to certain common dynamics that seem to recur in many different settings. These archetypes, consisting of various combinations of balancing and reinforcing loops, are the systems thinker’s “paint-by-numbers” set—users can take real-world examples and fit them into the appropriate archetype.










They serve as a starting point from which one can build a clearer articulation of a business story or issue. Specific archetypes include: “Drifting Goals,” “Shifting the Burden,” “Limits to Success,” “Success to the Successful,” “Fixes That Fail,” “Tragedy of the Commons,” “Growth and Underinvestment,” and “Escalation” (see “Systems Archetypes at a Glance,” p. 20).

STRUCTURAL THINKING TOOLS

Graphical Function Diagrams, Structure-Behavior Pairs, and Policy Structure Diagrams can be viewed as the building blocks for computer models. Graphical Functions are useful for clarifying nonlinear relationships between variables. They are particularly helpful for quantifying the effects of variables that are difficult to measure, such as morale or time pressure. Structure-Behavior Pairs link a specific structure with its corresponding behavior. Policy Structure Diagrams represent the processes that drive policies. In a sense, when we use these tools we are moving from painting on canvas to sculpting three-dimensional figures.

COMPUTER-BASED TOOLS

This class of tools, including computer models, management flight simulators, and learning laboratories, demands the highest level of technical proficiency to create. On the other hand, very little advance training is required to use them once they are developed. ■

DYNAMIC THINKING TOOLS	STRUCTURAL THINKING TOOLS	COMPUTER-BASED TOOLS
<p>Behavior Over Time Diagram</p>  <p>Can be used to graph the behavior of variables over time and gain insights into any interrelationships between them. (BOT diagrams are also known as reference mode diagrams.)</p>	<p>Graphical Function Diagram</p>  <p>Captures the way in which one variable affects another, by plotting the relationship between the two over the full range of relevant values.</p>	<p>Computer Model</p>  <p>Lets you translate all relationships identified as relevant into mathematical equations. You can then run policy analyses through multiple simulations.</p>
<p>Causal Loop Diagram</p>  <p>Used in conjunction with behavior over time diagrams, can help you identify reinforcing (R) and balancing (B) processes.</p>	<p>Structure-Behavior Pair</p>  <p>Consists of the basic dynamic structures that can serve as building blocks for developing computer models (for example, exponential growth, delays, smooths, S-shaped growth, oscillations, and so on).</p>	<p>Management Flight Simulator</p>  <p>Provides “flight training” for managers through the use of interactive computer games based on a computer model. Users can recognize long-term consequences of decisions by formulating strategies and making decisions based on those strategies.</p>
<p>Systems Archetype</p>  <p>Helps you recognize common system behavior patterns such as “Drifting Goals,” “Shifting the Burden,” “Limits to Growth,” “Fixes That Fail,” and so on—all the compelling, recurring “stories” of organizational dynamics.</p>	<p>Policy Structure Diagram</p>  <p>A conceptual map of the decision-making process embedded in the organization. Focuses on the factors that are weighed for each decision, and can be used to build a library of generic structures.</p>	<p>Learning Laboratory</p>  <p>A manager’s practice field. Is equivalent to a sports team’s experience, which blends active experimentation with reflection and discussion. Uses all the systems thinking tools, from behavior over time diagrams to MFSs.</p>

PART II

DYNAMIC THINKING TOOLS

BEHAVIOR OVER TIME DIAGRAM

CAUSAL LOOP DIAGRAM

SYSTEMS ARCHETYPES

As discussed in “Levels of Understanding: ‘Fire-fighting’ at Multiple Levels” on p. 8, we need to develop our capability to see beyond the event-to-event view of the world. The dynamic thinking tools provide the means to represent the patterns of events that occur over time and also map the structures that are producing those dynamics. This section begins by discussing reinforcing and balancing loops, which are the fundamental building blocks that help us represent the feedback loop structures responsible for generating the dynamic patterns that we observe.

Although the basic concept of reinforcing and balancing loops is simple, actually mapping out one’s own issues in a free-form causal loop diagramming session requires a fair amount of skill. “Guidelines for Drawing Causal Loop Diagrams” (p. 18) can give you some heuristics to follow in trying to construct your own diagrams.

Even with experience, it can be rather daunting to stare at a blank page and try to construct a systemic picture of your issue from scratch. This is where systems archetypes can be very helpful in providing the initial story line from which to elicit understanding of an issue. The archetypes represent generic story lines and structures that have been found to be prevalent in our systems. “Systems Archetypes at a Glance” (p. 20) offers descriptions and guidelines for each of the archetypes (for a more complete coverage of all the archetypes, see *Systems Archetypes I: Diagnosing Systemic Issues and Designing High-Leverage Interventions*, also published by Pegasus Communications). By developing an understanding of each of the archetypes, we can begin to enrich our intuitive sense of how complex structures work.



REINFORCING AND BALANCING LOOPS: BUILDING BLOCKS OF DYNAMIC SYSTEMS

In the book *The Double Helix*, James Watson describes the process through which he and Robert Crick “cracked” the DNA code. While other researchers were searching for complex structures to explain the diversity of life

forms, Watson and Crick explored more simple geometrical designs. They eventually received the Nobel Prize for revealing the double helix structure that is the genetic basis for all life. Through their research, Watson and Crick proved that

the infinite variations we see in nature can all be produced by one simple, elegant structure.

Similarly, two basic loops—reinforcing and balancing—can be seen as the equivalent building blocks of complex social and economic systems. These simple structures combine in an infinite variety of ways to produce the complex systems that we as managers are expected to control.


REINFORCING LOOPS: ENGINES OF GROWTH AND DECAY

Reinforcing loops produce both growth and decay. That is, they compound change in one direction with even more change. For example, in the “Employee-Supervisor Reinforcing Loop” diagram, encouragement from the supervisor is capable of producing good employee performance—that is, as the supervisor demonstrates supportive behavior, the employee’s performance will improve, which will lead the supervisor to be even more supportive. At the same time, unsupportive behavior can produce poor employee performance over time—if the supervisor is not supportive, performance will likely decrease, leading the supervisor to be even less supportive. The same loop can create either kind of reinforcing cycle.

BALANCING LOOPS: GOAL-SEEKING PROCESSES


Of course, most things in life cannot continue growing forever. There are other forces—balancing loops—that resist fur-

EMPLOYEE-SUPERVISOR REINFORCING LOOP



Reinforcing loops compound change in one direction with even more change. For example, encouragement can enhance an employee’s performance, while critical or unsupportive behavior can lead to poor employee performance over time.

INVENTORY CONTROL BALANCING LOOP



Balancing loops try to bring a system to a desired state and keep it there. In an inventory control system, the desired inventory is maintained by adjusting the actual inventory whenever there is too much or too little.



ther increases in a given direction. Balancing loops try to bring things to a desired state and keep them there, much like a thermostat regulates the temperature in a house.

An equivalent example in manufacturing involves maintaining buffer inventory levels between production stages. In this situation, there is a desired inventory level that is maintained by adjusting the actual inventory whenever there is too much or too little (see “Inventory Control Balancing Loop”).

USING THE BUILDING BLOCKS

To see how these two basic loops can combine to form more complex structure-behavior pairs, let’s revisit the employee-supervisor feedback loop. Clearly the employee’s performance will not improve indefinitely just because the supervisor is supportive. The employee may have been putting in longer hours in order to continue impressing the supervisor. Over a period of time, the increased work hours may begin to wear down the employee’s energy level (see “Reinforcing Loop Coupled with a Balancing Loop”

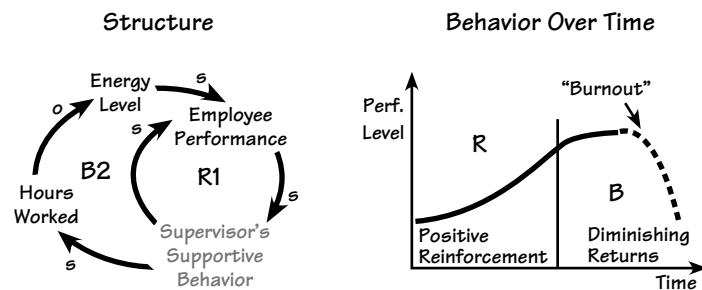
diagram). If this continues, at some point the supervisor’s supportive behavior will be eclipsed by the sheer energy drain of working long hours. Improved performance will gradually be offset by the effects of burnout, until finally the balancing loop connecting energy level and hours worked becomes dominant. At this point the employee’s performance will either plateau or decline.

SUMMARY POINTS

All complex dynamic behavior is produced by two loops: reinforcing and balancing. Behind every growth or decay is at least one reinforcing loop. For every goal-seeking behavior, there is a balancing loop.

A period of growth followed by a slowdown in growth is usually caused by a shift in dominance from a reinforcing to a balancing loop. ■

REINFORCING LOOP COUPLED WITH A BALANCING LOOP



Reinforcing and balancing loops can be combined to describe more complex behavior. For example, encouragement by the supervisor could lead the employee to work longer and longer hours in order to continue impressing the supervisor, eventually leading to burnout and a decrease in performance.



BALANCING LOOPS WITH DELAYS: TEETER-TOTTERING ON SEESAWS

Most of us have played on a seesaw at one time or another and can recall the up and down motion as the momentum shifted from one end to the other. The more equal the weights of both people, the smoother the ride. At a very basic level, a free market economy is a lot like a seesaw with supply at one end and demand on the other end. Prices indicate the imbalance between the two, like a needle positioned at the pivot point of the seesaw.

The goal of a seesaw ride is to always keep things in a state of imbalance (it would be pretty boring to sit on a perfectly balanced one). But the goal in the marketplace is exactly the opposite—to bring supply in balance with demand. Unfortunately, the supply and demand balancing process feels a lot more like a seesaw ride than a smooth adjustment to a stable equilibrium. As shown in a causal loop diagram, the dynamics of this adjustment

process are produced by two balancing loops that try to stabilize on a particular price. But the process is complicated by the presence of significant delays (see “Supply and Demand”).

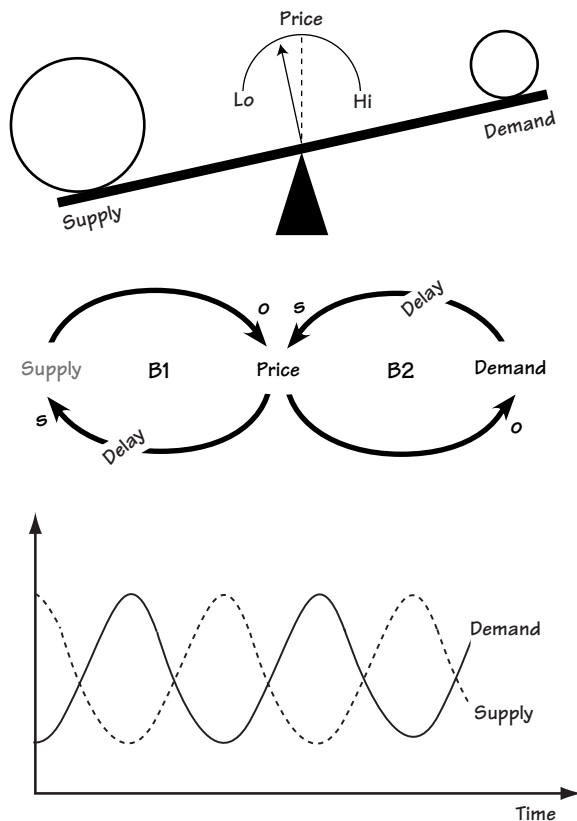
BALANCING SUPPLY AND DEMAND

Tracing through the loops you can see that if demand rises, price tends to go up (all else remaining the same), and as price goes up, demand tends to go down (Beanie Babies notwithstanding). If there is enough inventory or capacity in the system to absorb the increased demand, prices may not go up immediately. As demand outstrips supply, however, price will rise.

On the supply side of the seesaw, an increase in price provides a profit incentive for firms to produce more. Of course, it takes time for firms to expand. The length of the delay depends on how close they already are to full capacity and how quickly they can add capacity to produce more. Hiring new workers may take only a few days, while obtaining additional capital equipment or factory floor space may take months or even years. While firms are making supply adjustments, the gap between supply and demand widens and price goes even higher. The higher price spurs companies to increase their production plans even more.

As supply eventually expands and catches up with demand, price begins to fall. By this time, firms have overexpanded their production capacity and supply overshoots demand, causing price to fall. When the price falls low enough, the product becomes more attractive again and

SUPPLY AND DEMAND



A free market economy is a lot like a seesaw with supply at one end and demand on the other. The dynamics that result from trying to balance supply and demand are produced by two balancing loops that try to stabilize on a particular price. Due to the presence of significant delays, a cycle of overshoot and collapse occurs.



demand picks up, starting the cycle all over again.

AIRPLANES ON SEESAWS

The supply and demand seesaw is played out in all but the most tightly regulated markets. A good example of this balancing act is described in a *Forbes* article entitled “Fasten Seat Belts, Please” (April 2, 1990), about airplane leasing companies.

Leasing companies, which account for roughly 20 percent of all commercial jet aircraft currently on order, have enjoyed enormous profits during booms in air travel. At one time, one carrier alone put in an order to lease 500 planes. Based on leasing and buying rates in the industry, the total number of airplanes was expected to increase by 50 percent between 1990 and 1995. But in the meantime, air traffic growth had slowed in the late 1980s. The leasing companies, however, did not seem too worried.

According to the article, “eight years of unbroken prosperity have created the illusion that many cyclical businesses aren’t cyclical any longer.” But, as one airline executive warned, “This is a cyclical business. Always has been, always will be. With a small change in load factor, the airlines can go from spilling cash to bleeding red ink like the Mississippi River going through the delta.”

If you draw out a causal loop diagram of this industry operating in this way, you see the same supply and demand structure at work. An increase in air traffic growth

fueled a strong demand for airplanes. That in turn sparked an increase in airplane lease rates as airlines scrambled for additional airplanes. The high lease rates led to increased profits and a surge in airplane orders. Since airplanes take many months to build, the supply of leasable airplanes did not adjust right away, making lease rates go even higher. This led to higher profits, which attracted more capital, which was then plowed into even more orders for airplanes.


As the supply catches up to demand, however, the airplane lease rates will fall (the slowing of air traffic growth will accelerate this process). With so many airplanes in the pipeline, the supply will likely begin to outstrip demand and drive lease rates down even further. This puts a squeeze on profits and force marginal firms out of business. Some orders will be canceled; others will be renegotiated.

This example makes it clear that pieces of the airline leasing industry have operated within a seesaw structure. Although the extended period of air traffic growth kept demand ahead of supply for several years, it did not change the nature of the delays in the supply line. Whenever the supply adjustments bring the seesaw back down, airplane leasing companies will be in for a bumpy landing.

SUMMARY

The balancing loop with delay structure is at once simple and complex: simple, because it seems to be an innocuous single loop structure that is easy to comprehend; complex because the resulting behavior is neither simple nor easily predictable. The delays in a typical system are rarely consistent or well known in advance, and the cumulative effects are usually beyond the control of any one person or firm. ■

AIRPLANE LEASING INDUSTRY



A causal loop diagram of the airplane leasing industry shows the same seesaw structure at work.



GUIDELINES FOR DRAWING CAUSAL LOOP DIAGRAMS

The old adage “if the only tool you have is a hammer, everything begins to look like a nail” can also apply to language. If our language is linear and static, we will tend to view and interact with our world as if it were linear and static. Taking a complex, dynamic, and circular world and linearizing it into a set of snapshots may make things seem simpler, but we may totally misread the very reality we were seeking to understand. Making such inappropriate simplifications “is like putting on your brakes and then looking at your speedometer to see how fast you were going,” says Bill Isaacs of *DIA•logos*.

ARTICULATING REALITY

Causal loop diagrams provide a language for articulating our understanding of the dynamic, interconnected nature of our world. We can think of them as sentences that are constructed by linking together key variables and indicating the causal relationships between them. By stringing together several loops, we can create a coherent story about a particular problem or issue.

Following are some more general guidelines that should help lead you through the process:

- **Theme selection.** Creating causal loop diagrams is not an end unto itself, but part of a process of articulating and communicating deeper insights about complex issues. It is pointless to begin creating a causal loop diagram without having selected a theme or issue that you wish to understand better. “To understand the implications of changing from a technology-driven to a marketing-oriented strategy,” for example, is a better theme than “To better understand our strategic planning process.”

- **Time horizon.** It is also helpful to determine an appropriate time horizon for the issue—one long enough to see the dynamics play out. For a change in corporate strategy, the time horizon may span several years, while a change in advertising campaigns may be on the order of months.

Time itself should not be included as a causal agent, however. After a heavy rainfall, a river level steadily rises over time, but we would not attribute it to the passage of time. You need to identify what is actually driving the change. In computer chips, \$/MIPS (million instructions per second) decreased in a straight line in the 1990s. It would be incorrect, however, to draw a causal connection between time and \$/MIPS. Instead, increasing investments and learning curve effects were likely causal forces.

- **Behavior over time charts.** Identifying and drawing out the behavior over time of key variables is an important first step toward articulating the current understanding of the system. Drawing out future behavior means taking a risk—the risk of being wrong. The fact is, any projection of the future will be wrong, but by making it explicit, we can test our assumptions and uncover inconsistencies that may otherwise never get surfaced. For example, drawing projections of steady productivity growth while training dollars are shrinking raises the question, “If training is not driving our growth, what will?” The behavior over time diagram also points out key variables that should be included in the diagram, such as Training Budget and Productivity. Your diagram should try to capture the structure that will produce the projected behavior.

- **Boundary issue.** How do you know when to stop adding to your diagram? If you don’t stay focused on the issue, you may quickly find yourself overwhelmed by the number of connections possible. Remember, you are not trying to draw out the whole system—only what is critical to the theme being addressed. When in doubt, ask, “If I were to double or halve this variable, would it have a significant effect on the issue I am mapping?” If not, it probably can be omitted.

- **Level of aggregation.** How detailed should the diagram be? Again, the level should be determined by the issue itself. The time horizon also can help determine how detailed the variables need to be. If the time horizon is on the order of weeks (fluctuations on the production line), variables that change slowly over a period of many years may be assumed to be constant (such as building new factories). As a rule of thumb, the variables should not describe specific events (a broken pump); they should represent patterns of behavior (pump breakdowns throughout the plant).

- **Significant delays.** Make sure to identify which (if any) links have significant delays relative to the rest of the diagram. Delays are important because they are often the source of imbalances that accumulate in the system. It may help to visualize pressures building up in the system by viewing the delay connection as a relief valve that either opens slowly as pressure builds or opens abruptly when the pressure hits a critical value. An example of this might be a delay between long work hours and burnout: After sustained periods of working 60+ hours per week, a sudden collapse might occur in the form of burnout. ■



	GUIDELINE	EXAMPLE
SELECTING VARIABLE NAMES	<p>1 Use nouns when choosing a variable name. Avoid verbs and action phrases, because the action is conveyed in the loop's arrows. For example, "Costs" is better than "Increasing Costs," because a decrease in Increasing Costs is confusing. The sign of the arrow ("s" for same or "o" for opposite) indicates whether Costs increase or decrease relative to the other variable.</p>	
	<p>2 Use variables that represent quantities that can vary over time. It does not make sense to say that "State of Mind" increases or decreases. A term like "Happiness," on the other hand, can vary.</p>	
	<p>3 Whenever possible, choose the more "positive" sense of a variable name. For example, the concept of "Growth" increasing or decreasing is clearer than an increase or decrease in "Contraction."</p>	
LOOP CONSTRUCTION	<p>4 Think of the possible unintended consequences as well as the expected outcomes for every course of action included in the diagram. For example, an increase in "Production Pressure" may increase "Production Output," but it may also increase "Stress" and decrease "Quality."</p>	
	<p>5 All balancing loops are goal-seeking processes. Try to make explicit the goals driving the loop. For example, Loop B1 may raise questions as to why increasing "Quality" would lead to a decrease in "Actions to Improve Quality." By explicitly identifying "Desired Quality" as the goal in Loop B2, we see that the "Gap in Quality" is really driving improvement actions.</p>	
	<p>6 Distinguishing between perceived and actual states, such as "Perceived Quality" versus "Actual Quality," is important. Perceptions often change slower than reality does, and mistaking the perceived status for current reality can be misleading and create undesirable results.</p>	
	<p>7 If a variable has multiple consequences, start by lumping them into one term while completing the rest of the loop. For example, "Coping Strategies" can represent many different ways we respond to stress (exercise, meditation, alcohol use, etc.).</p>	
	<p>8 Actions almost always have different long-term and short-term consequences. Draw larger loops as they progress from short- to long-term processes. Loop B1 shows the short-term behavior of using alcohol to combat stress. Loop R2, however, draws out the long-term consequences of this behavior, showing that it actually <i>increases</i> stress.</p>	
GENERAL TIPS	<p>9 If a link between two terms requires a lot of explanation to be clear, redefine the variables or insert an intermediate term. Thus, the relationship between "Demand" and "Quality" may be more obvious when "Production Pressure" is inserted between them.</p>	
	<p>10 A shortcut to determining whether a loop is balancing or reinforcing is to count the number of "o's" in the loop. An odd number of "o's" indicates a balancing loop (i.e., an odd number of U-turns keeps you headed in the opposite direction); an even number or no "o's" means it is a reinforcing loop. CAUTION: After labeling the loop, you should always read through it to make sure the story agrees with your R or B label.</p>	



SYSTEMS ARCHETYPES AT A GLANCE

ARCHETYPE	DESCRIPTION	GUIDELINES
<p>Drifting Goals</p>	<p>In a “Drifting Goals” archetype, a gap between the goal and current reality can be resolved by taking corrective action (B1) or lowering the goal (B2). The critical difference is that lowering the goal immediately closes the gap, whereas corrective actions usually take time. (See <i>The Systems Thinker</i>, October 1990.)</p>	<ul style="list-style-type: none"> • Drifting performance figures are usually indicators that the “Drifting Goals” archetype is at work and that real corrective actions are not being taken. • A critical aspect of avoiding a potential “Drifting Goals” scenario is to determine what drives the setting of the goals. • Goals located outside the system will be less susceptible to drifting goals pressures.
<p>Escalation</p>	<p>In the “Escalation” archetype, one party (A) takes actions that are perceived by the other as a threat. The other party (B) responds in a similar manner, increasing the threat to A and resulting in more threatening actions by A. The reinforcing loop is traced out by following the outline of the figure-8 produced by the two balancing loops. (See <i>The Systems Thinker</i>, November 1991.)</p>	<p>To break an escalation structure, ask the following questions:</p> <ul style="list-style-type: none"> • What is the relative measure that pits one party against the other and can you change it? • What are the significant delays in the system that may distort the true nature of the threat? • What are the deep-rooted assumptions that lie beneath the actions taken in response to the threat?
<p>Fixes That Fail</p>	<p>In a “Fixes That Fail” situation, a problem symptom cries out for resolution. A solution is quickly implemented that alleviates the symptom (B1), but the unintended consequences of the “fix” exacerbate the problem (R2). Over time, the problem symptom returns to its previous level or becomes worse. (See <i>The Systems Thinker</i>, November 1990.)</p>	<ul style="list-style-type: none"> • Breaking a “Fixes that Fail” cycle usually requires acknowledging that the fix is merely alleviating a symptom, and making a commitment to solve the real problem <i>now</i>. • A two-pronged attack of applying the fix and planning out the solution will help ensure that you don’t get caught in a perpetual cycle of solving yesterday’s “solutions.”
<p>Growth and Underinvestment</p>	<p>In a “Growth and Underinvestment” archetype, growth approaches a limit that can be eliminated or pushed into the future if capacity investments are made. Instead, performance standards are lowered to justify underinvestment, leading to lower performance which further justifies underinvestment. (See <i>The Systems Thinker</i>, June/July 1992.)</p>	<ul style="list-style-type: none"> • Dig into the assumptions which drive capacity investment decisions. If past performance dominates as a consideration, try to balance that perspective with a fresh look at demand and the factors that drive its growth. • If there is potential for growth, build capacity in anticipation of future demand.



ARCHETYPE	DESCRIPTION	GUIDELINES
<p style="text-align: center;">Limits to Success</p>	<p>In a “Limits of Success” scenario, continued efforts initially lead to improved performance. Over time, however, the system encounters a limit which causes the performance to slow down or even decline (B2), even as efforts continue to rise. (See <i>The Systems Thinker</i>, December 1990/January 1991.)</p>	<ul style="list-style-type: none"> • The archetype is most helpful when it is used well in advance of any problems, to see how the cumulative effects of continued success might lead to future problems. • Use the archetype to explore questions such as What kinds of pressures are building up in the organization as a result of the growth? • Look for ways to relieve pressures or remove limits <i>before</i> an organizational gasket blows.
<p style="text-align: center;">Shifting the Burden/Addiction</p>	<p>In a “Shifting the Burden,” a problem is “solved” by applying a symptomatic solution (B1), which diverts attention away from more fundamental solutions (R3). (See <i>The Systems Thinker</i>, September 1990.) In an “Addiction” structure, a “Shifting the Burden” degrades into an addictive pattern in which the side-effect gets so entrenched that it overwhelms the original problem symptom. (See <i>The Systems Thinker</i>, April 1992.)</p>	<ul style="list-style-type: none"> • Problem symptoms are usually easier to recognize than the other elements of the structure. • If the side-effect has become the problem, you may be dealing with an “Addiction” structure. • Whether a solution is “symptomatic” or “fundamental” often depends on one’s perspective. Explore the problem from a differing perspective in order to come to a more comprehensive understanding of what the fundamental solution may be.
<p style="text-align: center;">Success to the Successful</p>	<p>In a “Success to the Successful” archetype, if one person or group (A) is given more resources, it has a higher likelihood of succeeding than B (assuming they are equally capable). The initial success justifies devoting more resources to A, and B’s success diminishes, further justifying more resource allocations to A (R2). (See <i>The Systems Thinker</i>, March 1992.)</p>	<ul style="list-style-type: none"> • Look for reasons why the system was set up to create just one “winner.” • Chop off one half of the archetype by focusing efforts and resources on one group, rather than creating a “winner-take-all” competition. • Find ways to make teams collaborators rather than competitors. • Identify goals or objectives that define success at a level higher than the individual players A and B.
<p style="text-align: center;">Tragedy of the Commons</p>	<p>In a “Tragedy of the Commons” structure, each person pursues actions which are individually beneficial (R1 and R2). If the amount of activity grows too large for the system to support, however, the “commons” becomes experiences diminishing benefits (B5 and B6). (See <i>The Systems Thinker</i>, August 1991.)</p>	<ul style="list-style-type: none"> • Effective solutions for “Tragedy of the Commons” scenario never lie at the individual level. • Ask questions such as: “What are the incentives for individuals to persist in their actions?” “Can the long-term collective loss be made more real and immediate to the individual actors?” • Find ways to reconcile short-term cumulative consequences. A governing body that is chartered with the sustainability of the resources limit can help.

PART III

STRUCTURAL THINKING TOOLS

GRAPHICAL FUNCTION DIAGRAM

STRUCTURE-BEHAVIOR PAIR

Structural thinking tools can help us become even more explicit about the structures that create the dynamic behaviors we are trying to understand. “From Causal Loops to Graphical Functions: Articulating Chaos” (p. 24) and “Graphical Functions: ‘Seeing’ the Full Story” (p. 26) describe how graphical function diagrams can easily represent nonlinear relationships. These relationships characterize the nature of most interconnections in complex systems (as opposed to the simple, linear relationships that are often assumed).

The second half of this section focuses on structure-behavior pairs. Accumulators and flows provide a rigorous framework for representing systemic structures in a more precise way than through causal loop diagrams alone. They can better represent a system’s nonlinearity, as well as distinguish between things that accumulate (water in a bathtub) and things that flow (water flowing through a faucet). The articles on accumulators and flows (p. 28–37) show how these concepts add further precision to our thinking and understanding about the link between structure and behavior.



FROM CAUSAL LOOPS TO GRAPHICAL FUNCTIONS: ARTICULATING CHAOS

In *Chaos: Making a New Science* (Penguin Books, New York), James Gleick describes a relatively new branch of science that has profound implications for how we view our world. Chaos, simply put, is the science of seeing order and pattern where formerly only the random, erratic, and unpredictable had been observed. In a way, systems thinking also deals in the science of chaos. Diagrams such as causal loops, accumulators and flows, and graphical functions are all ways of extracting the underlying structure from the “noise” of everyday life.

RELATING BEHAVIOR TO STRUCTURE

Both systems thinking and chaos insist that real-world phenomena need to be described in terms that reflect our intuition. Writing partial differential equations to describe clouds, for example, misses the point, because we don’t perceive clouds in

that way. It is not enough to have a model that reproduces some real-world phenomena if we cannot identify the structures that produce the behavior of the actual system. That is why systems thinking diagrams focus on capturing reality in a format that taps into our intuitive understanding of the systems in which we manage and live.

FROM CAUSAL LOOPS TO GRAPHICAL FUNCTION DIAGRAMS


To see how a range of systems thinking tools can help capture the structure of a system at increasing levels of detail, let’s look at a system we are all familiar with—a savings account. If we plot out the structure of a savings account using a causal loop diagram (see “Savings Loop”), we see that an increase in savings will lead to more interest earned, which increases our savings balance still further. The graph of the behavior over time

would look something like the exponential growth curve shown on the right of the diagram.

“Wait a minute,” you may protest, “I don’t know whose bank account that is, but it certainly doesn’t look like mine!” That’s true—rarely is a system so simple in real life; nor are bank accounts that well-behaved. There are usually many other factors


involved. The question of how many factors to include always depends on the purpose of examining the system. Since the details of any system are infinitely complex, it is futile to strive to “model the system.” In our sample case, the purpose is to represent as concisely as possible the important factors that affect the balance of a typical savings account, so we want to look at savings, income, interest earned, and spending (see “Savings and Spending Loops”). If we were only interested in capturing the fact that there is a balancing loop that explains the slowdown in the growth of our savings account, we could stop at this point. On the other hand, if we want to be more explicit about the structure behind the behavior, we need to translate our diagram into accumulators and flows.

SAVINGS LOOP



If we begin to explore our savings account “system” by drawing a causal loop diagram, we see that an increase in savings will lead to more interest earned, which increases our savings balance still further (left). The graph of this behavior over time will look something like the exponential growth curve (right).

SAVINGS AND SPENDING LOOPS



A balancing loop explains the slowing growth of the savings account: As our savings increases, we are more likely to increase spending, which will reduce our savings.

ACCUMULATORS AND FLOWS

When we translate CLDs into accumulators and flows, we are becoming even more precise about the structures producing the dynamics. The bathtub as a metaphor for accumulations helps us visualize how concepts as diverse as savings, pollution, customers, and corporate reputation share a similar underlying structure (see “Accumulators: Bathtubs, Bathtubs Everywhere,” p. 30).

Accumulators and flows add more detail and understanding to our causal loop diagram by differentiating between those variables in the diagram that “accumulate” (our savings balance) and those that just “flow” through the system (income and spending). In the “Savings as an Accumulator” diagram, we can visually see money flowing into and out of savings in the form of income and spending. More importantly, we can relate to this structure intuitively because we experience money in terms of flows and accumulators (or lack thereof!).

GRAPHICAL FUNCTIONS: MAPPING POLICIES

So now we have a pretty good idea of both the basic dynamic behavior of the savings account, and a feel for the important inflows and outflows. But our model is still pretty elementary. Suppose now you wanted to go a little further and use a systems diagram for describing your family’s policy for managing your savings. “Our discretionary spending depends on how much savings we have,” you explain. “If the balance in our savings account is below \$5000, we don’t spend a dime. As our savings rise above \$5000, we may increase discretionary spending by, say, \$15-20 per month. If our savings tops \$10,000, then we’re likely to spend several hundred dollars a month. But in any case, we don’t see ourselves spending more than \$500 per

month on discretionary expenses.”


Graphical functions allow us to expand our exploration of a system to include policies and interrelationships between variables. If we tried to capture the savings plan we described above in an analytical form, we would have to do quite a bit of work in order to come up with a suitable equation. And when we were done, it would be hard to tell if the equation represented our savings account or the number of widgets on sale at Wal-Mart. The truth is, most of us don’t think in abstract mathematical concepts, but in images and structures grounded in our everyday experience. That is why graphical functions are useful. They capture policies in an intuitive way through a simple graph that maps out the relationship between one variable in relation to another (see “Savings Policy Graphical Function Diagram”). In our savings policy plan, for example, we see at a glance that savings has no impact on our discretionary expenses until savings hits \$5000. After that, discretionary expenses rise until savings reaches \$20,000, at which point they level out at \$500.

ARTISTIC MANAGERS

Physicist Mitchell Feigenbaum suggests that art is a theory about the way the world looks to human beings. “It’s abundantly obvious that one doesn’t know the world around us in detail. What artists have accomplished is realizing that there’s only a small amount of stuff that’s important, and seeing what it is.”

Whether we recognize it or not, we are artists as well, selectively picking out details of the world that we choose to focus on. Those details appear


SAVINGS POLICY GRAPHICAL FUNCTION DIAGRAM



This Graphical Function Diagram captures our savings policy in an intuitive way by mapping out the relationship between savings and discretionary expenses.

as items on our production reports, financial statements, and customer surveys. To the extent that those details do not capture the core structures that are important, we may be the unwitting producers of our own chaos. As one systems thinking maxim warns, “It ain’t what you don’t know that hurts you, it’s what you DO know that ain’t so.” ■

SAVINGS AS AN ACCUMULATOR



Accumulators and flows add more detail and understanding to our causal loop diagram by differentiating between those variables in the diagram that “accumulate” (savings) and those that “flow” through the system (income and spending).



GRAPHICAL FUNCTIONS: “SEEING” THE FULL STORY

An executive of a large automotive company tells the story of two engineers who were arguing about the correct angle of an engine mount. The two had been at it for more than half an hour—one engineer swearing that the angle was 40 degrees while the other fumed that it was 50 degrees. After several civil attempts to correct each other’s viewpoint, they had just started attacking each other’s intelligence, ability, and character when the executive happened to walk by.

“What axis of reference are you using?” he asked.

“The vertical, of course!” exclaimed one engineer.

“The horizontal!” said the other.

Both stopped in amazement as they realized *they had been saying the same thing!* Because they had not established a common

frame of reference for their discussion, each had assumed the other’s viewpoint was wrong.

INDIVIDUAL WORLDS

The story of the two engineers points out an age-old communication problem. Each of us carries our own set of assumptions about reality—our own individual picture of the world. Oftentimes, we mistakenly assume that our viewpoint is the only way of looking at a situation. Both engineers, for example, believed the other person’s position was based on the same axis as their own—they never even questioned it. If we don’t acknowledge our assumptions at the outset of a discussion, we risk experiencing the same frustrations as the two engineers.

In many instances, spoken language can be a hindrance rather than a help in communicating our mental pictures of reality because words, unlike pictures, do not force us to be explicit when explaining our reasoning. Graphics, because they can represent ideas more clearly, can be a much more powerful and effective means of communication (see “Systems Thinking as a Language,” p. 6). Trite as it may sound, the saying “a picture is worth a thousand words” still holds true. Had the two engineers simply drawn two axes and a line, they would have saved a great many angry words. When the issue is more complex than a single angle, the use of graphics can become even more important for reaching a shared understanding.

Using graphical function diagrams (GFDs), it is much easier to capture how two variables relate in a format that is concise and invites others to share their own

perspectives. Graphical functions can help us go beyond merely observing correlational relationships (when X happens, Y happens) to exploring our understanding of the causal connection between two variables (X causes Y). In constructing GFDs, one should follow the 60 percent rule—it’s better to get it 60 percent right very quickly and spend time modifying it than spend a great deal of effort trying to get it 100 percent right the first time.

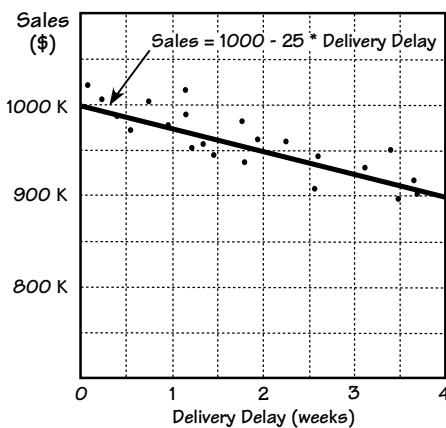
GRAPHICAL FUNCTIONS VS. SCATTER CHARTS

Graphical functions are best described by first establishing what they are *not*. Although they may look similar, graphical functions are not the same as scatter charts, which plot one variable’s data against another’s. If we were to look at the relationship between sales and delivery delay using a scatter chart, we would plot some data points and then draw a regression line through them (see “Sales vs. Delivery Delay”).

From the scatter chart, we can see that in weeks one through four, sales fall by \$25K for each one-week delay. We can then extrapolate beyond the historical data to predict that a five-week delay will result in an additional \$25K drop in sales. In general, scatter diagrams answer the question, “What happened historically, was there a correlation, and based on that information what can I expect to happen in the future?” They tend to be retrospective.

A graphical function, on the other hand, is very much prospective in nature. By including the full spectrum of possible values, GFDs can help you see beyond the

SALES VS. DELIVERY DELAY



Scatter charts plot one variable’s data against another and answer the question, “What happened historically?” The view is retrospective.



historical range of operating values and ask, “Given my understanding of the system, what do I think will happen at each possible point?”

CREATING A GFD

A graphical function diagram can help explicate your (or a team’s) mental model of the relationship between two critical variables. Unlike behavior charts, GFDs do not show how variables change over time, but how two variables interrelate. To create a GFD, it is best to begin by answering the following questions:

- What do we know from the outset about the causal relationship between these two variables?
- Are there any “neutral zones” where the variable on the y-axis is not affected by changes in the x-variable?
- What are the extreme values that both variables can assume?

If we looked at the sales and delivery delay example using a GFD, we would start by asking what we think the general nature of the relationship is between the two variables—is it flat, is it upward sloping, or is it downward sloping? With most products, longer delivery delays mean lower sales, so we can assume that the relationship would slope downward.

Using available historical data and past experience, we can then take a first cut at identifying a neutral zone where sales may be insensitive to differences in the length of the delay (see “Delivery Delay Graphical Function Diagram”). Past experience may suggest that sales will increase steadily as the delay falls below four weeks. A sampling of customer contacts may tell us that there is not a whole lot of difference between 3.5 and 4.5 weeks. On the other hand, past market research also tells us that if the delay grows greater than five weeks, sales will fall dramatically. Looking further, we realize that even in the extreme case of a 20-week delay, \$200K of sales will

still come from captive customers who have nowhere else to turn in the short term.

BUILDING SHARED UNDERSTANDING

The resulting diagram is a concise causal hypothesis that states that customers will reward shorter delays with slightly higher orders, but will severely penalize delays that extend beyond an acceptable range.

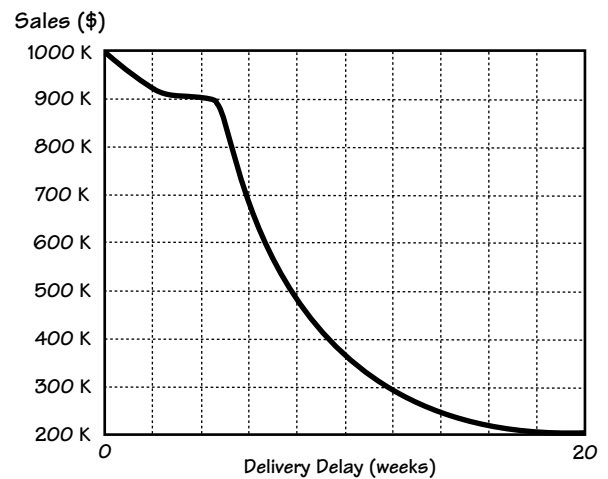
The GFD conveys a much richer description about the relationship between delivery delay and sales than a scatter chart based on historical data. The diagram helps visualize the full range of implications and minimizes the danger of remaining myopically focused on a narrow band of possible outcomes. Developing the diagram as a group can also help surface differing mental assumptions about the potential impact of deteriorating or improving delivery performance (remember the engineers!).

Sometimes it is helpful to convert the relationship into a more general form where the y-variable is converted to an “effect-of” variable. Instead of “Sales” on the y-axis, for example, we would have “Effect of Delivery Delay on Sales” (see graph), which shows that a 3.5 to 4.5 week delay has no effect, shortening the delay nets us a maximum gain of 5 percent (1.05 times the sales number we would have obtained if we were in the neutral zone), and lengthening the delay to 20 weeks can choke sales by

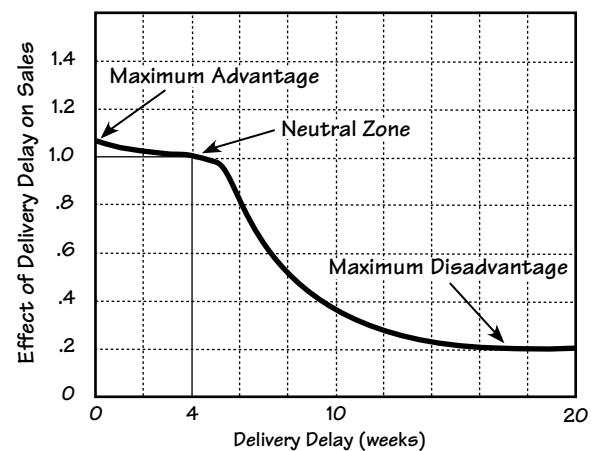
as much as 80 percent.

The “Effect-of” version of the GFD focuses attention on the relative impact of the delivery delay on sales instead of on the specific numbers themselves. In this way, we can compare across different variables, such as the relative effects of quality on sales vs. marketing spending on sales, and make explicit our understanding of which factor is the dominant driver. ■

DELIVERY DELAY GRAPHICAL FUNCTION DIAGRAM



EFFECT OF DELIVERY DELAY ON SALES



A graphical function diagram conveys a much richer description about the relationship between delivery delay and sales (top) than a scatter chart based on historical data. Looking at the effect of delivery delay on sales (bottom) allows you to compare the relative effects of different variables, to see which is the dominant driver.



STRUCTURAL THINKING: THE WORLD ACCORDING TO ACCUMULATORS AND FLOWS

A vice president of a major U.S. manufacturer once questioned whether today’s rapid pace of change means that all our old tools and ways of managing are now inadequate. “Are we doomed to keep on throwing out our current tools and practices as soon as the next wave of innovations comes along?” he asked.

The answer is . . . “it depends.” It depends on the underlying theory on which the current tools and methods are based. If one’s management practices are based on transient or situation-specific phenomena, they are likely to require revision whenever the circumstances change. If, on the other hand, they are based on a structural understanding, the situation may change, but the tools will still apply.

WHERE ARE THE COWS?

Barry Richmond of High Performance

Systems tells this story: “While perusing a well-known economic journal, I came across an article which described a model that had been constructed to forecast U.S. milk production. The model was of the $Y=f(X_i)$ form [$Y = Y_0 + a_1X_1 + a_2X_2 + \dots + a_nX_n$], where the X_i ’s included such things as: last year’s milk production, interest rates, spending on cattle feed, GNP growth, and other macroeconomic factors. As the article detailed, the model performed quite well as a predictive device—at least in terms of its ability to ‘track history.’ The obvious thing about this model, that would bother both dairy farmers and people who were partial to operational specifications, is: ‘where’s the cows?!’ Simply stated, if you’ve got no cows, you’ve got no milk! Crude, but true.”

How does all this talk about cows relate to our vice president’s question? Well, imagine that an epidemic swept over the

country and killed all the cows. What would the above model predict for next year’s milk production? The answer would most likely look a lot like the number for last year’s milk production, which is clearly incorrect. The model must be abandoned.

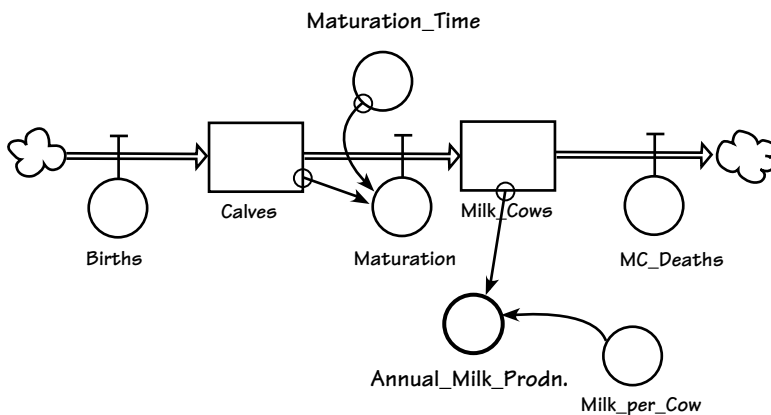
“Unfair,” you might say. “It’s not that the model is wrong. It’s just that the world has changed dramatically since the model was originally built and the changes must now be added.” But what has really changed? Yes, the cows are now dead, but the basic fact that milk comes from cows, and that without cows there can be no milk, is as true now as it was before the mass decimation. From a structural perspective, the nature of the world has not changed at all. The model was inadequate because it was based on situation-specific data that has now changed.

STRUCTURAL THINKING

When we look at the world through a structural lens, we are interested in understanding how things actually work. We are less interested in correlational relationships and more interested in the causal structures that produce the observed behavior. This is not to say that nonstructural models aren’t valuable. Regression models, for example, have many applications and are useful for identifying correlation, explaining sources of variance, and extrapolating from historical data. Those models are inadequate, however, for gaining insight into how a system actually operates.

If we were to look at the milk production model from a structural viewpoint, we would start with the basic fact that milk

MILK PRODUCTION MODEL



If we wanted to create a structural representation of milk production, we would begin with the central accumulator “milk cows.” Milk production is determined by the number of cows and the amount of milk per cow. To create our hypothetical scenario of an epidemic, we would simply enter zero for the number of milk cows. The resulting annual milk production would also be zero.



comes from cows. Therefore, cows are the central accumulator in the model—the number of cows accumulates over time, as cows are born, mature, and become milk cows (see “Milk Production Model”).

Depending on the scope of our study, we may be interested in representing the lifecycle of all cows, or just milk cows. In this case, we will focus our attention on the flow of cows from birth through maturity into the milk cow accumulator. The annual milk production is then determined by the number of milk cows at any one time and the amount of milk per cow. Of course, there are many other factors that affect milk production, such as food supplies, milk demand, and dairy farmers. These factors could also be added to our diagram in the form of additional accumulators and flows.

The resulting model can then be simulated on a computer to see how annual milk production behaves over time. To create our hypothetical epidemic scenario, for example, we would simply put zero for the stock of cows. In that event, the annual milk production would also equal zero. Because this model is tied to the structure of the system, not just historical data, it would not have to be thrown out even if all of the cows suddenly died.

LEVELS OF EXPLANATION

We live in the world of events. As a result, we encounter and navigate through the rapids of life on an event-by-event basis. But this does not mean that we must act on each event as if it were an isolated occurrence. We can look at patterns of behavior over time and try to glean lessons from the past that will improve our ability to handle present situations. That is the purpose of forecasting models.

Forecasting models, like the economist’s milk production model described above, attempt to provide information about the future by looking at the past. But in many

ways, managing on the basis of forecasts is a lot like trying to drive a car by looking through the rearview mirror. When does it work best? When the road is straight and there are no obstacles in the way. When does it fail? The rest of the time! When using a forecasting model, you only realize you have missed a turn once you see the cliff’s edge behind you and feel the sensation of free fall hit your stomach.

Forecasting provides very little insight into what actually produces the observed behavior. Consequently, it allows us to anticipate and react to changes only if they do not deviate too much from past behavior. Models, on the other hand, capture the structural forces at work and are therefore less situation-dependent. To come back to the vice president’s question, structural thinking provides a more stable basis of understanding that will last even through times of turbulent change.

GENERIC THINKING SKILLS

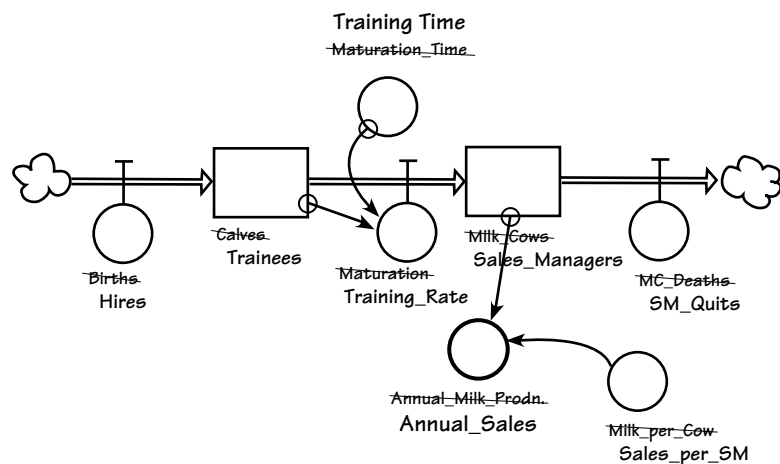
If we begin to view the world through a structural perspective, another benefit

emerges—the ability to transfer insight. This ability to see similar structures occurring in diverse settings is referred to as “generic thinking,” and the structures themselves are referred to as “generic structures.”

For example, if we take the “Milk Production Model” and substitute “hires” for “births,” “trainees” for “calves,” and “sales managers” for “milk cows,” we can transform the milk cow model into a model that can be used to explore the structural forces that influence annual sales (see “Sales Growth Model”). The same generic resource development structure underlies both models. Although we may debate whether it takes longer to produce a milk cow or a sales manager, we can both agree that the structure of both processes is fundamentally the same. ■

For further reading about structural thinking and the other critical thinking skills included under the systems thinking umbrella, see Barry Richmond’s The Thinking in Systems Thinking: Seven Essential Skills (Pegasus Communications, 2000).

SALES GROWTH MODEL



If we replace the names of the variables in the “Milk Production Model” with those listed above, we can create a model that explores sales growth. The same generic resource development structure can be used to describe both processes.



ACCUMULATORS: BATHTUBS, BATHTUBS EVERYWHERE

When's the last time you actually took a real, honest-to-goodness bath? If you are like most people, it has probably been quite a while. We live in the world of quick showers and instant breakfasts. Yet, it wasn't too long ago when taking baths was part of our normal daily routine. The shift from baths to showers marked a far more deeper change in our thinking than merely a change in personal hygiene habits.

When we run the bathwater, we can visually see the water accumulating in the tub (see "Bathtubs and Accumulators"). We know we have to keep an eye on the water level so it won't overflow. When we take showers, however, the accumulation process is virtually eliminated. Water flows out of the showerhead, over our bodies, and out the drain. Where does the water go? We hardly give it any thought.

SHOWERHEAD VS. BATHTUB THINKING

Taking showers disconnects us from experiencing one of nature's most basic structures—accumulators. Lakes and ponds are accumulators of various water flows. Global warming has been attributed to the cumulative effects of burning fossil fuels. Plants are accumulators of energy and nutrition. Displacement, velocity, and acceleration can be represented in terms of accumulators. That is, displacement represents the accumulation of past velocity, and velocity is an accumulation of past acceleration.

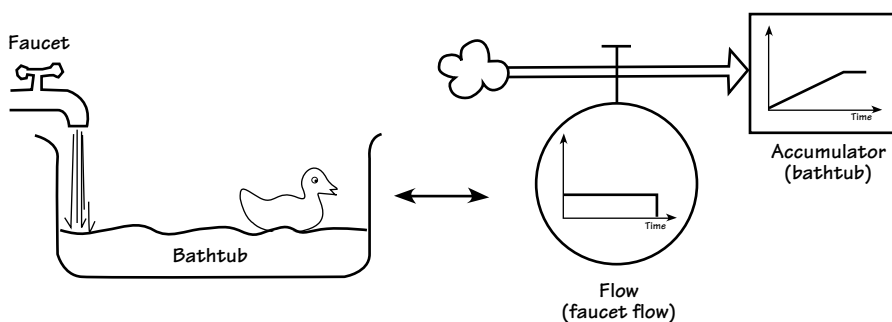
If we use showerhead thinking, we are less conscious of accumulations. Flows of materials such as water, fuel, or energy simply "go away" somewhere. But from a bathtub—or systems—perspective, there is no "away." Everything accumulates somewhere. Forgetting about that "somewhere"

can lead to disastrous results.

When Just-in-Time (JIT) manufacturing first hit the U.S., for example, many companies implemented it using a showerhead perspective. The basic concept of JIT is to manage a steady flow of materials through a factory with minimal accumulations of inventory at each step. Many companies that instituted JIT tried to minimize their own accumulations by demanding that their suppliers provide them with materials just when they needed them and not any sooner.

The problem with the above approach, of course, is that the flow of materials has to accumulate somewhere, and it was accumulating in the suppliers' warehouses. The JIT flow was accomplished by shifting the accumulations to suppliers, severely straining the relationship between suppliers and manufacturers. Bathtub thinking would have highlighted the fact that unless the entire flow from raw materials to final customer worked together, there would be undesirable accumulations for somebody in the system.

BATHTUBS AND ACCUMULATORS



The core building blocks of dynamic thinking tools are reinforcing and balancing loops. The analogous elements in the structural thinking set of tools are accumulators and flows. An accumulator (or stock) is represented by a rectangular box, and the flow (or rate) is represented by a pipe with a directional arrow, a valve, and a circle. The circle and the box each contain a timeline graph as a visual reminder that the dynamics of the two are intimately connected. For example, the constant flow from the circle to the box as indicated in the diagram must produce the straight linear rise in the water level. No other behavior is possible for that structure as it is drawn.

INVISIBLE BATHTUBS

When's the last time you actually let a bathtub overflow? Probably not in a long time. Of course, we all know not to let the water run indefinitely, because the tub has a limited capacity. The tub's dimensions are obvious and so is the rising water line. But suppose the bathtub is invisible, and so is the water once it leaves the faucet. And suppose you are not in the bathroom to keep an eye on the tub—you are off answering phone calls and dealing with the latest crisis at the office. How will you

know when the bathtub is getting full or already overflowing?

Flows are easy to keep track of because they involve action, and actions are easy to measure—how many products to ship, how many people to hire, for example. Some accumulations are also very visible, such as order backlogs or bulging inventories. There are, however, many accumulations that are not tangible but nonetheless very real. These possess the same behavioral characteristics as physical accumulators and flows, but they are like invisible bathtubs—we can never tell for sure whether they are overflowing or not.

IDENTIFYING ACCUMULATIONS

So how can you locate the “invisible bathtubs” lurking in your company? For every flow (action, decision, policy), try to figure out what, if anything, is accumulating and what are the implications of those accumulations.

For example, as workload outstrips capacity and work pressures become excessively high (see “Stress Accumulator”), you should question whether those pressures simply come and go or whether their effects are accumulating somehow. For example, extra pressure may generate more stressful events, which will accumulate into increasing levels of stress. High stress levels will then lead to lower productivity, which further reduces work capacity and leads to more stressful events. This reinforcing loop of accumulating stress is intangible, yet all too real for many people.

If you look at the situation from the accumulator viewpoint and trace out the reinforcing loop, it becomes clear why typical stress reduction efforts do not work very well. Each round of stressful events produces more stress, like compound interest in a savings account. And coping mechanisms are like savings with-

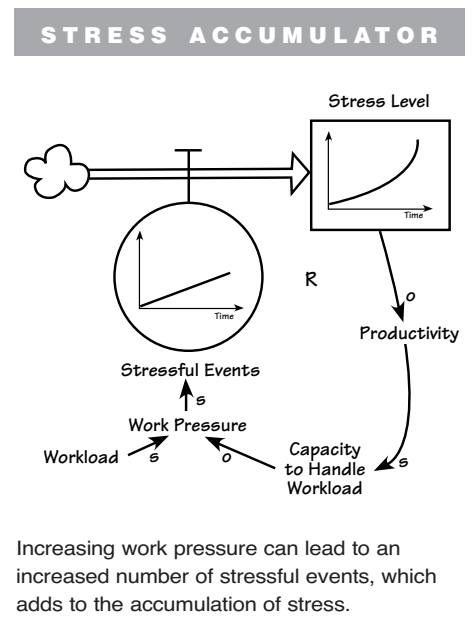
drawals—unless you withdraw more than you are earning in interest, the account balance never goes down. Likewise, if the stress “withdrawal” rate (coping mechanisms) are not exceeding the stress “interest” rate (stressful events), then the best you can do is learn to live with the higher stress level. From the accumulator perspective, the high-leverage action would be to “close the account” by reducing or eliminating the real source of stress.

LOOP DIAGRAMS VS. ACCUMULATORS AND FLOWS

If causal loop diagrams and systems archetypes are such powerful tools, why do we need to bother with accumulators and flows? Both tools have their unique strengths. Tools like systems archetypes capture and communicate dynamic issues in a concise way, but they do not provide a detailed representation of the structure producing the dynamics.

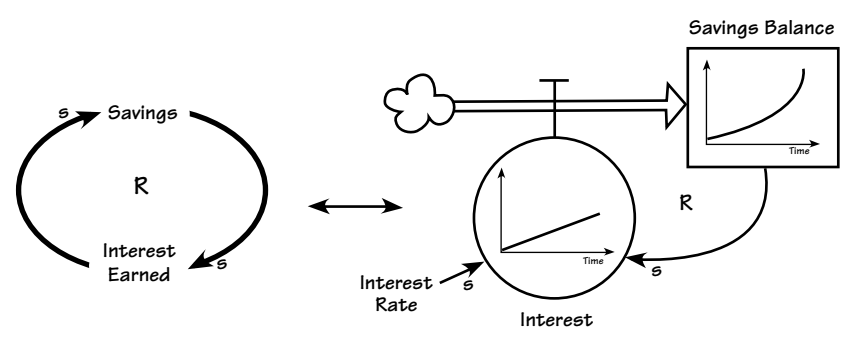
There are cases when tracing through a loop diagram can be confusing. For example: “Savings and interest form a reinforcing loop where higher savings balance leads to higher interest payments,

which leads to still higher savings (see “From Loop Diagrams to Accumulators and Flows”). If we start making withdrawals, the balance goes down and interest payments decrease, but savings does not decrease. It still increases but at a decreased rate.” Sound confusing? That’s where the accumulator and flow diagram can help you actually visualize how that loop works in terms of the flow of money into and out of the account. ■



Increasing work pressure can lead to an increased number of stressful events, which adds to the accumulation of stress.

FROM LOOP DIAGRAMS TO ACCUMULATORS AND FLOWS



The reinforcing loop of savings and interest can be represented as a causal loop diagram (left) or as an accumulator and flow diagram (right), where you can visualize the flow of interest into savings.



ACCUMULATION MANAGEMENT: AVOIDING THE “PACK RAT” SYNDROME

There is a story about a trivia “pack rat,” a man who had spent his entire life memorizing trivia. He knew baseball statistics of every player in the history of the major league. He had memorized the titles, directors, and actors of hundreds of movies. He knew the name of every television show that had ever aired.

But one day he found himself in an awkward predicament—no matter how hard he tried, he could not memorize any more trivia. He had finally taxed the limits of his rote memorization capacity. Although he had worked hard at acquiring his stock of trivia throughout his life, he had never considered how he might go about depleting it. He had not learned the fundamentals of

accumulator management.

PACK RATS AND NOMADS

Life can in some ways be viewed as a never-ending task of managing various accumulators. Our pantries, refrigerators, checking accounts, and closets are among the many accumulations we manage daily.

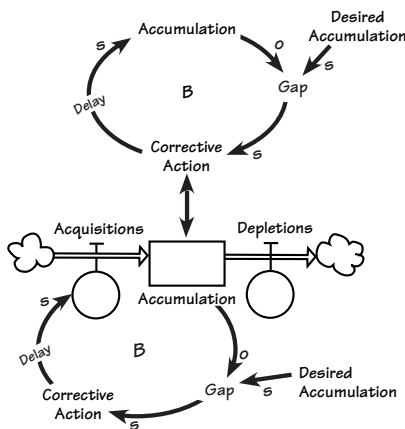
On one end of the accumulation management spectrum is the pack rat who throws nothing away. On the other end is the “nomad” who makes a virtue of owning no more than what can be packed into one suitcase. In between these two extremes lies the majority of the population who is constantly struggling to maintain the right balance between acquisitions and depletions.

ANATOMY OF ACCUMULATOR MANAGEMENT

A typical accumulator management structure (AMS) has the following elements: the Accumulation, Acquisitions, Depletions, Desired Accumulation, and a Corrective Action (see “Accumulator Management”). In addition, there is almost always some delay between the Corrective Action and the Acquisition, because it takes time to actually memorize data or clear out the closet once we have decided to do so.

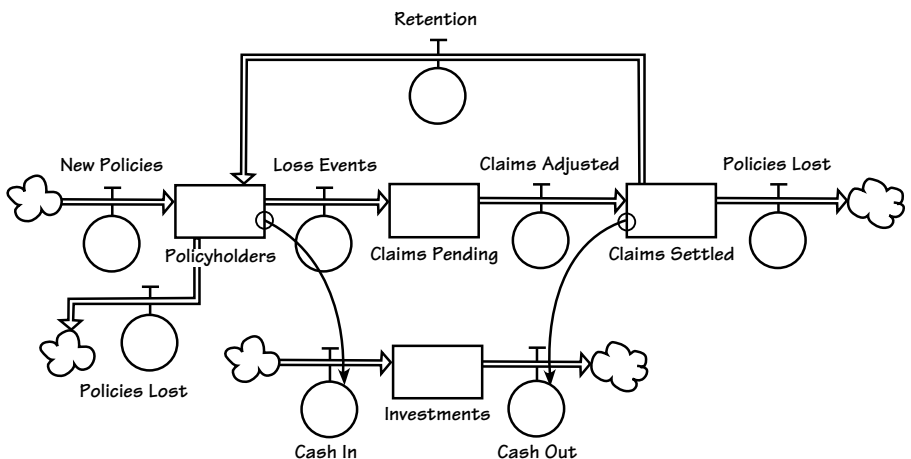
The accumulator management structure is a generic structure that can represent a wide range of business settings where accumulation management is

ACCUMULATOR MANAGEMENT



Accumulator management can be viewed simply as a balancing loop with delay (top). A structural diagram (bottom) reveals that the flows controlling the accumulation are acquisitions and depletions.

INSURANCE BUSINESS AS ACCUMULATION MANAGEMENT



The insurance business can be mapped into a relatively simple diagram that highlights the major accumulators and flows. If we assign numbers next to each accumulator or flow indicating the percentage of organizational resources devoted to it, the diagram can help reevaluate the organization’s current emphasis.

important. For example, the insurance business can be mapped into a relatively simple diagram by focusing on the basic accumulators and flows (see “Insurance Business as Accumulation Management”). Insurance revolves around managing two main accumulators—policyholders and investments.

If managers assign a number next to each accumulator and flow in the diagram to represent the percentage of organizational resources devoted to each, the diagram can highlight which areas receive the largest organizational focus. This exercise can point out any weaknesses in the current organizational emphasis—for example, spending too little time trying to retain current policyholders—and reveal ways in which the company can better serve its customers.

SUPPLY LINES AND DELAYS

If we had direct and immediate control over all the elements in the AMS diagram, managing accumulations would be simple: We would calculate the depletion rate, set our desired accumulations accordingly, and implement actions that will immediately result in acquisitions. In our home life we already pretty much follow this pattern. For example, we plan our meals, decide on an appropriate amount of food to have on hand, figure out how long it will be before we run out of certain staples, and go to the grocery store as needed. Unfortunately, things are not that straightforward when we move into the organizational context.

One of the most challenging aspects of managing accumulations within organizations is captured in one word—delays. Identifying and characterizing the nature and source of delays often plays a critical role in managing accumulations effectively. A big part of the problem is that we usually have very little control over the supply line delay.

MANAGING THE “BEER GAME”

In a production distribution system game fondly known as the Beer Game, participants are given the task of managing their own inventory (accumulation) of beer. Each team is composed of four players linked together in a structure similar to that represented in the AMS diagram (see “Supply Line and Delay in the Beer Game”). Within that team, each participant must make ordering decisions in order to maintain his desired level of inventory.

According to MIT professor John Sterman, when participants try to manage accumulations in the Beer Game they usually run into three common problems. First, they typically underestimate the true length of the delay from the time they order to when they receive the beer and then overadjust their orders—even when they are given full information about the supply line delays. They do not appear to recognize that their ordering decisions affect the length of the supply line delay—that is, the more they order, the longer it takes to receive the beer.

In addition, he found that when people find it difficult to determine their optimal inventory level, they simply anchor their desired inventory on the initial inventory and adjust from there.

This finding highlights the more general tendency people have to anchor on past goals or standards rather than search for better ones.

The third observation is that people generally point to factors outside the system as being responsible for the instabilities they observe in the game. That is, people offer open loop expla-


nations rather than connecting the dynamics back to their own decision-making. In fact, the wide oscillations in inventory are actually generated by the decisions they make.

AVOIDING THE “PACK RAT” SYNDROME

If you want to avoid the “pack rat” syndrome, you need to manage the whole accumulator management structure and not just focus on one piece of it. The observations about the difficulties of managing the Beer Game suggests that you should think through the following questions when confronting a typical accumulator management situation: (1) Where are the supply line delays and how are they changing? (2) What factors are determining what Desired Accumulation should be? (3) How do current policies and decisions feed back into this system to produce the results we have observed? The accumulation management structure diagram is a useful starting point to begin addressing these questions. ■

Further Reading: “Modeling Managerial Behavior: Misperceptions of Feedback in a Dynamic Decision Making Experiment,” by John D. Sterman, Management Science, Vol. 35, No. 3, March 1989.

SUPPLY LINE AND DELAY IN THE BEER GAME



The structure of the inventory management system in the Beer Game is similar to the AMS diagram. Understanding the nature and source of delays in a systems—such as the supply line delay above—often plays a critical role in managing accumulations without overcorrecting.



DELAYS: ACCUMULATORS IN DISGUISE

Imagine a new manager at a beef packaging plant who knows nothing about the birthing process of calves. On his first day, his workers show him a newborn calf. The dollar signs go off in his head as he calculates: More calves mean more beef; more beef means more sales; more sales mean more profits. He points to the mother cow and barks, “I want you to get two more calves out of that sucker by Monday morning, and that’s an order!”

Of course, the workers will find a way to fulfill his request, either by bringing two calves from another part of the plant or perhaps slaughtering the mother cow to produce the extra pounds of beef. The workers will have successfully executed their task, and the manager will continue to believe that his orders control the production cycle of calves.

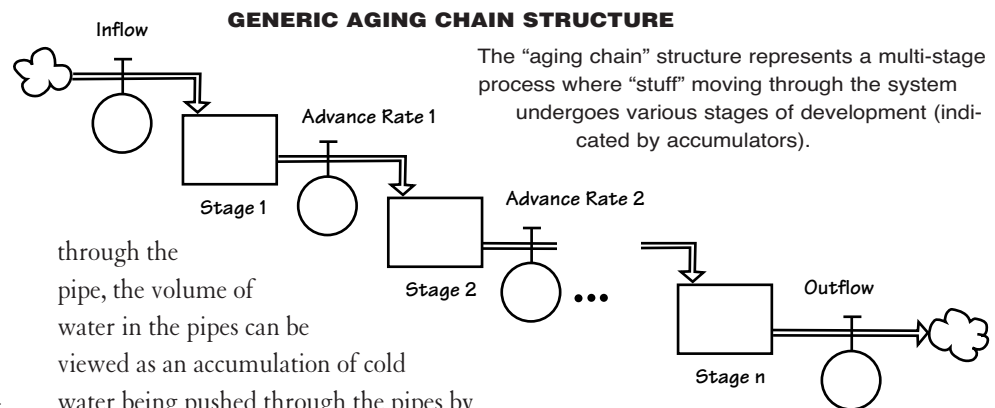
The story is obviously far-fetched. No one would expect a cow to produce a calf in one weekend. But how do we know that equally ridiculous demands are not being made every day on processes where we have less understanding of the time dynamics? Are there such “calving” equivalents in manufacturing, for example, where arbitrary quarterly sales targets given to investment analysts translate into marching orders for the production line?

AGING CHAIN STRUCTURE

Why is it that whenever we want something right away, it seems to take forever? Then when we do get it, it often is more than we ever wanted? Chances are, delays played a large part in our mis-timing. We realize that “things take time,” but unless we have a clear understanding of the struc-

ture that produces the delays, it is difficult to know how long to wait before we should take further action.

Whenever there are significant delays in a system, you can bet that there are accumulators involved. In some cases the accumulators are less obvious than in others, but they almost always play a role. In the case of the packaging plant, the accumulator is the cow (or stock of cows). When you take a shower, the delay in getting hot water is due to accumulation in the length of pipe from the hot water heater to the shower. Even though the water “flows”



through the pipe, the volume of water in the pipes can be viewed as an accumulation of cold water being pushed through the pipes by the subsequent hot water.

The “aging chain” structure is the name given to a whole class of processes that includes the birthing cycle of cows, production-distribution systems, and the spread of infectious diseases (see “Generic Aging Chain Structure”). The key principle in the aging chain is that there are delays in the system that depend on the inherent nature of that system, and those delays cannot be shortened except within some narrow bounds. For calf production that delay is the gestation period. For new ideas it may be an incubation period, for new products it may be the development

time, and for production it is the manufacturing cycle time. Trying to shorten the inherent time delay by pushing things through the accumulators faster can wreak havoc on the system.

The aging chain represents a multi-stage process where “stuff” (ideas, products, calves) moving through the system undergoes various stages of development. Each stage can be represented by an accumulator, where the stuff “ages” before moving on to the next stage. The “aging” time at any stage can vary.

PRODUCTION SYSTEMS

One way to get a better idea of the delays involved in a system is to create a structural map. A typical production system is shown in the “Production Chain” diagram. In the diagram, different stages of production process are represented by accumulators. (For dimensional consistency, this diagram represents accounting numbers and not the actual physical stuff moving through the system.) The accumulator and flow diagram is very much like a process flow chart, showing how production can be mapped out as a series of stages. Each

accumulator, in effect, adds a delay to the system.

If we wanted to add greater detail, each accumulator could be further broken out into smaller stages. For example, the Work-in-Process accumulator may be broken into various production stages such as assembly, paint, bake, test, and inspection. Each stage has its own time delay or cycle time associated with it.

These time delays have an important aspect: They usually do not stay constant—a point that is not captured in a typical process flow chart. In the first stage, for example, a rising backlog of production requests increases the scheduling load, which decreases the rate at which the requests can be scheduled. This keeps the backlog high, further exacerbating the schedule load problem (R1). Once a schedule gets behind, the system can actually reinforce the tendency to fall further behind.

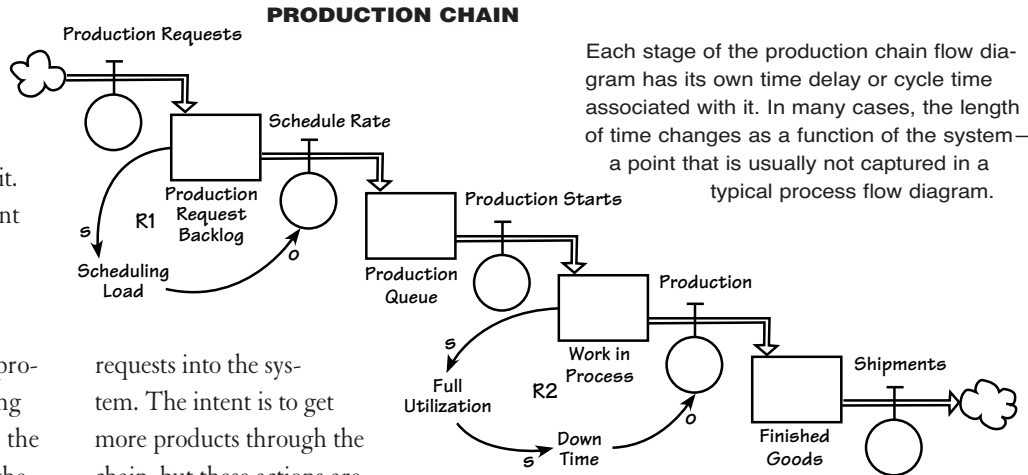
Similarly, at the work-in-process (WIP) stage, if we load up on production starts, the WIP inventory rises. As the WIP rises, the sheer amount of extra “stuff” in the works can slow things down. Another consequence of pumping up the system with additional WIP is that it creates pressure to run equipment at full capacity, which can lead to increased down time (no time for maintenance work) and lower yields (higher scrap rates). All this will ultimately lead to lower production and an accumulation of more WIP (R2).

DON'T JUST DO SOMETHING, STAND THERE!

The “Production Chain” diagram does not show the complete picture of what goes on in a typical manufacturing setting (see “Expediting Loops”). When shipments fall short of customer orders, customers are

left waiting for their goods. One response to such a gap is to go into the production system and expedite some of the more “important” orders and/or push more production

starts and pushing orders into production will exacerbate the full utilization loop (R2), leading to further delays and more pressure to expedite (R4).



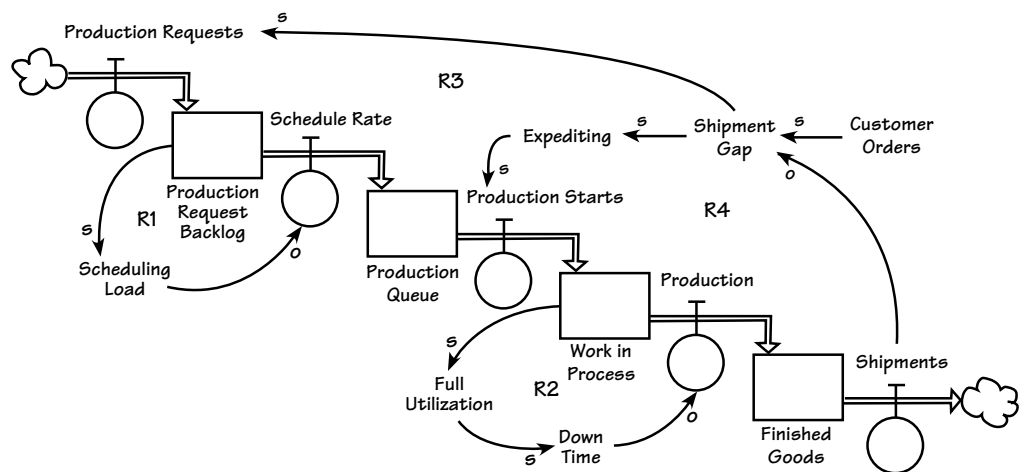
Each stage of the production chain flow diagram has its own time delay or cycle time associated with it. In many cases, the length of time changes as a function of the system—a point that is usually not captured in a typical process flow diagram.

requests into the system. The intent is to get more products through the chain, but these actions are likely to produce the opposite effect.

Given that the system was already running at full tilt, the expediting actions are likely to create additional reinforcing sources of delays. Pushing on production requests will kick in the scheduling load loop (R1), which will further delay shipments and intensify pressure to expedite (R3). Similarly, rearranging the production

In general, the aging chain structure tells us that there are structural limits to how fast you can force a system to respond. Unless you can somehow change the inherent delays built into the various stages of the system, the best expediting action one can take may be to simply do nothing—and wait. ■

EXPEDITING LOOPS



Efforts to expedite orders can backfire by actually lengthening the average production delay and reinforcing the need to expedite (loops R3 and R4).



S - SHAPED GROWTH AND THE LAW OF DIMINISHING RETURNS

Most of us are familiar with the story of Sir Isaac Newton sitting underneath an apple tree and “discovering” the law of gravity when he saw an apple fall (some say on his head). Had Newton been an entrepreneur, he might have discovered another law. With apple sales—and profits—in mind, he might have shaken the tree vigorously, causing more apples to drop to the ground. The harder he shook, the more apples would fall, meaning more sales and more profits. After a while, however, each shaking would produce fewer and fewer apples.

We can almost picture the scene: Sir Isaac wipes the sweat from his brow, then climbs the tree to knock down the tenacious few apples left. Precariously perched, he strains to reach one of the last remaining apples. The limb gives way, and he falls. As he lands on the ground, another

discovery hits him—the Law of Diminishing Returns.

THE LAW OF DIMINISHING RETURNS

The phenomena of diminishing returns—when more effort yields fewer results—is ubiquitous. Oil recovery and mining operations exhibit this behavior. Companies experience rapid new product sales followed by decreasing demand. At a personal level, we see that working longer hours, jogging more miles, and eating less food lead to diminishing returns in productivity gains, health benefits, and weight loss.

The Law of Diminishing Returns can be considered the law of gravity for the business world. Launching a marketing campaign, for example, is like the trajectory of a cannonball—the returns climb higher and higher, until the “force” of diminishing returns kicks in and pulls the rate of return down. When traced out over time, the cumulative returns of the marketing effort produce an S-shaped curve.

The Law of Diminishing Returns is essentially about saturation effects—reaching the limits of a particular system. The characteristic S-shaped behavior of this process can be produced by two different structures: the well-known Bass Diffusion Model and the more general “Limits to

Success” archetype. At the heart of both structures is a pair of reinforcing and balancing loops that interact to produce the S-shaped pattern.

DIFFUSION DYNAMICS

The basic Bass Diffusion Model is usually given as a set of equations:

$$f(t) = dF(t)/dt = [p + qF(t)] [1 - F(t)]$$

$$p, q \geq 0$$

p = coefficient of advertising,
 q = coefficient of interaction.

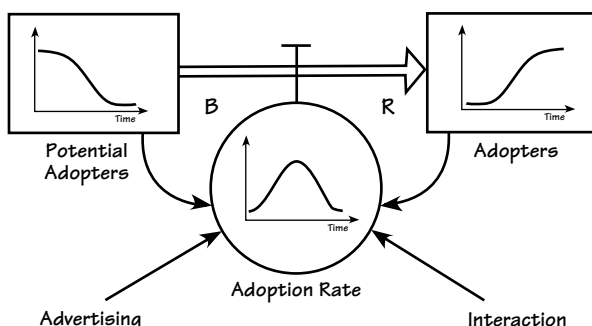
Integration yields an S-shaped growth curve of diffusion.

Although the equations may offer an elegant way to represent such dynamics, most of us don’t view the world as a set of equations. From an accumulator and flows perspective, we see diffusion dynamics as a flow of people from a pool of potential adopters to adopters (see “Bass Diffusion Model—A Structural Viewpoint”).

Instead of p ’s and q ’s, we talk about an advertising effect and a word-of-mouth effect. This structural view makes the dynamics more explicit, closer to the way we actually think about and experience the world.

Faddish products, such as hula hoops and Cabbage Patch dolls, usually exhibit classic S-shaped growth. Remember the sudden popularity of “pet rocks” in the late 1970s? They were just plain old rocks repackaged in a pet-carrier style box. At first, driven by strong advertising, they were viewed as a popular novelty item. Sales began to grow, increasing the demand for the rocks and spreading word-of-mouth endorsements. Soon sales—and the rocks’

BASS DIFFUSION MODEL — A STRUCTURAL VIEWPOINT



The Bass Diffusion Model is an example of S-shaped dynamics. From a structural perspective, it can be represented by a flow of people from a pool of Potential Adopters to Adopters. If we substitute Potential Pet Rock Owners for Potential Adopters, Purchase Rate for Adoption Rate, and Pet Rock Owners for Adopters, the same structural diagram describes the dynamics of the Pet Rock fad.

popularity—began to skyrocket. But eventually the pool of potential adopters (or potential pet rock owners, in this case) was drained, and there was no one left to buy them.

CAPACITY LIMITS

In general, diminishing returns occur whenever we hit a capacity limit. In the Bass Model, capacity is the number of people who can ultimately become adopters of a particular product, technology, or idea. The adoption rate falls to zero when the potential adopters accumulator is depleted, or, in the pet rock case, when everyone comes to their senses (whichever comes first!).

The “Limits to Success” archetype (*THE SYSTEMS THINKER*, December 1990/January 1991) is another way of describing the capacity limit dynamics that produce S-shaped behavior. In a “Limits to Success” structure, a system’s performance improves owing to certain efforts. Better performance results in more efforts, leading to further improvement (loop R1 in “Limits to Success”—From Archetype to Accumulators”). Over time, however, performance begins to plateau despite increased efforts—the system has reached some limit or resistance that is preventing further improvements (loop B2).

If we look at the “Limits to Success” archetype from an accumulator and flow perspective, we can see more clearly the structures producing the unwanted behavior. In a service organization, for example, service capacity may become the limiting factor if it does not keep up with increasing demand. In the beginning, growth in customers will lead to higher revenues, increased marketing, and further growth in customers (R3). This reinforcing loop drives the initial growth of the customers accumulator.

As the number of customers grows, however, so does the service load on the company. If the service capacity does not grow at least as fast as the service load,

capacity adequacy decreases. This leads to lower quality and produces a downward pressure on customer growth. That is, capacity constraints eventually diminish the effectiveness of the marketing efforts.

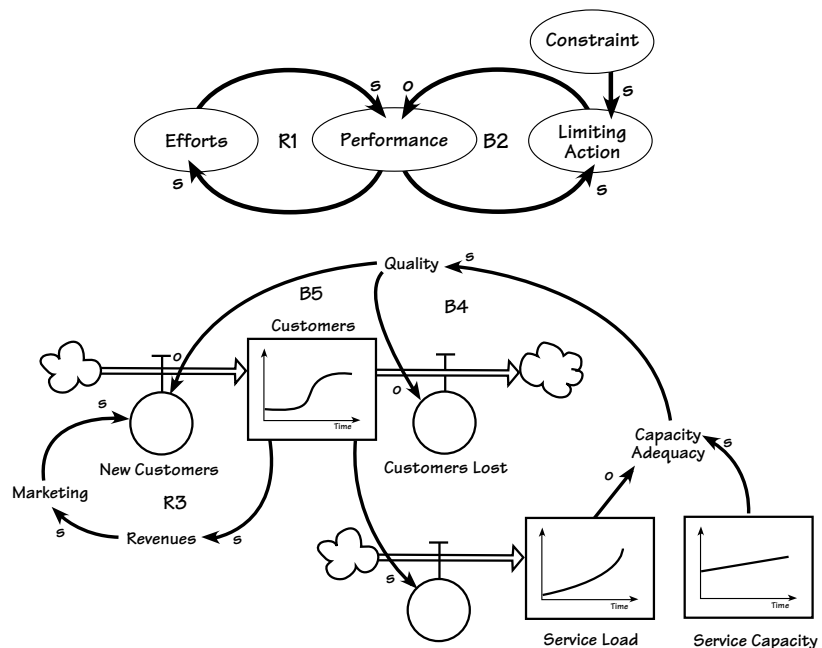
In both the Bass Diffusion Model and the “Limits to Success” archetype, the S-shaped curve is produced by a reinforcing loop coupled with a balancing loop. The reinforcing loop creates the initial growth in demand, while the balancing loop is generally responsible for the diminishing returns. The balancing loops do not suddenly “appear.” They are almost always present from the very start. When the dynamic changes from one of rising growth to a slowing pace, the force driving the system has simply shifted from a reinforcing to a balancing loop. “Saturation” occurs in both cases—whether it is the saturation of a given market or the full utilization of a specific capacity.

BREAKING THE LAW

If you find yourself “caught” by the Law of Diminishing Returns, using a structural diagram may help you identify the critical factors and find a way to break out of it. In a diffusion dynamics case, for example, quantifying and measuring each of the different pieces of the diagram may help decide whether you should try to expand the pool of potential adopters, segment adopters into different categories, beef up the advertising budget, or push on direct sales efforts.

In the more general case of capacity limits, breaking out of the diminishing returns phenomena requires identifying the accumulator(s) that are operating at or near full capacity and calculating the true workload demand. Eliminating any gaps between demand and capacity is likely to produce more results than simply pushing harder on the reinforcing loops in the system. ■

“LIMITS TO SUCCESS”—FROM ARCHETYPE TO ACCUMULATORS



The “Limits to Success” archetype (top) is one way of describing the capacity limitation dynamics that produce S-shaped behavior. Looking at the archetype from an accumulator and flow perspective (bottom), we can see more clearly the structures producing the unwanted behavior.

PART IV

COMPUTER - BASED TOOLS

COMPUTER MODEL

MANAGEMENT FLIGHT SIMULATOR

LEARNING LABORATORY

Many of the systems we are charged with managing are so dynamically complex that they are almost incomprehensible. Complex social systems frequently exhibit counterintuitive behavior, where actions that provide short-term relief often result in greater long-term pain. When actions and consequences are greatly separated in space and time, making effective decisions for the long-term well-being of a system becomes extremely difficult.

Causal loop diagrams, archetypes, and structure-behavior pairs can help us build a better conceptual understanding of the key relevant structures of a system and perhaps even predict the general behavior of the system over time. “Modeling for What Purpose” (p. 40), however, describes times when we need even greater precision about the ramifications of certain actions at specific points in time.

The rest of this section describes how we can translate our pen-and-paper representations into computer-based models that can be simulated, converted into interactive decision-making games (management flight simulators), and embedded in a rich learning environment (learning laboratories).



MODELING FOR WHAT PURPOSE?

BY JAY W. FORRESTER

System dynamics does not impose models on people for the first time—models are already present in everything we do. One does not have a family or corporation or city or country in one’s head. Instead, one has observations and assumptions about those systems. Such observations and assumptions constitute mental models, which we then use as a basis for action.

The ultimate success of a system dynamics investigation depends on a clear initial identification of an important purpose and objective. Presumably a system dynamics model will organize, clarify, and unify knowledge. The model should give people a more effective understanding about an important system that has previously exhibited puzzling or controversial

behavior. In general, influential system dynamics projects are those that change the way people think about a system. Mere confirmation that current beliefs and policies are correct may be satisfying but hardly necessary, unless there are differences of opinion to be resolved. Changing and unifying viewpoints means that the relevant mental models are being altered.

UNIFYING KNOWLEDGE

Complex systems defy intuitive solutions. Even a third-order, linear differential equation is unsolvable by inspection. Yet, important situations in management, economics, medicine, and social behavior usually lose reality if simplified to less than fifth-order nonlinear dynamic systems.

Attempts to deal with nonlinear dynamic systems using ordinary processes of description and debate lead to internal inconsistencies. Underlying assumptions may have been left unclear and contradictory, and mental models are often logically incomplete. Resulting behavior is likely to be contrary to that implied by the assumptions being made about underlying system structure and governing policies.


System dynamics modeling can be effective because it builds on the reliable part of our understanding of systems while compensating for the unreliable part. The system dynamics procedure untangles several threads that cause confusion in ordinary debate: underlying assumptions (structure, policies, and parameters), and implied behavior. By considering assumptions independently from resulting behavior, there is less inclination for people to differ on assumptions (on which they actu-

ally can agree) merely because they initially disagree with the dynamic conclusions that might follow.

If we divide knowledge of systems into three categories, we can illustrate wherein lie the strengths and weaknesses of mental models and simulation models (see “Three Categories of Information”). The top of the figure represents knowledge about structure and policies; that is, about the elementary parts of a system. This is local non-dynamic knowledge. It describes information available at each decision-making point. It identifies who controls each part of a system. It reveals how pressures and crises influence decisions. In general, information about structure and policies is far more reliable, and is more often seen in the same way by different people, than is generally assumed. It is only necessary to dig out the information by using system dynamics insights about how to organize structural information to address a particular set of dynamic issues.

The middle of the figure represents assumptions about how the system will behave, based on the observed structure and policies in the top section. These beliefs are, in effect, the assumed intuitive solutions to the dynamic equations described by the structure and policies in the top section of the diagram. They represent the solutions, arrived at by introspection and debate and compromise, to the high-order nonlinear system described in the top part of the figure. In the middle lie the presumptions that lead managers to change policies or lead governments to change laws. Based on assumptions about how behavior is expected to change, policies and laws in the

THREE CATEGORIES OF INFORMATION



There are three categories of information about a system: knowledge about structure and policies; assumptions about how the system will behave based on the observed structure and policies; and the actual system behavior as it is observed in real life. The usual discrepancy is across the boundary a-a: expected behavior is not consistent with the known structure and policies in the system.

top section are altered in an effort to achieve assumed improved behavior in the middle section.

The bottom of the figure represents the actual system behavior as it is observed in real life. Very often, actual behavior differs substantially from expected behavior. In other words, discrepancies exist across the boundary b-b. The surprise that observed structure and policies do not lead to the expected behavior is usually explained by assuming that information about structure and policies must have been incorrect. Unjustifiably blaming inadequate knowledge about parts of the system has resulted in devoting uncouthed millions of hours to data gathering, questionnaires, and interviews that have failed to significantly improve the understanding of systems.

A system dynamics investigation usually shows that the important discrepancy is not across the boundary b-b, but across the boundary a-a. When a model is built from the observed and agreed-upon structure and policies, the model usually exhibits the actual behavior of the real system. The existing knowledge about the parts of the system is shown to explain the actual behavior. The dissidence in the diagram arises because the intuitively expected behavior in the middle section is inconsistent with the known structure and policies in the top section.

These discrepancies can be found repeatedly in the corporate world. A frequently recurring example in which known corporate policies cause a loss of market share and instability of employment arises from the way delivery delay affects sales and expansion of capacity (see “Underinvestment in Capacity”). Rising backlog (and the accompanying increase in delivery delay) discourage incoming orders for a product (B1) even while management favors larger backlogs as a safety buffer against business downturns. As manage-

ment waits for still higher backlogs before expanding capacity, orders are driven down by unfavorable delivery delay until orders equal capacity (R3). The awaited signal for expansion of capacity never comes because capacity is controlling sales, rather than potential demand controlling capacity (B2).

When sales fail to rise because of long delivery delays, management may then lower price in an attempt to stimulate more sales (B4). Sales increase briefly but only long enough to build up sufficient additional backlog and delivery delay to compensate for the lower prices. In addition, price reductions lower profit margins until there is no longer economic justification for expansion (R5). In such a situation, adequate information about individual relationships in the system is always available for successful modeling, but managers are not aware of how the different activities of the company are influencing one another.

Lack of capacity may exist in manufacturing, product service, skilled sales people, or even in prompt answering of telephones. For example, airlines cut fares to attract passengers. But how often, because of inadequate telephone capacity, are potential customers put on “hold” until they hang up in favor of another airline?

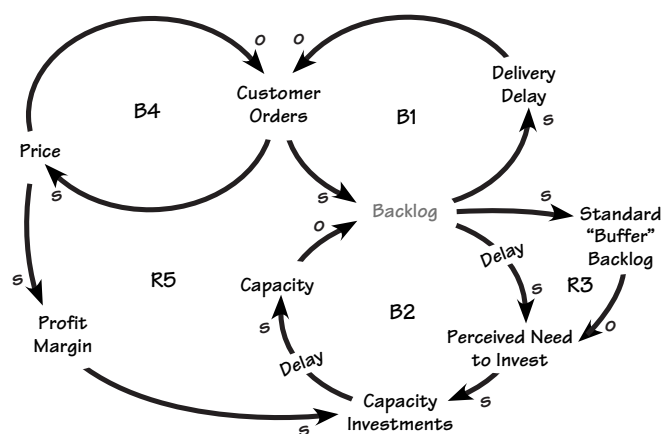
System dynamics models have little impact unless they change the way people perceive a situation. A model must help to organize information in a more understand-

able way. A model should link the past to the present by showing how present conditions arose, and extend the present into persuasive alternative futures under a variety of scenarios determined by policy alternatives. In other words, a system dynamics model, if it is to be effective, must communicate with and modify the prior mental models. Only people’s beliefs—that is, their mental models—will determine action. Computer models must relate to and improve mental models if the computer models are to fill an effective role. ■

Jay W. Forrester, professor emeritus at the Massachusetts Institute of Technology and former director of the MIT System Dynamics Group, is the founder of the field of system dynamics. Since his retirement in 1989, he has been working toward bringing system dynamics into K through 12th grade schools as the basis for a new kind of education.

This article is a selection from “System Dynamics and the Lessons of 35 Years,” in Kenyon B. De Greene (ed.) Systems-Based Approach to Policymaking, (Kluwer Academic Publishers, 1993).

UNDERINVESTMENT IN CAPACITY



Rising backlog dampens customer orders because of increasing delivery delay (B1). However, if management is reluctant to invest in capacity expansions until the backlog reaches a certain level (Standard “Buffer” Backlog), orders will be driven down until demand equals capacity (R3). The awaited signal to expand capacity never comes, because capacity is controlling sales rather than potential demand controlling capacity (B2). If management tries lowering price to stimulate demand (B4), the resulting lower profit margins will further justify a delay in capacity investment (R5).



MANAGEMENT FLIGHT SIMULATORS: FLIGHT TRAINING FOR MANAGERS (PART I)

Imagine you're leaving on a six-hour flight from Boston to Los Angeles. As the plane pulls away from the gate, the pilot comes on over the loudspeaker: "Hi, I'm Captain Bob, and I want to thank you for choosing to fly with us today. . . . Just wanted to let you know I've recently completed ground school training, and I have read all the manuals, but this is my first time in the cockpit. So sit back, relax, and enjoy the flight, as we learn together. . . ."

Of course this scenario is ludicrous—a pilot is allowed into a cockpit only after hundreds of hours of experience in a flight simulator. Then he or she spends many additional hours as a co-pilot, assisting in the operation of an aircraft. The result of this careful system of education and training is an industry with the highest safety record of any mode of transportation.

FLIGHT TRAINING FOR MANAGERS

Imagine if we trained pilots like we do managers; how many people would be willing to take a flight? Managerial train-

ing, in the traditional business-school sense, is the equivalent of ground-school for pilots. Managers-to-be read textbooks and solve already formulated problems, but they don't get much real experience before they have to perform on-line.

MANAGEMENT FLIGHT SIMULATORS

Management Flight Simulators (MFSs) provide a simulated environment in which managers can "learn from experience" in a controlled setting. The simulator captures the interconnections between the different parts of the system under study and provides a computer interface that allows managers to interact with the model through a familiar "lens" (reports, graphs, and spreadsheets).

Similar to a pilot's flight simulator cockpit, an MFS puts managers in control of a realistic environment where they are in charge of making key decisions similar to the ones they face in their actual work settings. MFS's are particularly useful for getting away from the details of day-to-day

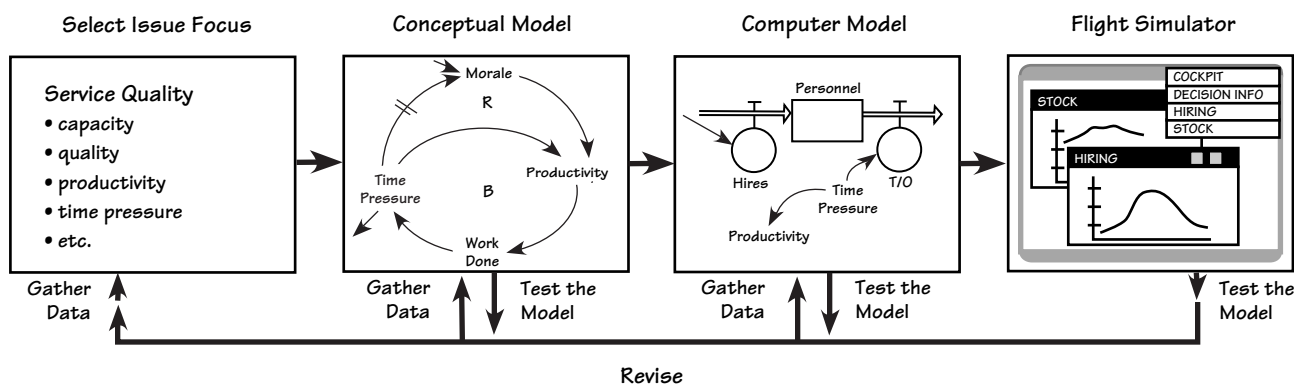
operations and focusing on the long-term dynamics of managerial decisions.

CREATING A FLIGHT SIMULATOR

There are four stages involved in creating a management flight simulator: (1) selecting an issue focus, (2) developing a conceptual model, (3) constructing a computer model, and (4) translating the computer model into an interactive simulator (see "Management Flight Simulator Development Stages"). These four stages involve integrating many of the tools of systems thinking into a single, powerful learning tool (see "A Palette of Systems Thinking Tools," p. 10, for a description of each of the tools).

1. Select issue focus. The first step in designing a flight simulator is to choose an issue to explore. To select a topic, look for a problem symptom that has been around for a long time or a puzzling dynamic you want to investigate (see "The Do's and Don'ts of Systems Thinking on the Job," p. 52, for guidelines on identifying good

MANAGEMENT FLIGHT SIMULATOR DEVELOPMENT STAGES



systems problems). The goal at this stage is to gather relevant data through interviews, company records, and the experience base of those involved in developing the MFS.

For example, let's say we are puzzled by a pattern of oscillating quality levels in the customer service department.

Interviews with people in the department reveal a pattern of tremendous time pressure that repeats in a regular cycle.

Company data provide a record of steadily rising sales and irregular levels of customer satisfaction. This process grounds the project in real data from which to build a causal theory.

2. Build conceptual model. After selecting an issue focus, you can begin to build a conceptual model that organizes the data into a coherent dynamic theory. Systems archetypes and causal loop diagrams (CLDs) can be very helpful for trying to understand what is going on (see "Systems Archetypes at a Glance," p. 20, and "Guidelines for Drawing Causal Loop Diagrams," p. 18).

In the customer service quality example, we can start building causal structures that provide plausible explanations for the observed data. When customer demand increases, we know our service people feel added pressure. If the increase in time pressure is not addressed, quality tends to drift downward and eventually dampens demand (B1 in "Time Pressure Loops").

Our service people tell us people initially respond to the time pressure by working harder, thereby increasing productivity and getting more work done (B2). If the time pressure persists, however, morale declines and begins to hurt productivity even though people continue to work harder. As time pressure escalates, morale spirals downward (R3). By adding loops, we can continue building a dynamic theory about our customer service setting.

3. Construct computer model. The dynamic theory developed in the conceptual model helps guide the construction of

a computer model. It provides a framework for people to distill their experience into explicit statements that can be represented in a computer model. Just as the pilot's flight simulator is created based on the laws of physics and aerodynamics, the computer modeling process uses a set of fundamental building blocks (e.g., accumulators and flows) to represent a coherent set of theories about the interconnections in an organization.

In the customer service quality example, we could model the number of "personnel" as an accumulator and "hires" and "turnover" as inflows and outflows, respectively. The effect of time pressure on turnover may be modeled using a graphical function diagram representing a nonlinear link between the two variables. That is, there may be little or no negative effects at low levels of pressure, but beyond a certain threshold, there may be a sudden dramatic increase in turnover.

4. Translate to flight simulator. Pilots first learn about the principles and concepts of aviation in school and then use the simulator to gain a better understanding of how those principles actually play out in real life. Likewise, a management flight simulator is created by translating the "principles" captured in the computer model into a form that allows managers to interact with it in a realistic way.

A good simulator interface should provide managers with a set of decisions that either they control, or that directly affect them. The main criteria should be that the decisions are directly relevant or easily transferable from the simulator to the workplace.


In the service quality example, the simulator

could require managers to make decisions about hiring/firing, monthly production numbers, and quality standards. By implementing a "Quality First" policy, for example, we may discover that if we raise quality standards but don't adjust capacity, we actually end up with lower quality in the long run. Quality improves in the short run, but as time pressure persists, morale decreases, turnover increases, which in turn increases time pressure, resulting in more turnover. The overall dynamic is a vicious reinforcing cycle in which capacity continually erodes and quality suffers.

The simulator should also provide managers with the same types of reports, spreadsheets, and graphs they use to make decisions. There are many issues involving the design of the simulator's management information system that are entwined with the intended use of the MFS as a whole. These and other issues are covered in Part II (p. 44). ■

For help on converting conceptual maps to computer models, see "Accumulators: Bathtubs, Bathtubs Everywhere," p. 30; "Structural Thinking: The World According to Accumulators and Flows," p. 28; and "Graphical Functions: 'Seeing' the Full Story," p. 26. To learn more about constructing computer models, see Introduction to System Dynamics Modeling (Pegasus Communications), and Academic User's Guide to STELLA (Hanover, NH: High Performance Systems).

TIME PRESSURE LOOPS



Customer demand puts pressure on service personnel. Initially, people may work harder, thus increasing productivity and reducing the pressure (B2). Over time, however, morale can suffer, hurting productivity and increasing the time pressure (R3).



MANAGEMENT FLIGHT SIMULATORS: FLIGHT TRAINING FOR MANAGERS (PART II)

A management flight simulator, along with causal loop diagrams and systems archetypes, allows you to see more clearly the connections between your decisions and future consequences. As simulated months pass in a matter of minutes and the consequences of your actions unfold, an MFS provides a means for making sense of the short-term and long-term effects of your decisions.

Management flight simulators can be most useful for understanding situations in which causality is distant in time and space. When the inherent time lag is particularly long (on the order of months or years) and organizational complexity is high (see “Organizational Complexity”), learning from experience can be fraught with pitfalls. An MFS allows you to leverage your ability to learn from experience in a complex environment.

MANAGING VS. LEARNING

There are two fundamentally different uses for a computer model and simulator—managing and learning. Simulators and models designed to support decision-making in a real operational setting must focus on capturing the operational reality precisely because operational or strategic decisions will be based on those numbers.

Simulators that are designed for learning, on the other hand, are much more concerned with surfacing the tacit mental models that drive managers’ decision-making. Accuracy of specific numbers is not as important as the relevancy of the issues and concepts captured in the simulator; in other words, simulators for learning are idea-rich versus data-rich.

There are several different simulator design criteria to keep in mind when designing an MFS:

- A clear, real-world context provides a real operational focus that engages line managers in learning more about their own issues.

- Face validity: Make the MFS real enough so the simulator grounds people in their own real-life experiences.

- A strong conceptual framework helps make systemic sense out of the complex dynamics (e.g., systems archetypes).

- Conventional and unconventional information systems provide a familiar information environment, as well as an opportunity to explore and experiment with new ones.

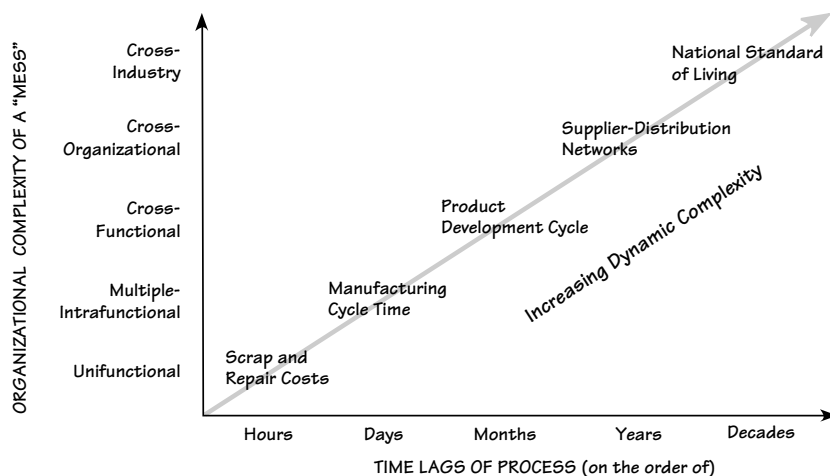
- Surface and challenge mental models to break through individual mental straight-jackets and corporate sacred cows and advance team learning.

DESIGNING MFS’S AS TRANSITIONAL OBJECTS

Designing an MFS for learning requires an interface that maintains a careful balance between realism and comprehensibility. It needs to be real enough to serve as a transitional object, which Seymour Papert says allows managers to “play out” a scenario, learn about the system, and explore how they interact with that system. However, it also needs to be manageable—if the model tries to capture every little detail of reality, it can become just as complex and incomprehensible as the real system.

In a learning setting, it is also important not to position the model as an answer-generator, but rather as an exploratory tool for gaining a better understanding about one’s environment. The MFS acts like a

ORGANIZATIONAL COMPLEXITY



The greater the complexity or the longer the time lags, the more an MFS is beneficial for helping people learn from experience.



mirror that reflects mental models in a way that helps us understand current reality better.

There are three major elements of an MFS: decisions, reports, and a management information system (see below for a sample interface produced with *MicroWorld Creator*).

Decisions. The kinds of decisions made in the simulator should be those the participants would either make themselves in real life or those someone else in the organization would make that affect the participants. The decisions should be directly relevant or easily transferrable from the simulator to the workplace. If the decisions are too far removed, the simulator becomes more of an academic exercise or a game, even when a meaningful context is built around it.

Although the participant might not be the one who makes hiring and staffing decisions at his/her level, for example, these decisions can be included in the MFS because they are still part of the real environment in which the participants manage. In fact, putting that manager or supervisor into the decision-making role can be an illuminating experience: He or she will learn what role they play in the system and realize the challenge of managing from the level above.

Reports. As far as the actual physical design of the simulator interface, there are some general guidelines. The reports should look similar to what people typically receive—they should not provide additional data that is normally not accessible, such as time pressure, or perceived quality by the customer. If these variables need to be included, they should not be as prominent as more typical day-to-day data.

Information systems. Designing the information system provides a lot more flexibility in reporting variables that are normally inaccessible. For example, a critical variable like “time pressure” can help you experiment to see how people may

manage differently when provided with such information.

OUTCOMES FROM THE SIMULATOR

Once participants work with the simulator and understand the theory behind it, they can make connections between the simulator and their real work situation more easily. Participants can also explore what interventions they might make in order to better manage the process: What kinds of adjustments need to be made? What controls do participants need to monitor?

Theories-in-use. The simulator can be a powerful tool for surfacing tacit assumptions, for it reflects participants’ understanding of the system. When someone makes a decision and then explains it with data or information that are not in the model, they explicate their own theories and understanding of what’s going on. For example, in a management flight simulator created for insurance claims managers, participants assumed that settlement dollars were rising because of inflation. However, when they discovered that inflation was not included as part of the simulator, they had to rethink their own understanding of what causes settlement dollars to rise.

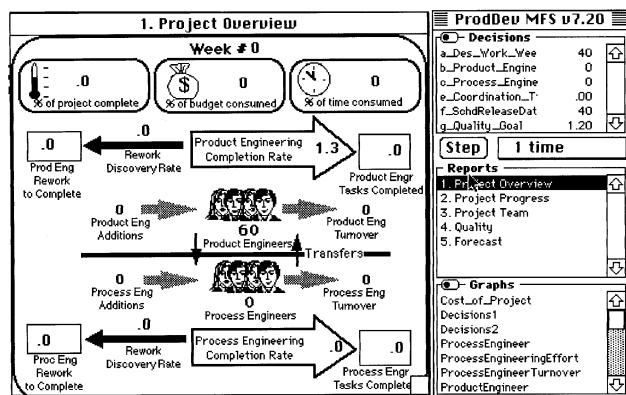
The simulator can also reveal the gap that often exists between espoused theories (what we conceptually believe is the right course of action) and our theories-in-use (what we choose to do, given the surrounding circumstances). For example, in a session with a product development flight simulator, the participants all agreed that investing in coordination and communication between upstream and downstream activities is important. But when they were placed in the simulator and given the objective of meeting tim-

ing, quality and cost, most of them actually chose to invest very little in coordination. Instead, they focused on trying to get the tasks done in each area.

Team learning. A simulator can be even more useful if used in groups. The interplay between the participants, as they propose new strategies and explain their reasoning, helps them to surface and clarify their assumptions. The simulator can be structured to require participation and coordination among a group of people to encourage team learning. For example, in a product engineering case, the team could be made up of a product and a process engineer. Each one would be responsible for staffing and workweek decisions for their particular function, but together they would decide on a schedule completion date and manage the coordination between the two functions. The use of the simulator can therefore be designed to provide a richer practice field for a team to manage.

SUMMARY

MFS’s provide managers with a simulated experience of working through issues or implementing a strategy. The practice field element also enables the simulator to provide an experiential “feel” for the dynamics of decision-making. Participants gain practice in the art of decision-making; reflecting on the consequences, exploring the causal connections, and understanding the underlying structure producing the behaviors. ■





LEARNING LABORATORIES: PRACTICING BETWEEN PERFORMANCES

Imagine you are walking across a tightrope stretched between the world trade towers in New York City. The wind is blowing and the rope is shaking as you inch your way forward. One of your teammates is sitting in the wheelbarrow you are balancing in front of you, while another perches on your shoulders. There are no safety nets, no harnesses. You are thinking to yourself, “One false move and the three of us will be taking an express elevator straight down to the street.” Suddenly your trainer yells from the other side, “Try a new move! Experiment! Take some risks! Remember, you are a learning team!”

Sound ludicrous? No one would be crazy enough to try something new in a situation like that. And yet that is precisely what many companies expect management teams to do—experiment and learn in an environment that is risky, turbulent, and unpredictable. Unlike a high-wire act or sports team, however, management teams do not have a practice field in which to learn; *they are always on the performance field.*

DESIGNING MEANINGFUL PRACTICE FIELDS

A learning lab can be viewed as a manager’s equivalent of a sports team’s practice field. The goal of a learning lab design is to provide a “real” enough practice field so that the lessons are meaningful, but safe enough to encourage experimentation and learning. In the tightrope example, a practice field could be a rope stretched across two pillars six feet off the ground. There may be mats below to cushion the fall, but also a large fan to simulate the kinds of winds you would

encounter in the actual setting.

A managerial practice field should also have its own sets of equipment and tools for making the practice sessions meaningful. The purpose of a “learning laboratory” is to provide an environment in which managers can experiment with alternative policies, test assumptions, and practice working through complex issues productively. It should allow managers to practice working together as a team on issues of real significance to them. To be effective, the learning lab must provide (1) an environment conducive to learning, (2) a way of surfacing deep-rooted assumptions that affect the way we think and act, (3) tools for understanding our reality in a way that highlights the interconnections and the systemic consequences of our actions, and (4) a management flight simulator that allows us to speed up or slow down time, experiment with different strategies, and see the long-term consequences of our actions (see “A Sample Learning Laboratory Design”).

CREATING A SAFE LEARNING SPACE

Learning usually involves making mistakes because we are trying things we have never done before. It requires us to approach things from a place of “not knowing.” It involves risk. How, then, can we create a safe space where people feel free to learn?

There are some ground rules that can help create such safe spaces. One ground rule is to hold each person’s viewpoint as valid. That requires taking the position that “If I could stand in the other person’s shoes, I too could see what the other person sees.”

It does not mean you agree or disagree with that person’s view; you simply acknowledge the right of that person to hold that view. A second rule is to suspend one’s own assumptions and the other person’s and hold them equally in our minds, without judging ours to be superior or “right.” Creating such a learning space also means engaging in dialogue rather than discussion—operating in a spirit of inquiry rather than advocacy.

MAPPING MENTAL MODELS

Along with the proper environment, we need tools for helping people surface and share their assumptions. For example, the “Ladder of Inference,” developed by Chris Argyris, distinguishes between directly observable data, shared cultural meanings, judgments, conclusions, and values and assumptions. Argyris uses the ladder to illustrate the “leaps of inference” that occur when people take a little bit of observed data (a person walks into a 2:00 meeting at 2:15) and go straight up the ladder to the level of values and assumptions (“He’s late and doesn’t care about the project or the other players”) without even being conscious of it. The ladder provides a useful framework for helping people “walk back down the ladder” to understand what is really happening and begin managing by facts, not opinions.

Systems archetypes also provide a powerful set of tools for mapping out a person’s understanding of a problem or issue in a form that invites others to inquire and clarify the picture together. Having one’s assumptions captured in terms of archetypes and causal loop diagrams helps depersonal-



ize the issue and focuses everyone's energy on the diagram, not the person. These diagrams also explicate the assumptions behind the connections, and clarify the points of agreement or contention.

MANAGEMENT FLIGHT SIMULATORS

When practicing a concerto, an orchestra has the ability to slow down or speed up time to practice certain sections. Through management flight simulators (computer models that have been turned into interactive decision-making games), managers can also accelerate time to see the long-term consequences of decisions, or slow down the flow

of time at each decision point. With a simulator, a manager can test out new strategies and policies, reflect on the outcomes, and discuss pertinent issues with others in the team.

By providing quick and accurate feedback, the computer simulator can facilitate learning by shortening the delay between action and outcome. Managers can chart a strategy and implement it over a simulated number of years in a matter of minutes. They can try scenarios that bankrupt the company or lose market share without risking a single dollar or job. As they explore the systemic reasons for their results, managers can begin to understand the underlying

forces that produce a given set of outcomes.

PERFORMING ON THE TIGHTROPE

In order to foster learning among teams of managers, we must look for alternate ways to help them deal with the increasing complexity of today's business world. Providing safe, yet meaningful learning environments where they can continually alternate between practice and performance is one approach. Whether we are walking on a tightrope stretched across two buildings or across two competing product strategies, practice is bound to improve our performance. ■

A SAMPLE LEARNING LABORATORY DESIGN

In designing a Learning Lab (LL), the goal is to create an environment that is of operational relevance. The lab should help managers step out of day-to-day demands to:

- reflect on their decision-making
- develop a common language
- learn new tools for thinking systemically
- discuss operational objectives and strategies in an open forum
- test operating assumptions
- experiment with new policies and strategies for managing
- have fun.

1. The First Crucial Hour—Buy-In

Explaining the context of the LL to participants (the history of its development, the original intent or purpose) is critical for establishing a common understanding: The LL is not meant to provide “the answers,” but to serve as a useful vehicle for illuminating and communicating issues of importance. The facilitators are positioned as enablers, not authority figures, and the participants are encouraged to question the assumptions behind the LL design.

2. Current Reality—Where Are We?

This exercise helps the participants construct a group picture of current problems and issues they face in their jobs. Working in small groups, they are then asked to brainstorm and come up with a list of operational objectives, strategies required to achieve them, and obstacles that need to be overcome in order to reach their goals (e.g., reduce settlement expenses by 10 percent). The idea is to get everyone thinking in terms of their own operational reality.

3. Introducing the Tools

Causal loop diagrams (CLDs) and systems archetypes are introduced in a “storytelling” format in which participants begin to tell systemic stories about their issues. Facilitators then describe a small portion of the CLDs that were in the game model to connect the tool to the issues at hand. The underlying purpose is to get people to immediately begin to connect each structure to corresponding patterns of behavior over time.

4. Using the Tools—Conceptualizing

In small groups, the participants are asked to focus on a particular issue, such as one of the decision variables in the management flight simulator, and (1) determine the key factors that affect that variable, (2) sketch patterns of behavior, (3) provide structural explanation (using CLDs), and (4) identify intervention points. By having the group explore these variables, the participants can replicate part of the model-building process and accept the predeveloped model. The overall objective in this section is to have the group cover all the major issues contained in the model and have a chance to challenge and test the interrelations that different people within the group may propose.

5. Introducing the Computer Simulator

The facilitators begin by showing a simplified CLD that contains all the major variables in the model. They trace through the major loops and explain the dynamic consequences of a particular action or incident, and then sketch a corresponding pattern of behavior and connect it back to the structure. This is followed by a hands-on introduction to the computer and simulator.

Continued on the next page



Continued from the previous page

A SAMPLE LEARNING LABORATORY DESIGN

6. Planned Scenarios—Holding the Reins

In this section, it is best if people work in groups of two at each computer. The teams are instructed to pursue a single-minded strategy where they are accountable for meeting one particular goal (e.g., hiring freeze). Each two-person team is responsible for doing the following:

(1) plan a strategy and commit to it on paper, (2) predict the consequences of executing the strategy by sketching in behavior over time of some key variables, (3) play the game, and (4) debrief game results and explain to the rest of the group. This stage allows participants to begin to address particular organizational issues within a carefully controlled setting.

7. Free Plays—Cutting the Reins

This time, the participants are free to choose their own objectives and time tables. Again, each team strategizes, and then explains to the rest of the group how they plan to achieve their goals. For the designers of the LL, this section provides the opportunity to challenge deep-rooted norms and assumptions, address specific “hot topics,” or re-create various historical behavior modes for further exploration.

PART V

REFERENCE GUIDE



THE VOCABULARY OF SYSTEMS THINKING: A POCKET GUIDE

Systems thinking can serve as a language for communicating about complexity and interdependencies. To be fully conversant in any language, you must gain some mastery of the vocabulary, especially the phrases and idioms unique to that language. This glossary lists many terms that may come in handy when you're faced with a systems problem.

Accumulator—Anything that builds up or dwindles; for example, water in a bathtub, savings in a bank account, inventory in a warehouse. In modeling software, a stock is often used as a generic symbol for accumulators. Also known as **Stock** or **Level**.

Balancing Process/Loop—Combined with reinforcing loops, balancing processes form the building blocks of dynamic systems. Balancing processes seek equilibrium: They try to bring things to a desired state and keep them there. They also limit and constrain change generated by reinforcing processes. A balancing loop in a causal loop diagram depicts a balancing process.

Balancing Process with Delay—A commonly occurring structure. When a balancing process has a long delay, the usual response is to overcorrect. Overcorrection leads to wild swings in behavior. *Example:* real estate cycles.

Behavior Over Time (BOT) Diagram—One of the 10 tools of systems thinking. BOT diagrams capture the history or trend of one or more variables over time. By sketching several variables on one graph, you can gain an explicit understanding of how they interact over time. Also called **Reference Mode**.

Causal Loop Diagram (CLD)—One of the 10 tools of systems thinking. Causal loop diagrams capture how variables in a system are interrelated. A CLD takes the

form of a closed loop that depicts cause-and-effect linkages.

Drifting Goals—A systems archetype. In a “Drifting Goals” scenario, a gradual downward slide in performance goals goes unnoticed, threatening the long-term future of the system or organization.

Example: lengthening delivery delays.

Escalation—A systems archetype. In the “Escalation” archetype, two parties compete for superiority in an arena. As one party's actions put it ahead, the other party “retaliates” by increasing its actions. The result is a continual ratcheting up of activity on both sides. *Examples:* price battles, the Cold War.

Feedback—The return of information about the status of a process.

Example: annual performance reviews return information to an employee about the quality of his or her work.

Fixes That Fail—A systems archetype. In a “Fixes That Fail” situation, a fix is applied to a problem and has immediate positive results. However, the fix also has unforeseen long-term consequences that eventually worsen the problem. Also known as “Fixes That Backfire.”

Flow—The amount of change something undergoes during a particular unit of time. *Example:* the amount of water that flows out of a bathtub each minute, or the amount of interest earned in a savings account each month. Also called a **Rate**.

Generic Structures—Structures that can be generalized across many different settings because the underlying relationships are fundamentally the same. Systems archetypes are a class of generic structures.

Graphical Function Diagram (GFD)—One of the 10 tools of systems thinking. GFDs show how one variable, such as delivery delays, interacts with another, such as sales, by plotting the relationship between

the two over the entire range of relevant values. The resulting diagram is a concise hypothesis of how the two variables interrelate. Also called **Table Function**.

Growth and Underinvestment—A systems archetype. In this situation, resource investments in a growing area are not made, owing to short-term pressures. As growth begins to stall because of lack of resources, there is less incentive for adding capacity, and growth slows even further.

Learning Laboratory—One of the 10 tools of systems thinking. A learning laboratory embeds a management flight simulator in a learning environment. Groups of managers use a combination of systems thinking tools to explore the dynamics of a particular system and inquire into their own understanding of that system. Learning labs serve as a manager's practice field.

Level—*See* **Accumulator**.

Leverage Point—An area where small change can yield large improvements in a system.

Limits to Success—A systems archetype. In a “Limits to Success” scenario, a company or product line grows rapidly at first, but eventually begins to slow or even decline. The reason is that the system has hit some limit—capacity constraints, resource limits, market saturation, etc.—that is inhibiting further growth. Also called “Limits to Growth.”

Management Flight Simulator (MFS)—One of the 10 tools of systems thinking. Similar to a pilot's flight simulator, an MFS allows managers to test the outcome of different policies and decisions without “crashing and burning” real companies. An MFS is based on a system dynamics computer model that has been changed into an interactive decision-making simulator through the use of a user interface.



Policy Structure Diagram—One of the 10 tools of systems thinking. Policy structure diagrams are used to create a conceptual “map” of the decision-making process that is embedded in an organization. It highlights the factors that are weighed at each decision point.

Rate—See *Flow*.

Reference Mode—See *Behavior Over Time Diagram*.

Reinforcing Process/Loop—Along with balancing loops, reinforcing loops form the building blocks of dynamic systems. Reinforcing processes compound change in one direction with even more change in that same direction. As such, they generate both growth and collapse. A reinforcing loop in a causal loop diagram depicts a reinforcing process. Also known as vicious cycles or virtuous cycles.

Shifting the Burden—A systems archetype. In a “Shifting the Burden” situation, a short-term solution is tried that successfully solves an ongoing problem. As the solution is used over and over again, it takes attention away from more fundamental, enduring solutions. Over time, the ability to apply a fundamental solution may decrease, resulting in more and more reliance on the symptomatic solution. *Examples:* drug and alcohol dependency.

Shifting the Burden to the Intervener—A special case of the “Shifting the Burden” systems archetype that occurs when an intervener is brought in to help solve an ongoing problem. Over time, as the intervener successfully handles the problem, the people within the system become less capable of solving the problem themselves. They become even more dependent on the intervener. *Example:* ongoing use of outside consultants.

Simulation Model—One of the 10 tools of systems thinking. A computer model that lets you map the relationships that are

important to a problem or an issue and then simulate the interaction of those variables over time.

Stock—See *Accumulator*.

Structural Diagram—Draws out the accumulators and flows in a system, giving an overview of the major structural elements that produce the system’s behavior. Also called flow diagram or accumulator/flow diagram.

Structure-Behavior Pair—One of the 10 tools of systems thinking. A structure-behavior pair consists of a structural representation of a business issue, using accumulators and flows, and the corresponding behavior over time (BOT) diagram for the issue being studied.

Structure—The manner in which a system’s elements are organized or interrelated. The structure of an organization, for example, could include not only the organizational chart but also incentive systems, information flows, and interpersonal interactions.

Success to the Successful—A systems archetype. In a “Success to the Successful” situation, two activities compete for a common but limited resource. The activity that is initially more successful is consistently given more resources, allowing it to succeed even more. At the same time, the activity that is initially less successful becomes starved for resources and eventually dies out. *Example:* the QWERTY layout of typewriter keyboards.

System Dynamics—A field of study that includes a methodology for constructing computer simulation models to achieve better understanding of social and corporate systems. It draws on organizational studies, behavioral decision theory, and engineering to provide a theoretical and empirical base for structuring the relationships in complex systems.

System—A group of interacting, interrelated, or interdependent elements forming a

complex whole. Almost always defined with respect to a specific purpose within a larger system. *Example:* An R&D department is a system that has a purpose in the context of the larger organization.

Systems Archetypes—One of the 10 tools of systems thinking. Systems archetypes are the “classic stories” in systems thinking—common patterns and structures that occur repeatedly in different settings.

Systems Thinking—A school of thought that focuses on recognizing the interconnections between the parts of a system and synthesizing them into a unified view of the whole.

Table Function—See *Graphical Function Diagram*.

Template—A tool used to identify systems archetypes. To use a template, you fill in the blank variables in causal loop diagrams.

Tragedy of the Commons—A systems archetype. In a “Tragedy of the Commons” scenario, a shared resource becomes overburdened as each person in the system uses more and more of the resource for individual gain. Eventually, the resource dwindles or is wiped out, resulting in lower gains for everyone involved. *Example:* the Greenhouse Effect. ■

The above glossary is a compilation of definitions from many sources, including:

- *Innovation Associates’ and GKA’s Introduction to Systems Thinking coursebooks*
- *The Fifth Discipline: The Art and Practice of the Learning Organization, by Peter Senge*
- *High Performance Systems’ Academic User’s Guide to STELLA*
- *The American Heritage Dictionary and The Random House Dictionary.*



THE DO'S AND DON'T'S OF SYSTEMS THINKING ON THE JOB

BY MICHAEL GOODMAN

So, you've taken a systems thinking course—or maybe you've read a few issues of *The Systems Thinker*—and now you want to start using systems thinking on the job. How do you begin? Your best bet is to approach this endeavor in the spirit of “learning to walk before you run.” Here are some suggestions:

OVERALL GUIDELINE

The tools of systems thinking are best used as vehicles to promote team learning in the organization. Whether you are doing “paper and pencil” models or creating full-fledged microworlds, the process of constructing and using models is primarily about exploring and examining our “mental models”—the deeply held assumptions that influence the way we think and act.

IDENTIFYING A SYSTEMS PROBLEM

The problem should have ALL of the following characteristics:

1. The issue is important to me and my business.
2. The problem is chronic, rather than a one-time event.
3. The problem has a known history that I can describe.

Example: Profits were steady for 2 years, but have been declining for the last 6 months. Or: Productivity rose rapidly until about a year ago, when it leveled off.

4. People have tried to solve this problem before, with little or no success.

If your problem does not have all of these characteristics (especially the first three), it may not be appropriate for a systems thinking analysis. Try redefining it for a different approach.

GENERAL GUIDELINES

DON'T use systems thinking to further your own agenda. Systems thinking is most effective when it is used to look at a problem in a new way, not to advocate a predetermined solution. Strong advocacy will create resistance—both to your ideas, and to systems thinking. It should be used in the spirit of inquiry, not inquisition.

DO use systems thinking to sift out major issues and factors.

Benefit: Systems thinking can help you break through the clutter of everyday events to recognize general patterns of behavior and the structures that are producing them. It also helps in separating solutions from underlying problems. Too often we identify problems in terms of their solution; for example, “the problem is that we

have too many _____ (fill in the blank: people, initiatives, steps in our process),” or “the problem is that we have too little _____ (resources, information, budget . . .).”

DON'T use systems thinking to blame individuals. Chronic, unresolved problems are more often the result of systemic breakdowns than individual mistakes. Solutions to these problems lie at the systemic, not the individual, level.

DO use systems thinking to promote inquiry and challenge preconceived ideas.

Cues that non-systemic thinking is going on: Phrases such as “We need to have immediate results,” “We just have to do more of what we did last time,” or “It’s just a matter of trying harder.”

GETTING STARTED

DON'T attempt to solve a problem immediately. Don't expect to represent, much less understand, persistent and complex systemic problems overnight. The time and concentration required should be proportional to the difficulty and scope of the issues involved.

More realistic goal: to achieve a fuller and wider understanding of the problem.

DO start with smaller-scale problems.

DON'T attempt to diagram the whole system—otherwise you'll quickly become overwhelmed.

Better: Try to focus on a problem issue and draw the minimum variables and loops you'll need to capture the problem.

DON'T work with systems thinking techniques “on line” under pressure, or in front of a group that is unprepared for or intolerant of the learning process.

Additional danger: If the audience is not familiar with the concepts and methods of systems thinking, they might not understand that the process reveals mental models, can be controversial, and is highly iterative in nature. It is far more beneficial to have the group engage in their own loop building after appropriate instruction and foundation have been given.

DO develop your diagrams gradually and informally, in order to build confidence in using systems thinking.



Good practice: Look at newspaper articles and try to draw a few loops that capture the dynamics of a problem being described. Even better: Try matching a template to the article.

DON'T worry about drawing loops right away. One of the strongest benefits of the systems thinking perspective is that it can help you learn to ask the right questions. This is an important first step toward understanding a problem.

DRAWING DIAGRAMS

DO start with the process of defining variables. **DO** encourage airing of assumptions.

Benefit: better shared understanding of a problem. Diagramming is a very effective tool for promoting group inquiry into a problem or issue.

DO start with a central loop or process. Then add loops to “fill in” detail.

Example: The central loop may show how the system is supposed to work, and the additional loops can explore what is pushing it out of whack.

DON'T get bogged down in details. Start simply, at a high level of generalization, but with enough detail to sum up the observed behavior.

Example: If you are exploring the causes of missed delivery dates in a factory, lump together the types of products that are experiencing similar delays.

DO begin by looking for templates or general structures that might clarify the problem.

Advantage: Systems archetypes provide a focal point or a storyline to begin the process of understanding a problem.

DO work with one or more partners.

Advantage: Multiple viewpoints add richness and detail to the understanding of a problem.

DO check with others to see if they can add some insight or improve upon your diagram—especially people in other functional areas who might have a different

perspective on the problem.

Example: With a manufacturing delay problem, you might check with finance to see if there are any dynamics in the finance arena that are affecting the manufacturing delays (capital investments and purchases, etc.). The same can be done for marketing, sales, etc.

DO work iteratively. There is no “final” model (set of loops). Looping is a learning process that should continue to evolve with new data and perspectives.

DON'T present “final” loop diagrams as finished products.

Better: Present as a tentative and evolving picture of how you are seeing things. To get buy-in and maximize learning, the audience needs to participate in the modeling process.

DO learn from history. When possible, check data to see if your diagram correctly describes past behavior.

INTERVENTIONS

DO get all stakeholders involved in the process. This will help ensure that all viewpoints have been considered, and will improve the acceptance rate for the intervention.

DON'T go for vague, general, or open-ended solutions such as “Improve communications.”

Better: “Reduce the information delay between sales and manufacturing by creating a new information system.”

DO make an intervention specific, measurable, and verifiable.

Example: “Cut the information delay between sales and manufacturing down to 24 hours.”

DO look for potential

unintended side-effects of an intervention.

General principle: “Today’s problems often come from yesterday’s solutions.” Any solution is bound to have trade-offs, so use systems thinking to explore the implications of any proposed solution before trying to implement it.

DON'T be surprised if some situations defy solution, especially if they are chronic problems. Rushing to action can thwart learning and ultimately undermine efforts to identify higher leverage interventions. Resist the tendency to “solve” the issue and focus on gaining a deeper understanding of the structures producing the problem. Be wary of a symptomatic fix disguised as a long-term, high-leverage intervention. ■

Michael Goodman is a principal of Innovation Associates, Framingham, MA, an Arthur D. Little company. The material in this article was drawn from his 20 years of experience in the field, as well as business courses developed by Innovation Associates.

GUIDELINES FOR INTERVENTIONS

1. To be effective, an intervention must be self-sustaining, self-correcting and long-lasting. It must make long-term changes in the performance trend.
2. Types of interventions in a causal loop diagram:
 - Add a link.
 - Break a link.
 - Shorten a delay.
 - Make a goal explicit.
 - Slow down a growth process; relieve a limiting process.
3. The best intervention is likely to be a combination of interventions applied gently and patiently.
4. Avoid pushing on a structure from the outside.
5. Look for variance between long- and short-term impacts, to anticipate unexpected effects.

FURTHER READING ON THE 10 TOOLS

1. Double-Q Diagram

Based on TQC tool “Cause-and-Effect Diagram.” See Ishikawa, Kaoru (1982) *Guide to Quality Control*, Ann Arbor, MI: UNIPUB.

2. Behavior Over Time Diagram

Based on diagrams referred to as “reference modes” in system dynamics literature. See Richardson, George and Alexander Pugh (1981) *Introduction to System Dynamics Modeling*, Waltham, MA: Pegasus Communications.

3. Causal Loop Diagram

See Richardson, George and Alexander Pugh (1981) *Introduction to System Dynamics Modeling*, Waltham, MA: Pegasus Communications.

4. Systems Archetypes

See Senge, Peter (1990) *The Fifth Discipline*, New York: Doubleday. Also covered regularly with current business applications in a management newsletter, *THE SYSTEMS THINKER*, published by Pegasus Communications, Inc., Waltham, MA.

5. Graphical Function Diagram

Based on diagrams referred to as “table functions” in system dynamics literature. See Richardson, George and Alexander Pugh (1981) *Introduction to System Dynamics Modeling*, Waltham, MA: Pegasus Communications.

6. Structure-Behavior Pairs

Referred to as “Atoms of Structure” in *Academic User's Guide to STELLA* by Barry Richmond, published (as part of software documentation) by High Performance Systems, Hanover, NH. Also, see Goodman, Michael (1974) *Study Notes in System Dynamics*, Waltham, MA: Pegasus Communications.

7. Policy Structure Diagram

Contact Professor John Morecroft at the London Business School.

8. Computer Model

One of the best software for building system dynamics computer models (Macintosh) is *ithink*TM and *STELLA*TM by High Performance Systems, Hanover, NH. For IBM-compatibles, there is *Vensim* by Ventana Systems, and *PowerSim Studio Enterprise 2000* by PowerSim Corp.

9. Management Flight Simulators

Contact Professor John Sterman at the M.I.T. Sloan School of Management (617-253-1951) for copies of computer simulators on People Express, managing product lifecycles, real-estate management, and super tanker management.

10. Learning Laboratory

Kim, Daniel (1989) “Learning Laboratories: Designing a Reflective Learning Environment,” Proceedings of the 1989 International System Dynamics Conference, Stuttgart, Germany: Springer-Verlag.

INDEX TO THE SYSTEMS THINKER

PART I: AN OVERVIEW

- 6 Systems Thinking as a Language *by Michael R. Goodman* V2N3, April 1991
8 Levels of Understanding: “Fire-Fighting” at Multiple Levels V4N5, June/July 1993
10 A Palette of Systems Thinking Tools V1N3, August 1990

13 PART II: DYNAMIC THINKING TOOLS

- 14 Reinforcing and Balancing Loops: Building Blocks of Dynamic Systems V1N1, April/May 1990
16 Balancing Loops with Delays: Teeter-Tottering on Seesaws V1N2, June/July 1990
18 Guidelines for Drawing Causal Loop Diagrams V3N1, February 1992
20 Systems Archetypes at a Glance V3N4, May 1992

23 PART III: STRUCTURAL THINKING TOOLS

- 24 From Causal Loops to Graphical Functions: Articulating Chaos V2N8, October 1991
26 Graphical Functions: “Seeing” the Full Story V2N7, September 1991
28 Structural Thinking: The World According to Accumulators and Flows V2N2, March 1991
30 Accumulators: Bathtubs, Bathtubs, Everywhere V2N1, February 1991
32 Accumulation Management: Avoiding the “Pack Rat” Syndrome V2N4, May 1991
34 Delays: Accumulators in Disguise V2N5, June/July 1991
36 S-Shaped Growth and the Law of Diminishing Returns V2N3, April 1991

39 PART IV: COMPUTER-BASED TOOLS

- 40 Modeling for What Purpose? *by Jay W. Forrester* V4N4, May 1993
42 Management Flight Simulators: V3N9, November 1992
 Flight Training for Managers—Part I
44 Management Flight Simulators: V3N10, Dec 1992/Jan 1993
 Flight Training for Managers—Part II
46 Learning Laboratories: Practicing Between Performances V3N8, October 1992

49 PART V: REFERENCE GUIDE

- 50 The Vocabulary of Systems Thinking: A Pocket Guide V2N10, Dec 1991/Jan 1992
52 The Do’s and Don’t’s of Systems Thinking on the Job V3N6, August 1992
 by Michael R. Goodman

ABOUT THE TOOLBOX REPRINT SERIES

Systems Thinking Tools: A User's Reference Guide is a volume in the Toolbox Reprint Series. Other volumes include *Systems Archetypes I: Diagnosing Systemic Issues and Designing High-Leverage Interventions*, *Systems Archetypes II: Using Systems Archetypes to Take Effective Action*, *Systems Archetypes III: Understanding Patterns of Behavior and Delay*, and *The "Thinking" in Systems Thinking: Seven Essential Skills*. All volumes are available for \$16.95 each. As these booklets are often used in training and introductory courses, volume discounts are available. See the copyright page or call 1-800-272-0945 for details.

The Toolbox Reprint Series has been compiled from *The Systems Thinker*[®] Newsletter, which presents a systems perspective on current issues and provides systems tools for framing problems in new and insightful ways. *The Systems Thinker* includes articles by leading systems thinkers, case studies of systems thinking implementation, software and book reviews, and numerous other columns geared to different levels of systems thinking ability. Individual subscription rates to *The Systems Thinker* are as follows: A special rate of 20 issues for \$169, or 10 issues (one year) for \$109. Library subscriptions are available for \$189 for 10 issues (one year). Visit www.pegasus.com for more information.

Anthologies

Managing the Rapids: Stories from the Forefront of the Learning Organization
Reflections on Creating Learning Organizations
The New Workplace: Transforming the Character and Culture of Our Organizations
Organizational Learning at Work: Embracing the Challenges of the New Workplace
Making It Happen: Stories from Inside the New Workplace

The Pegasus Workbook Series

Systems Archetype Basics: From Story to Structure *Systems Thinking Basics: From Concepts to Causal Loops*
The Learner's Path: Practices for Recovering Knowers

The "Billibonk" Series

Billibonk & the Thorn Patch *Frankl's "Thorn Patch" Fieldbook*
Billibonk & the Big Itch *Frankl's "Big Itch" Fieldbook*

Learning Fables

Outlearning the Wolves: Surviving and Thriving in a Learning Organization
Shadows of the Neanderthal: Illuminating the Beliefs That Limit Our Organizations
The Lemming Dilemma: Living with Purpose, Leading with Vision
The Tip of the Iceberg: Managing the Hidden Forces That Can Make or Break Your Organization
Listening to the Volcano: Conversations That Open Our Minds to New Possibilities

Other Titles

Human Dynamics: A New Framework for Understanding People and Realizing the Potential in Our Organizations
When a Butterfly Sneezes: A Guide for Helping Kids Explore Interconnections in Our World Through Favorite Stories

The Innovations in Management Series

From Mechanistic to Social Systemic Thinking: A Digest of a Talk by Russell L. Ackoff
Applying Systems Archetypes
Toward Learning Organizations: Integrating Total Quality Control and Systems Thinking
Designing a Systems Thinking Intervention: A Strategy for Leveraging Change
The Natural Step: A Framework for Achieving Sustainability in Our Organizations
Anxiety in the Workplace: Using Systems Thinking to Deepen Understanding
The Soul of Corporate Leadership: Guidelines for Values-Centered Governance
Creating Sustainable Organizations: Meeting the Economic, Ecological, and Social Challenges of the 21st Century
Creating Value: Linking the Interests of Customers, Employees, and Investors
Relinking Life and Work: Toward a Better Future
Facing the Competition: An Organization Mobilizes for Large-Scale Change
Organizational Change at Philips Display Components: Reflections on a Learning Journey
Introduction to Systems Thinking
Rebounding, Rebuilding, Renewing at Shell Oil: A Former CEO Reflects on Large-Scale Change
Reinventing Human Resources at L.L. Bean: Lessons for Learning and Change
The Essentials of Servant-Leadership: Principles in Practice
Dialogue at Work: Skills for Leveraging Collective Intelligence

The Toolbox Reprint Series

Systems Archetypes I: Diagnosing Systemic Issues and Designing High-Leverage Interventions
Systems Archetypes II: Using Systems Archetypes to Take Effective Action
Systems Archetypes III: Understanding Patterns of Behavior and Delay
Systems Thinking Tools: A User's Reference Guide
The "Thinking" in Systems Thinking: Seven Essential Skills

E-Newsletters

THE SYSTEMS THINKER® *LEVERAGE POINTS® for a New Workplace, New World*

Pegasus Communications, Inc. is dedicated to providing resources that help people explore, understand, articulate, and address the challenges they face in managing the complexities of a changing world. Since 1989, Pegasus has worked to build a community of practitioners through newsletters, books, audio and video tapes, and its annual *Systems Thinking in Action*[®] Conference and other events. For more information, contact us at:

Pegasus Communications, Inc. • One Moody Street • Waltham, MA 02453-5339 USA • www.pegasuscom.com
Order Phone: 800-272-0945 / 781-398-9700 • Fax: 781-894-7175