



An Introduction to Systems Thinking

ithink[®] Software

iseesystems[™] • 46 Centerra Parkway, Suite 200 • Lebanon, NH 03766
Phone: (603) 643.9636 • Fax: (603) 643.9502
Technical Support: support@iseesystems.com • Pricing & Sales: orders@iseesystems.com
Training: workshops@iseesystems.com • To Order: orders@iseesystems.com
Visit us on the Web at: <http://www.iseesystems.com>
ISBN 0-9704921-0-3

Introduction to Systems Thinking, *ithink* ® ©1992-1997, 2000, 2001, 2004
isee systems, inc. All rights reserved.

No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording, or otherwise, without prior written permission from isee systems, inc.

ithink is a registered trademark of isee systems, inc. Macintosh is a trademark of Apple Computer, Inc. Windows is a trademark of Microsoft Corporation. Other brand names and product names are trademarks or registered trademarks of their respective companies.

isee systems, inc.'s Licensor makes no warranties, express or implied, including without limitation the implied warranties of merchantability and fitness for a particular purpose, regarding the software. isee systems, inc.'s Licensor does not warrant, guaranty, or make any representations regarding the use or the results of the use of the software in terms of its correctness, accuracy, reliability, currentness, or otherwise. The entire risk as to the results and performance of the software is assumed by you. The exclusion of the implied warranties is not permitted by some states. The above exclusion may not apply to you.

In no event will isee systems, inc.'s Licensor, and their directors, officers, employees, or agents (collectively isee systems, inc.'s Licensor) be liable to you for any consequential, incidental, or indirect damages (including damages for loss of business profits, business interruption, loss of business information, and the like) arising out of the use of, or inability to use, the software even if isee systems, inc.'s Licensor has been advised of the possibility of such damages. Because some states do not allow the exclusion or limitation of liability for consequential or incidental damages, the above limitations may not apply to you.

Contents

Part 1.	Setting the Stage	1
Chapter 1.	A Pressing Need: <i>Improving Performance</i>	3
Chapter 2.	Systems Thinking & the itthink Software <i>Better Mental Models, Simulated More Reliably</i>	13
Part 2.	Learning to “Write” Using the Language of Systems Thinking	29
Chapter 3.	Nouns & Verbs <i>Operational Thinking</i>	31
Chapter 4.	Writing Sentences <i>Operational Thinking</i>	43
Chapter 5.	Linking Sentences <i>Operational Thinking</i>	49
Chapter 6.	Constructing Simple Paragraphs <i>Closed-Loop Thinking</i>	59
	Appendix: Generic Flow Templates	70
Chapter 7.	Constructing “More Interesting” Paragraphs <i>Closed-loop & Non-linear Thinking</i>	75
	Appendix: Formulating Graphical Functions	86
Chapter 8.	Storylines, Part 1: <i>Main Chain Infrastructures</i>	91
Chapter 9.	Storylines, Part 2: <i>Support Infrastructures</i>	105

Part 3.	“Writing” Short Stories	127
Chapter 10.	An Overview of the “Writing” Process	129
Chapter 11.	Illustrating the “Writing” Process	135
Chapter 12.	Guidelines for the “Writing” Process	155
	Appendix: <i>Initializing Your Models in Steady-State</i>	176
Chapter 13.	Adding Texture to Your Compositions <i>Modeling “Soft” Variables</i>	179
	List of Figures	185
	Index	189

Part 1

Setting the Stage

The two chapters in this Part of the Guide provide context for what follows in the remainder of the Guide.

Chapter 1 surfaces a pressing challenge that virtually all organizations face: *How to create performance-improvement initiatives capable of achieving their intended impacts*. A look at the record suggests that...be it reengineering a set of processes, seeking to realize synergies inherent in a merger or acquisition, developing a successful growth strategy, implementing a change effort capable of sustaining change, creating an effective set of operating policies, or devising a useful Balanced Scorecard, the “fixes” too often fail—frequently, in fact, often exacerbating the very situations they were intended to improve!

The Chapter argues that the cause of our disappointing record within the performance-improvement arena is the poor quality of our underlying mental models, and the unreliability of the associated mental simulations. The conclusion is that finding ways to improve both is the key to meeting the performance improvement challenges we face.

Chapter 2 offers Systems Thinking as a framework, and the *ithink*® software as an associated key tool, that can significantly contribute to improving the quality of our mental models and the reliability of the associated simulations. In Chapter 2, a core set of eight Systems Thinking skills is identified. Each of the chapters in Parts 2 and 3 of the Guide then focuses on helping you develop one or more of these skills in the context of learning to use the *ithink* language to construct progressively better mental models.

Chapter 1

A Pressing Need:

Improving Performance

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

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

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

The first thing to note about thinking is that when we ponder something, we do not actually have that “something” in our head. Think about it...You're trying to figure out whether you should let your kid drive to the party. You're struggling to decide whether to quit your steady, but relatively unchallenging day job, to pursue the wild and woolly challenges of a start-up. You're wondering about the best way it is you are thinking about, you do not have it in your head.

Getting to Root Cause

**“What’s up”
with our Mental
Model
Construction
& Simulation
Processes?**

Then, what do you have in there? What are you working with when you’re “thinking?”

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

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

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

So, let’s review the bidding...When we create any sort of performance-improvement initiative, we think. And, when we think, we construct, and then simulate, a mental model. Therefore, if our performance-improvement initiatives come up short of the mark, it is reasonable to suspect that something is awry in the processes by which we construct and simulate our mental models.

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

What follows is a passage that describes a very simple supply chain. Use it to construct a mental model. Then, simulate the model in order

to predict how the system will perform in response to the “disturbance” to which it will be exposed.

A retailer maintains an inventory of product that is shipped to customers on demand. Upon shipping, the retailer orders more product (to re-stock inventory) from the firm that supplies it. The retailer always emails an order to the supplier for an amount of product exactly equal to what was shipped in a given day. If ten units go out in a day, the retailer emails an order for ten units at the end of the day. The retailer never delays in placing the order, and always orders *exactly* the amount of product that was shipped in a given day.

The supplier also is very regular. The supplier always processes the retailer’s order immediately upon receipt, then ships the requested amount of product to the retailer. Product always arrives six days after the retailer places the order. The supplier has never been out-of-stock (and never will be!), and has always (and will always) be able to get product to the retailer exactly six days after the retailer’s order is placed. Furthermore, no product shipped by the supplier is ever, or will ever be, defective, damaged or lost in transit.

This simple supply chain has been in steady-state for some time. This means that the volume of product being demanded at retail by customers has been constant at some level for a long time, as has therefore the volume of product the retailer has been ordering from the supplier, as well as the amount the supplier has been shipping to the retailer. Everything is in perfect, constant balance. Now suppose, all of a sudden, the volume of demand from customers coming into the retailer steps up to a new higher level, and then remains there (i.e., a one-time, step-increase occurs). On the axes provided in Figure 1-1, sketch the *pattern* you think will be traced by the level of the retailer’s inventory, over time, following the one-time step-increase in customer demand.

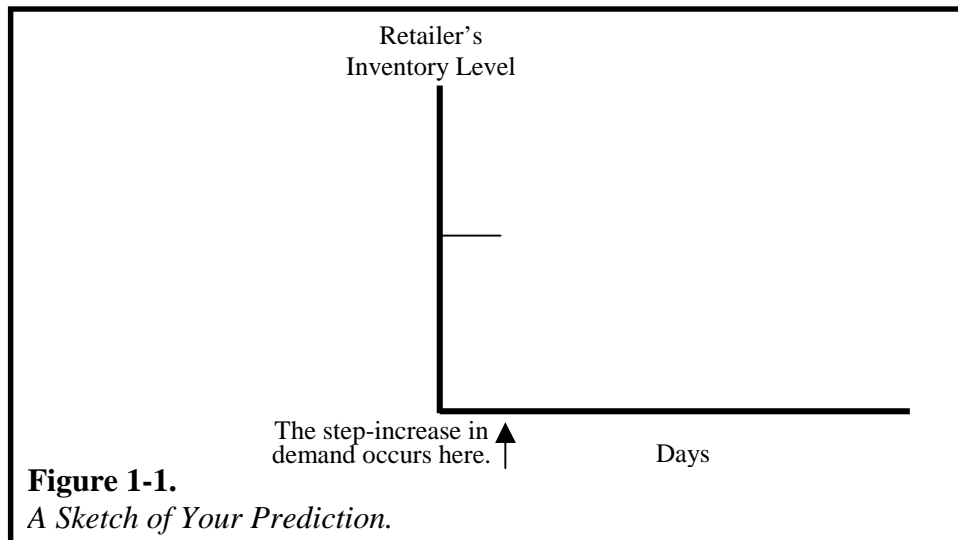


Figure 1-1.
A Sketch of Your Prediction.

Typically, upwards of 80% of any group who is asked to conduct this type of thought experiment traces an incorrect pattern! The correct pattern is that: *following the step-increase in demand, the Retailer’s inventory will decline in a straight-line manner for six days; it then*

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

*Our Simulation
Machinery*

will level off and remain at the new, lower level. (You'll develop an understanding of why in the next chapter). The relatively small percentage of people who *do* trace the correct pattern has proven to be independent of culture, education level, or experience with supply chains. These results strongly suggest that human beings, in general, either are not very good at constructing mental models (of even very simple systems!), performing mental simulations of these models, or both!

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

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

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

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

*Reason 1 for
Poor Quality
Mental Models:
Content*

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

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

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

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

So, while almost any action to improve performance taken today has extensive ramifications, both spatial and temporal, the *contents* of our mental models (i.e., the associated boundaries) do not allow us to

“think through” these ramifications! As a result, we get “surprised” a lot—and usually the surprises are not pleasant. In addition, because we don’t capture the ramifications, it’s not possible to *learn* from them! Hence, we are destined to re-live past mistakes. Figure 1-2 depicts the situation...

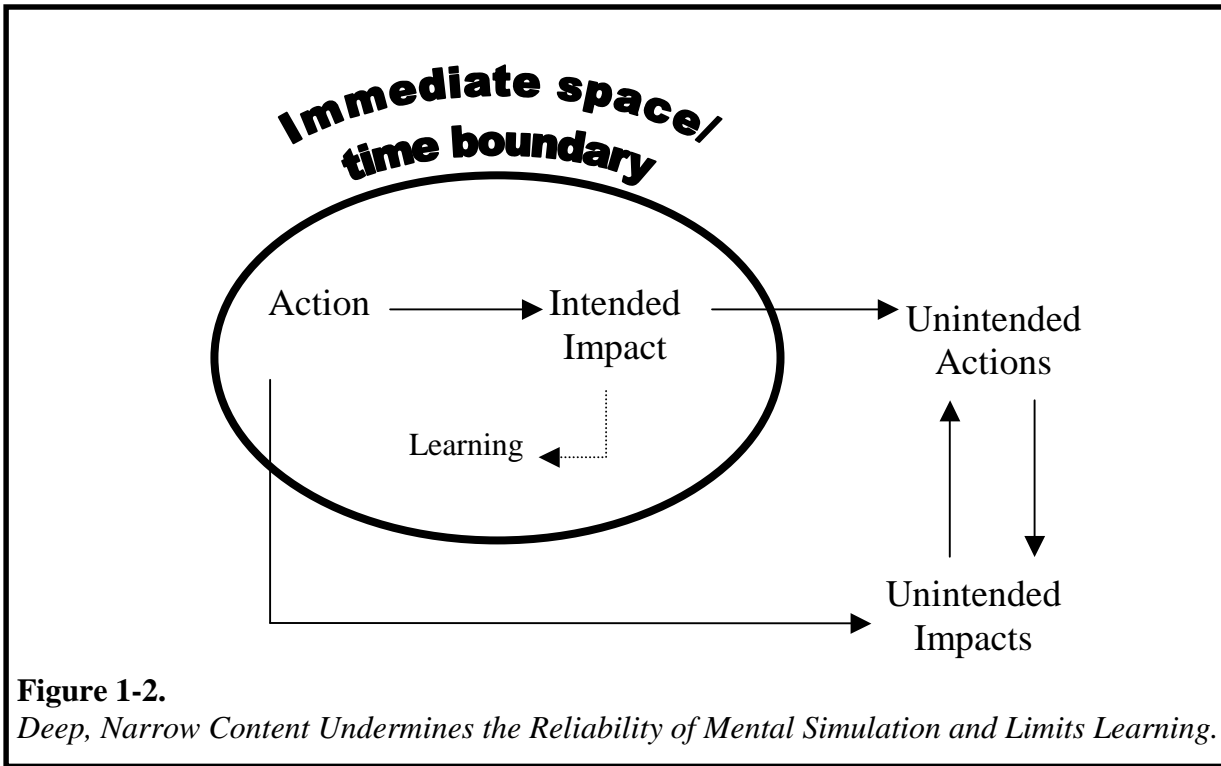


Figure 1-2.
Deep, Narrow Content Undermines the Reliability of Mental Simulation and Limits Learning.

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

*Reason 2 for
Poor Quality
Mental Models:
Representation
of Content*

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

Because we make such extensive use of “meta assumptions,” they submerge...outright disappear from consciousness! They become so “obviously true,” they’re no longer subject to scrutiny or question. But

if we are to have any hope of improving upon these assumptions, we must first bring them back into view. One way to surface them is to identify conceptual frameworks and analytical tools that are in widespread use in diverse arenas. The fact that they are widely used suggests they mask a set of commonly embraced “meta assumptions.”

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

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

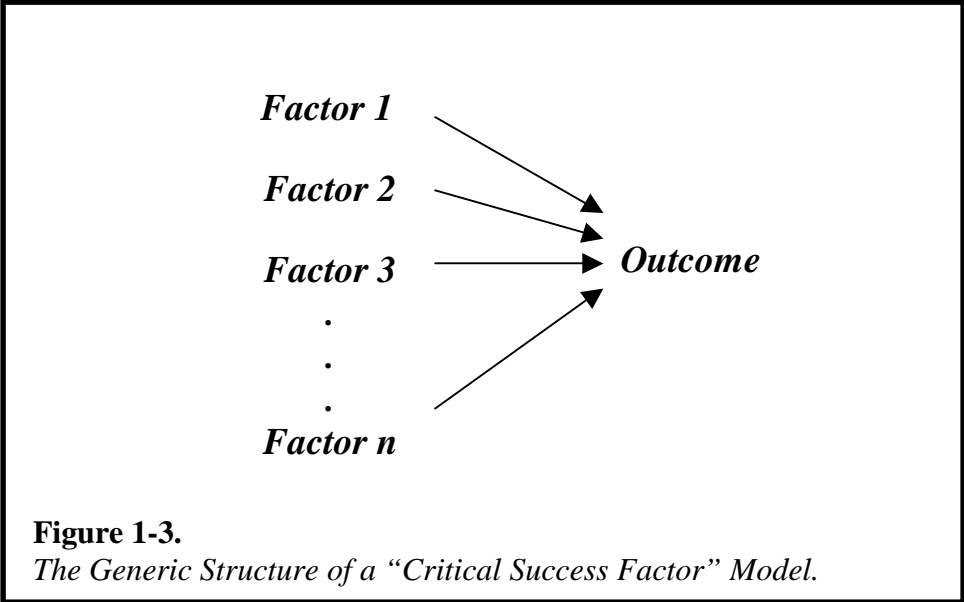


Figure 1-3.
The Generic Structure of a “Critical Success Factor” Model.

Okay, so what “meta assumptions” does this structure reveal? Two obvious ones suggest themselves. The first is that the “Factors” operate *independently*. Each “impacts” the outcome, but it does so, independently. The second is that the “Outcome” does not *cycle back*

to influence any of the Factors. That is, *causality is assumed to run one-way*—from Factor to Outcome, but not back again.

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

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

And so, isn’t there really a reciprocal, or closed-loop, causal relationship between top-quality people (or any of the other “factors”) and organizational success? Success is not just an *outcome*, something that is “driven” by a set of factors. Success is, itself, a driver! Causality runs *both* ways, not *one-way*! That’s “meta assumption” number two you hear landing with a thud!

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

One of these tools is the spreadsheet. Another is *The Balanced Scorecard* bubble diagram. A third is “root cause” or “fishbone” diagrams. In the artifacts created by each tool, like the CSF framework, we find “logic trees” with associated causality running

only one way. We also often find more independent than interdependent factors. But these popular tools also generally reflect two other “meta assumptions,” as well. The first of these is that impacts are felt *instantaneously* (i.e., delays are largely ignored). The second is that impacts are *linear* and *constant* (i.e., an x% change in input always results in a y% change in output).

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

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

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

*Reason 3 for
Poor Quality
Mental Models:
The Honing
Process*

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

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

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

What’s Next

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