Chapter 1

From Firefighting to Systematic Action

This book is about surprises that challenge evaluation. But surprise does not mean random. There is rhyme and reason to unanticipated occurrences. There is pattern. There is sense. The purpose of this book is to help evaluators appreciate the rhyme, understand the reason, see the pattern, and apply the sense. I hope to provide a theoretical understanding of surprise that affects evaluation, a social understanding of how unanticipated changes appear, and practical advice on how to act. My hope is to start a movement in our discipline toward collective thinking about how to incorporate consideration of surprise into our routine of planning and conducting evaluation.

ADDING "SURPRISE" TO THE MIX

When we think about how to conduct an evaluation we ponder topics such as evaluation models, evaluation theory, stakeholder relationships, information use, evaluators' roles, research design, data quality, funding, deadlines, and the logistics of evaluation implementation. We may place more or less emphasis on any of these. Sometimes our reliance on these intellectual tools is explicit and carefully planned. Sometimes the potential of these tools exerts an implicit pull on our consciousness and inclinations. But always, in one way or another, these concepts shape our view of how an evaluation should be done.

I believe that the science, technology, and craft of evaluation would be strengthened if another set of considerations were added to the mix, that is, a systematic understanding of how to contend with surprise,

with situations in which programs and evaluations do not behave as expected. Those surprises may be welcome or unwelcome, but either way, we need to be able to evaluate them.

Until recently, evaluators have not stepped back to look at surprise in a systematic manner, as a field of inquiry in its own right. This is beginning to change. The framework for this change is the rapidly growing influence of the "systems perspective" in evaluation, as evidenced by the enthusiastic acceptance of a Systems Topical Interest Group in the American Evaluation Association, discussions about simulation as an evaluation tool, and a growing awareness that evaluation must look not just at single programs in isolation, but at the context in which those programs are embedded. The systems view is important because much surprise comes from relationships, fluctuations, and uncertainties in how parts of a whole affect each other. The next step is to focus specifically on methods and procedures that evaluators can use to anticipate surprise, to ameliorate its impact on the integrity of evaluation designs, and to provide stakeholders with an enhanced understanding of the consequences of their actions.

HISTORICAL ROOTS:
EVALUATION, PLANNING, AND SYSTEM BEHAVIOR

The notion of unintended effects in evaluation has a long history. The notion of an unintended effect is critical in Scriven's advocacy of goal-free evaluation, an approach that focuses attention on what a program actually does rather than its stated goals (Scriven, 1991). Goal-free evaluation considers a wide range of change, but does not delve into why so many different program impacts occur, whether there are different categories of impacts, or what evaluators can do to improve their methods of detection and assessment. The theme of unintended effects is also echoed in the literature on how the act of evaluating a program may change the program that is being evaluated (Ginsberg, 1984; Glenwick, Stephens, & Maher, 1984). While the literature on evaluation–program reactivity does treat the question of how and why unintended effects occur, it is limited to a focus on the relationship between measurement and program action, rather than on the more general question of unintended consequences. The one article I could find that dealt with how evaluators might treat unintended consequences was titled "Identifying and Measuring Unintended Outcomes" (Sherrill, 1984). Had the ideas

in that article become mainstream in our field, I would probably not have felt a need to write this book. Sherrill touched on many of the ideas I elaborate on these pages—the difference between foreseeable and unforeseeable consequences, the value of theory and past experience, system-related reasons for surprise, and using diverse points of view to discern a program's consequences.

The notion of surprise is prominent in the planning literature of a wide variety of fields. Examples from a brief literature review include: industry sponsorship of university research (Behrens & Gray, 2001), marketing (Fry & Polonsky, 2004), tobacco restrictions (Hoek, 2004), drinking age regulation (DiNardo & Lemieux, 2001), speed and quality relationships in new product development (Lukas & Menon, 2004), welfare (Courtney, Needell, & Wulczyn, 2004), national fiscal reform (Kildegaard, 2001), teacher empowerment (Pugh & Zhao, 2003), nongovernmental organization (NGO) activity in developing countries (Stiles, 2002), and workplace safety (Kaminski, 2001).

Stepping outside discussions of unexpected change in particular fields, we find more general explanations that cut across specific domains of policy and planning. These explanations are typified by the works of Meyers (1981), Tenner (1996), and Dorner (1996). Their explanations focus on the principle that complex systems by their nature can yield unpredictable behavior because of the interplay of factors such as uncertain environments, cross-linkages, self-organized behavior, ecological adaptation, and feedback loops of different lengths. I do not want to leave the impression that systems exist in a perpetual state of flux. Under the right conditions they can be exceedingly stable over protracted periods of time. Stability, instability, and the boundary between them is one of the major concerns in the field of complex adaptive systems (CAS) (Kauffman, 1995; Marion, 1999). For our purposes, though, we must accept the fact that the potential for unpredictable behavior is inherent in the settings where we work.

In addition to the principles of complex systems, three behavioral/ organizational dynamics are at play. First, our decision making is always based on less information than we are able to collect. But to say that we can get "more relevant information" is to say that we know what information is relevant and how much information is enough. We can only know this in retrospect. Second, we are not as vigilant as we could be in scouting for developing changes in the system we are working with. But here, too, only a backward-looking view will tell us whether we were

vigilant enough and whether we were looking in the right places. Third, the nature of the planning process is such that opportunity for major intervention occurs only infrequently along a program's life cycle. It is only so often that budgets are set, requests for proposals are published, contracts are modified, staff are hired, or strategic plans made. Thus it is almost certain that knowing that action is needed will not synchronize with the ability to act.

In sum, our efforts to change systems take place during intermittent windows of opportunity, at which time we scramble to organize relationships among a finite (and often small) number of components (e.g., staff, money, time, client characteristics, material, procedure, information, and treatments). Because we are inescapably enmeshed in this process of systems change, we are destined to build models that are incomplete in their own right, and even more so when embedded in the larger system that we call the real world. To say that there will be no unexpected occurrences is akin to saying that a finite model can fully specify a complex system. That is impossible. It is akin to violating the First Law of Thermodynamics, and as we know, there cannot be a perpetual motion machine.

The difficulty with all these explanations for surprise is that while they help understand why unexpected events occur, they do not say anything about what evaluators should do about them. They do not inform our methodology. They do not help us resolve two competing design requirements—the requirement to plan as carefully as possible in advance, and the need to be flexible enough to provide useful information as circumstances change.

FROM EXPLAINING SURPRISE TO DEALING WITH IT

As evaluators our problem is not that the unexpected occurs. Our problem is that we react only when surprise falls upon us, when we need to put out a fire. At those times we exercise our craft. What we need is a framework for anticipation and for systematic action. We need to move beyond crisis management. For that, we need an answer to three questions.

1. When is the likelihood of surprise high?
2. Under what circumstances will surprise disrupt evaluation?

3. When the probability of surprise and disruption is high, what can we do about it?

A word of caution is in order. No matter how well we answer these questions we will always find ourselves in trouble. One problem we cannot escape is that any solution we conjure will entail overhead. Any commitment of resources will incur opportunity costs, and any tactic we deploy to deal with surprise will have its own drawbacks (much more on this topic to come in Chapter 7). Moreover, our evaluations are often themselves composed of many tightly linked elements. As we tinker with them, other surprises will occur. Our situation is like the problem faced by safety managers who must confront what Perrow calls a "normal accident" (Perrow, 1999). These kinds of accidents occur as a result of interactions among many tightly linked elements of complicated systems. Error in system design is not the problem. Rather, the very existence of those elements and their dependencies creates conditions where small changes in parts of the system can cause major disruptions. No matter what redesign is tried there will inevitably be many elements and many dependencies to consider, combining in unpredictable ways to bring about more mishaps.

We may not be able to escape surprise but we can appreciate how it works, and by so doing, we can develop strategies that will leave us better off. Unfortunately, no matter how much better off we become, there is no magic elixir that will turn the invisible visible, or that will make known all that cannot be predicted, or that will always allow us to react to change in a timely fashion. The best we can do is to increase the range of what we can see. We can give ourselves longer lead times to react. We can find ways to glimpse the hazy outlines of what was previously invisible.

So far I have used the term "surprise" as if it was simple and unitary; that is, either something happened that was expected, or something happened that was not expected. Either we are surprised, or we are not. In fact, surprise comes in different flavors. "Unexpected" can be divided into events that might have been foreseen had proper mechanisms been in place, and events that can never be foreseen. Different procedures are needed to buffer against each of these two types of surprise.

Respect for the difficulty of prognostication is needed when talking about surprise that might have been anticipated. Our problem is

that we are looking into a future in which multiple causal chains can lead to the same outcome, and in which chance events can change what seemed like a sure thing. In the discussion that follows I try to walk a fine line between claiming that sometimes eventualities might reasonably be anticipated in some way, shape, or form, and respecting the unknowability of how processes combine and interact.

DEVELOPMENT PATH OF THIS BOOK

This book began as an exercise in putting a foundation under a castle in the air. In 2005 I published a theoretical article on why unintended consequences exist, why some are truly unforeseeable while others might have been dimly foreseen, and what tactics evaluators might use to deal with each (Morell, 2005). While everything I said in that article made sense (to me and a few reviewers, at least), I was left with a discomfort over my lack of knowledge about what kind of unintended consequences evaluators actually face, and what they really do when faced with those circumstances. I set out on a quest to find real-world examples, to map them into my thinking about the evaluation of unintended consequences, and to use the synthesis to extend my understanding. From this effort came an expanded framework to categorize unintended consequences, and the rest of this book's content.

My data collection plan employed a 2-year-long snowball sampling methodology that worked through different avenues of outreach to the evaluation community. I reached out to members of the American Evaluation Association by distributing notices at their annual meetings and posting on their Listserv. I also used my connections as editor of *Evaluation and Program Planning* to query authors, reviewers, and advisors. I included my search request in any conversations or e-mail exchanges I had with friends and colleagues. Throughout, my message was always the same: Do you have cases? Do you have colleagues who have cases? Do you know people who may know others who may have cases?

The fruits of these labors were consistent. Almost every conversation I had with almost every evaluator elicited a response like: "What a great idea. This happens to me all the time. Do I ever have material for you!" But responses were few when it came to asking people to submit cases. I have several explanations for this outcome, but no data to support any of them. One possibility is that surprise threatens evaluation less often than I thought, despite people's verbal assurances that

the phenomenon was ubiquitous. Another possibility is the difficulty of writing on this topic without admitting failure as an evaluator, that is, copping to being an evaluator who executed an inadequate evaluation plan. The third possibility I can think of is that it is hard to discuss this topic without casting one's sponsors in a negative light, as a group of people who could not make up their minds, or who meddled in the evaluation, or who could not implement the program they had planned. Finally, there is the possibility that the discovery of unintended effects of program action requires a long-term view of program operations, and as we shall see, few of those kinds of evaluations were revealed with the sampling plan I used. Whatever the explanation, I am convinced that we will all be better evaluators if we amass as many cases as possible and use the collection to build a corpus of knowledge. My hope is that the small number of examples presented in this book will stimulate others to add to the collection.

Although I started with a plan to map real cases into my original typology, the mapping exercise forced me to extend my thinking about what "surprise" meant and how it should be handled. One major extension was the realization that "evaluations" and "programs" are similar social phenomena in that they are both collections of resources and processes that are embedded in a social/organizational setting for the purpose of accomplishing specific objectives. I came to understand that the same dynamics that generate surprise in programs also affect the evaluations themselves. (Chapter 8 is the start of a detailed explanation of these surprises.) A second realization was the extent to which efforts to buffer evaluation against surprise can become counterproductive as each buffering tactic incurs overhead costs in time, money, human capital, and management attention. (Chapter 7 delves into this topic in detail.) I realized how important it is to implement a great many methods to deal with surprise, but also how important it is to choose wisely.

As I reviewed cases and developed frameworks, I realized that I needed two different methods to use cases to illustrate the points I wanted to make. The first was to draw on the contributed cases. These were valuable because they highlighted the actual experience of evaluators working in real-world settings. To aid in drawing lessons from the cases and comparing them, I asked the contributors to present their cases in three main sections: (1) description of the case, (2) unexpected events, and (3) responses to the unexpected events.

The second method of using cases was needed because often no single case juxtaposed all the aspects of evaluation that I needed to make a point. For instance, in order to make a point in Chapter 5 I needed to illustrate the relationship between a program in a single department of an organization and similar programs that sprang up independently in a different part of the same organization. No such situation arose in the cases, so I constructed an example involving a safety training program for managers and its interaction with a companywide Lean Six Sigma quality improvement initiative.[1] I don't know of any evaluation like this, but I know of evaluations of safety programs and other evaluations of quality improvement programs. It is not much of a stretch to imagine the evaluation consequences of such programs operating simultaneously. The scenario has verisimilitude to real experience. Many of us would be able to think of that scenario and say, "I can imagine having to evaluate a situation like this." Or, "I have not done this kind of evaluation, but it makes sense to me. I have done things like it."

GUIDING PRINCIPLES

I am writing this book from a distinct point of view. It seems only fair to articulate that point of view to help you judge what is coming and its applicability to your own work. The arguments that follow emanate from a variety of perspectives that have guided my work for a long time: complex systems, innovation and organizational behavior, life cycle change, the dictates of practical action, and my role as an evaluator.

Complex Systems

As I have designed and implemented evaluations I have come to see my work as an exercise that takes place in a universe where control and authority are distributed, where the course of seemingly unambiguous action is uncertain, and where relationships are constantly changing. How many bosses have you met who can say: "I can tell people how to run this organization and they do as they are told"? Beyond immediate subordinates, almost nobody has this level of control. Organizations are just too big and too complicated for tight centralized control to be possible or desirable, particularly over extended periods of time across multiple subgroups. I believe that we all know this intuitively. The prin-

ciples of complex system behavior provide a theoretical underpinning to this intuitive understanding, as can be seen from some of the core principles of the adaptive system approach. (This list is adapted from Wikipedia's excellent overview of the essential elements of complex systems; *en.wikipedia.org/wiki/Complex_system*).

1. Complex systems show nonlinear relationships among their elements. Thus, small perturbation may cause a large effect, a proportional effect, or even no effect at all.
2. Complex systems contain multiple damping and amplifying feedback loops.
3. Complex systems have memory; that is, they change over time, and prior states may have an influence on present states.
4. Complex systems may be nested; that is, components of a complex system may themselves be complex systems.
5. Boundaries are difficult to determine.
6. Emergent phenomena may be present. (Think of a beehive. Its architecture, functioning, and output result from the behavior of many bees, but the hive cannot be understood by any summation of the actions of particular bees.)

Innovation in Organizational Settings

Evaluation is intimately bound up with innovation both because innovation is often the object of evaluation, *and* because the act of evaluation can itself be regarded as an innovation. Thus the behavior and characteristics of innovation are important. For instance, we know that characteristics of an innovation affect how it is implemented and its chances for success. We act differently depending on the innovation's degree of compatibility with existing values, complexity, trialability, and observability (Rogers, 1983). To illustrate the similarity between programs and their evaluations, consider a mental health program and a plan to evaluate that program with a randomized design. With respect to the program, we may observe that a behaviorist approach may not jibe with the values of service providers who are committed to psychodynamic therapies. With respect to evaluation, we may observe that a design based on randomization and control groups may not fit the values of either the service providers, or the evaluators, or both. In both scenarios, implementation of the innovation—whether the program or

its evaluation—suffers from the same problem, in this case, a values disconnect between various stakeholders.

For the most part, the innovations we evaluate are embedded in organizational settings. Thus we must appreciate organizational behavior in order to implement evaluations and also in order to apply evaluation in the service of understanding change. I tend to see organizations as complicated entities that pursue multiple (and often conflicting) goals, which must allocate scarce resources, are beholden to multiple stakeholders, have formal and informal structures, shifting stores of intellectual capital, changing environments, tensions between goals of service and goals of self-preservation, and interorganizational relationships. Into this setting are cast innovations and their evaluations, which are themselves partially nested and partially overlapping.

Life Cycle Change

Life cycle changes are driven by developmental processes that produce highly predictable consequences. Life sciences are one obvious field where life cycle change is an important concept, but the notion can be found in many other fields as well. A few examples are: organizational development (Sherman & Olsen, 1996), research and development (R&D) project management (Pillai & Joshi, 2001), innovation management in companies (Koberg, Uhlenbruck, & Sarason, 1996), organizational growth and financial performance (Flamholtz & Hua, 2002), strategy formation (Gupta & Chin, 1993), alliance behavior (Lambe & Spekman, 1997), software development (Boehm & Egyed, 1999), and innovation adoption (Adner, 2004).

Two aspects of a life cycle are noteworthy. First, stages are predictable and invariant. Once we know what stage an entity or process is in, we can predict what the next stage will be. (This is not *always* so, but when stages are not followed the results are rare and noteworthy.) Second, the way in which an entity is managed, and how it interacts with its environment, is affected by life cycle stage. For instance, programs in a start-up phase are likely to undergo fast-paced change that calls for rapid feedback evaluation methodologies that are not appropriate for assessing outcome once the program is stable. In this book I take the perspective that both innovations and evaluations go through life cycles, and that interactions between evaluation and innovation depend on where each is in the life cycle.

Practical Action

I believe that, above all, evaluation must be practical. It must inform action in real-world settings. I do not mean that evaluation use must only be instrumental, but I do believe that evaluation must guide practical action. In this sense evaluation is a technological endeavor (Morell, 1979, Ch. 5). As Jarvie (1972) puts it: "The aim of technology is to be effective rather than true, and this makes it very different from science." Evaluation should focus on program theory and design elements that will be powerful enough to make a difference in real-world settings, and what can either be manipulated by a program or taken into consideration by a program in order to take action that will have practical consequence. I believe that evaluation can be practical because while change is ever present and the future is always unknowable, time horizons can be set at which a reasonable degree of certainty exists and within which evaluation can provide a reliable guide for further action. John Maynard Keynes recognized the importance of setting appropriate time frames in a famous and eloquent statement. What was true for economics in 1923 is true for evaluation today.

> The long run is a misleading guide to current affairs. In the long run we are all dead. Economists set themselves too easy, too useless a task if in tempestuous seasons they can only tell us that when the storm is past the ocean is flat again. (*A Tract on Monetary Reform*, retrieved from *www.quotationspage.com/quote/38202.html*.)

My Role as an Evaluator

My view is that programs represent investments in a course of action that are designed to achieve specific objectives. Investors have a right (and an obligation) to ask whether their investment has paid off. Evaluators have the obligation to answer that question. I believe that the best approach to getting the needed information is almost always to begin with careful identification of measures and rigorous advance planning. An important source of evaluation surprise is the inescapable fact that evaluation must conform to a basic assumption that is made by program planners: that their programs as designed will make a difference. We may be able to influence planners' thinking about what those differences may be, and we may have some wiggle room to include outcomes that were not envisioned by the planners. We must do so whenever we can, and we should practice the art and craft of constructing as much

of that wiggle room as possible. But in the main, we are contractually bound (if not duty) bound to measure program behavior relative to the goals of the program's designers and funders. That said, programs do change in their structure, function, intent, and impact, and evaluators do need to stay abreast of those changes.

HOW TO READ THIS BOOK

Working through the Chapters

This book consists of 13 chapters that I have written and 18 cases contributed by people working across a wide array of methodological proclivities, programs, and geographical areas. Case titles were chosen to reflect these attributes. Table 1.1 contains the case numbers and case titles. Each time I refer to a case I provide the case number and the case's beginning page.

Chapters 2 though 7 are heavy on theory and draw on the cases to illustrate various points I tried to make. Chapters 8, 9, 10, and 11 focus on the cases and show how the previously discussed theory can be applied. Chapter 12 draws from all the previous chapters in an effort to explain how evaluators can handle unexpected program outcomes. Chapter 13 contains my concluding remarks.

Chapter 2 categorizes different kinds of surprise. It can be read in stand-alone fashion. Chapter 3 addresses two questions: (1) When is the probability of surprise high? and (2) When is surprise disruptive to evaluation? It, too, can be read as a stand-alone. Chapters 4, 5, and 6 should be read together because they progress through different degrees of surprise, all the while suggesting tactics that may help evaluators conduct their business. Chapter 7 is a cautionary tale about the perils of heeding too much of my advice. It argues that any tactic for dealing with surprise brings its own risks, and that the more that is done, the greater the risk. Although Chapter 7 draws on the previous chapters, knowing what went before is not a requirement for raising one's consciousness about the risk of making research design too complicated. Chapters 8, 9, 10, 11, and 12 discuss cases, but draw heavily from the earlier material. These chapters have tables and figures to illustrate various points or to summarize information. Just reading the titles of those tables and figures may help convey a sense of what this book is about. Table 1.2 provides numbers and titles for all the tables and figures that are to come.

TABLE 1.1. Overview of Cases

Case	Page	Title
1	197	Grasping at Straws and Discovering a Different Program Theory: An Exercise in Reengineering Analysis Logic in a Child Care Evaluation Setting
2	200	Shifting Sands in a Training Evaluation Context
3	204	Evaluating Programs Aimed at Promoting Child Well-Being: The Case of Local Social Welfare Agencies in Jerusalem
4	210	Assessing the Impact of Providing Laptop Computers to Students
5	214	Quasi-Experimental Strategies When Randomization Fails: Propensity Score Matching and Sensitivity Analysis in Whole-School Reform
6	219	Unexpected Changes in Program Delivery: The Perils of Overlooking Process Data When Evaluating HIV Prevention
7	224	Evaluating Costs and Benefits of Consumer-Operated Services: Unexpected Resistance, Unanticipated Insights, and Déjà Vu All Over Again
8	231	Keep Up with the Program!: Adapting the Evaluation Focus to Align with a College Access Program's Changing Goals
9	235	Assumptions about School Staff's Competencies and Likely Program Impacts
10	241	Mixed Method Evaluation of a Support Project for Nonprofit Organizations
11	244	Evaluating the Health Impacts of Central Heating
12	249	Recruiting Target Audience: When All Else Fails, Use the Indirect Approach for Evaluating Substance Abuse Prevention
13	253	Unintended Consequences of Changing Funder Requirements Midproject on Outcome Evaluation Design and Results in HIV Outreach Services
14	258	Generating and Using Evaluation Feedback for Providing Countywide Family Support Services
15	263	Trauma and Posttraumatic Stress Disorder among Female Clients in Methadone Maintenance Treatment in Israel: From Simple Assessment to Complex Intervention
16	270	From Unintended to Undesirable Effects of Health Intervention: The Case of User Fees Abolition in Niger, West Africa
17	277	Unintended Consequences and Adapting Evaluation: Katrina Aid Today National Case Management Consortium
18	281	Evaluation of the Integrated Services Pilot Program from Western Australia

Getting Acquainted with the Content of the Cases

The 18 cases that evaluators so graciously contributed to this book contain a great deal of information about evaluations that used many different methodologies to support a wide variety of stakeholder needs in a variety of substantive domains and geographical areas. As I draw on specific cases I will provide snippets of contextually appropriate material about that case. As you read, you may find in these snippets all you need to maintain a sense of how the case illustrates the point I am trying to make. Or, you may find it more comfortable to use either of two other approaches. The first is to take the time to read (or skim) through the whole case when I refer to it. Second, you may want to read (or skim) through all the cases now, before descending into the detail that is to follow. Any of these approaches, in various combinations, may work for you. Do what feels best.

IN SUM

In this chapter I tried to make the case that we would be better off if we moved our efforts to deal with surprise from the realm of crisis management to the realm of systematic inquiry. I summarized what the fields of planning and systems tell us about where surprise comes from and how it behaves. I tried to show that what is known about surprise from those disciplines is a useful foundation for us, but that knowledge from those other disciplines is not sufficient because it does not touch on evaluation's particular need to collect and analyze empirical data. Finally, I laid out the intellectual foundations that have forged my sense of how a systematic understanding of surprise can be integrated with evaluation methods. These are the intellectual threads that I spin in the rest of this book.

NOTE

1. By "Lean Six Sigma," I refer to the collection of process improvement methodologies and tools that populate the field of continuous process improvement.

TABLE 1.2. Index to Tables and Figures

[a]Also applies to Chapter 11.