© *Academy of Management Review*
2009, Vol. 34, No. 4, 733–756.

THE DYNAMICS OF ACTION-ORIENTED PROBLEM SOLVING: LINKING INTERPRETATION AND CHOICE

JENNY W. RUDOLPH
Harvard Medical School and Massachusetts General Hospital

J. BRADLEY MORRISON
Brandeis University

JOHN S. CARROLL
Massachusetts Institute of Technology

We offer a theory of action-oriented problem solving that links interpretation and choice, processes usually separated in the sensemaking literature and decision-making literature. Through an iterative, simulation-based process we developed a formal model. Three insights emerged: (1) action-oriented problem solving includes acting, interpreting, and cultivating diagnoses; (2) feedback among these processes opens and closes windows of adaptive problem solving; and (3) reinforcing feedback and confirmation bias, usually considered dysfunctional, are helpful for adaptive problem solving.

Understanding and improving action-oriented problem solving under time pressure is a crucial aspect of organizing for high performance (Eisenhardt, 1989; Rudolph & Repenning, 2002; Starbuck, Greve, & Hedberg, 1978; Sutcliffe & Weber, 2003; Weick, Sutcliffe, & Obstfeld, 1999). Professionals in many settings must formulate interpretations and make choices quickly with imperfect understanding as problems and opportunities evolve rapidly and as their actions both generate information and change the situation (Hodgkinson & Healey, 2008). To take effective action, executives in "high-velocity" or "high-hazard" environments, entrepreneurs seeking first-mover advantage, team leaders in military and clinical settings, firefighters confronting smoke and flames, and professionals troubleshooting manufacturing or computer programming challenges are each wondering, "What is going on here?" (Arthur, 1990; Carroll, Rudolph, & Hatakenaka, 2002; Eisenhardt, 1989; Klein, Pliske, Crandall, & Woods, 2005; Repenning & Sterman, 2002; Weick, 1993). As they pursue solutions, their diagnoses and actions coevolve; feedback links together their sensemaking and decision-making processes.

In this paper we clarify and articulate a model of action-oriented problem solving that integrates processes of interpretation and choice. We developed a system dynamics simulation model to represent theoretical concepts and relationships and to test emergent properties of our theory. The model was motivated by the empirical example of doctors coping with an operating room emergency. Patterns of problem-solving behavior produced by the model provide implications for both theory and practice.

INTERPRETATION AND CHOICE IN ACTION-ORIENTED PROBLEM SOLVING

Mapping and understanding the processes that link meaning making and choice in action-oriented problem solving have been difficult because the sensemaking literature and decision-making literature have partitioned problem solving by focusing on different prototypical situations. The sensemaking literature assumes

We gratefully acknowledge the help of Marlys Christianson, Jason Davis, Mark Healey, Zur Shapira, Karl Weick, the *AMR* reviewers and editors, and especially Olav Sorenson and Ron Adner. Thanks also to Fulton 214 (Erica Foldy, Pacey Foster, Danna Greenberg, Tammy MacLean, Peter Rivard, and Steve Taylor) for multiple trenchant reviews of this paper.

the environment is challenging if not hostile to problem solving: information abounds but meaning is ambiguous; the best information may emerge only after taking action. Indeed, the purpose of interpretation is not so much to be right (often a post hoc rationalization) as to guide action toward a more effective response. Some of the most compelling sensemaking studies analyze problem solvers in extreme and harrowing situations; they could be killed when sensemaking "collapses" (Weick, 1993). In contrast, classical decision-making research assumes that meaning is given or readily calculated, as in a laboratory gamble. Although decision makers with limited rationality take error-prone shortcuts (Simon, 1957; Tversky & Kahneman, 1974), a decision is properly approached as a cool, rational, consistent, and comprehensive search for the right answer (Fischhoff & Beyth-Marom, 1983; Langley, Mintzberg, Pitcher, Posada, & Saint-Macary, 1995).

While in most important real-world settings there is an overlap of processes of evaluation and choice associated with classic decision-making models and processes of interpretation and construction of meaning associated with sensemaking, the two bodies of literature have tended to make advances by addressing these problems separately. The lack of cross-fertilization between these two fields means that the interactions among processes of interpretation and choice are relatively uncharted. Whether problem solving is primarily a choice-based or rule-based process, whether it is primarily an instrumental and positivist or interpretive activity, or whether inputs and outputs to problem solving are characterized by clarity or ambiguity are all contested issues (March, 1994: viii–ix). We know little about how pacing, punctuation, and relative effort given to interpretation and meaning making versus evaluation and choice influence problem solving. Echoing March, we agree that "the largest problem is not to choose among the [theoretical] alternatives but to weave them together in a way that allows each to illuminate the others" (March, 1997: 10).

Insights about how interpretation and choice are woven together in action-oriented problem solving emerged from our attempt to understand variations in problem solving in a clinical crisis. We developed a formal model that includes feedback loops that the separate sensemaking and decision-making literature rely on implicitly but have not fully articulated. The model includes three basic processes—acting, interpreting, and cultivating new diagnoses—derived from the example of diagnostic problem solving in an operating room crisis and from theoretical constructs in both bodies of literature. Interactions among the three processes produce adaptive problem solving and various problem-solving failures. The speed and strength of acting, interpreting, and cultivating new diagnoses facilitate and constrain opportunities for choice between competing diagnoses and open and close windows of opportunity for adaptive problem solving. For example, either acting slowly when one is unsure or generating alternatives too quickly can make it difficult to collect enough supporting information to confirm a diagnosis and resolve the crisis. The model reveals leverage points for converting dysfunctional problem-solving modes into adaptive ones, including the counterintuitive result that reinforcing feedback and confirmation bias can be beneficial.

Although the initial focus of the model is on medical diagnoses in a crisis setting, the model is intended to apply more generally to problem solving in a broad range of situations characterized by (1) action-based inquiry—acting is the only way to generate new information to update explanations and action strategies (Sutcliffe & Weber, 2003; Torbert & Taylor, 2008); (2) temporal dynamism—the situation changes on its own if no action is taken (Rudolph & Repenning, 2002; Snook, 2000; Weick, 1993), and, indeed, failure to act promptly will result in adverse consequences; and (3) action endogeneity—actions change the problem-solving environment (Perlow, Okhuysen, & Repenning, 2002; Sutcliffe & Weber, 2003). Settings that exhibit these characteristics include high-velocity entrepreneurial environments where useful information becomes available over time as a strategy is executed, business competitors are moving on their own, and the market changes partly in response to an individual's own actions (Eisenhardt, 1988; Perlow et al., 2002; Sutcliffe & Weber, 2003). They also include firefighting emergencies in which firefighters may have to gather information about the fire by placing themselves in harm's way or trying various countermeasures (Klein, 1998). These examples involve knowledgeable professionals facing problems that are neither completely novel nor completely routine, in which they must act and diagnose iteratively in a dynamic, unfolding situation.

We begin our theory development process in the next section by presenting findings from a motivating example: an in-depth observational study of action-oriented problem solving by thirty-nine anesthesiology residents in a high-fidelity simulated[1] operating room facing an acute care crisis with the patient's life in the balance (Rudolph, 2003; Rudolph & Raemer, 2004). We then describe our iterative method of theory development and review and synthesize the relevant research literature in the exposition of our model. This exposition explains the key theoretical constructs of acting, interpreting, and cultivating new diagnoses and the dynamic interactions among these constructs. Next, we discuss simulation results that provide insights on the origins of adaptive and dysfunctional problem solving. We end with a discussion of the simulation results, highlighting the mechanisms that produce variation in action-oriented problem solving and their implications for problem-solving theory and practice.

ACTION-ORIENTED PROBLEM SOLVING: AN EXAMPLE

The starting point for our theorizing was Rudolph's study of action-oriented problem solving by thirty-nine advanced anesthesia residents facing the same simulated acute care crisis with a diagnostic challenge (Rudolph, 2003; see also Rudolph & Raemer, 2004). Rudolph observed problem solving by tracking doctors' concurrent verbal statements regarding diagnoses, observing their treatments and diagnostic tests, and conducting post hoc video reviews with participants. An operating room crisis is a particularly good context for elaborating the features of action-oriented problem solving because it requires handling both novel and routine demands. The novel demands require skillful exploration and interpretation of ambiguous cues to diagnose a problem, giving researchers a chance to examine cognitive processes. The routine demands require rapid and efficient action using exploitation of known procedures (Rudolph & Repenning, 2002), providing an opportunity to study action.

In Rudolph's study an anesthesia resident is called to take over anesthesia during an urgent appendectomy for a 29-year-old woman. This scenario presents a common but serious problem in anesthesia: difficulty with the process of ventilating—that is, breathing for the patient using a mechanical bellows. A variety of diagnoses for the ventilation problem are plausible, such as an asthma attack, a collapsed lung, or insufficient paralyzing agent, but contradictory evidence is present for each, except one: the patient has exhaled some mucous into the breathing tube, partially blocking it. This is not uncommon, and trainees have been acquainted with this problem in their training, but presentations of the problem can vary. Treatments addressing diagnoses other than the mucous plug in the breathing tube will not result in any sustained improvement in the patient's status. With a slowly dwindling level of oxygen in her blood, the patient can have an uneven heartbeat and even go into cardiac arrest if the problem is not rectified. The cues doctors have available include vital signs displayed on a monitor, the breath sounds of the patient, and new cues generated by pursuing standard operating procedures for treating and diagnosing the patient.

Rudolph found that the doctors fell into four modes of problem solving as they attempted to address the ventilation problem: stalled, fixated, vagabonding, and adaptive (see Table 1). The doctors who were *stalled* problem solvers had difficulty generating any diagnoses around which to organize action (on average, just 1.5 diagnoses) and pursued few or no treatments and tests (one treatment on average). In contrast, those in the *fixated* mode quickly established a plausible but erroneous diagnosis, to which they stuck despite countervailing cues (Table 1 shows that they returned to their favorite diagnosis on average ten times—double any other mode). Rather than advancing through multiple steps of a treatment algorithm to rule out diagnoses, they repeated the same step or got stuck.

[1] We use the term *simulation* in two ways in this paper. The first use refers to the source data for Rudolph's (2003) study of clinical problem solving. These data were provided by a full-field, high-fidelity simulation—that is, the research participant was in a fully equipped and staffed operating room [OR] with a computer-controlled mannequin patient. This is similar to war games or aircraft flight simulators. The second use of the term refers to the computer-based simulation we conducted to analyze the behavior of our formal mathematical model.

TABLE 1
Summary of Source Data[a]

	Problem-Solving Mode				
Variable	Stalled	Fixated	Vagabonding	Adaptive	Test of Difference
N	2	11	17	9	—
Subjects who resolved the airway problem	0	0	0	7	χ^2 (3) = 28.4***
Different treatment steps for a diagnosis	1.0 (0.0)	2.0 (1.1)	1.5 (0.5)	3.6 (0.7)	$F(3, 35) = 17.0$***
Considerations of favorite diagnosis	3.0 (0.0)	10.0 (5.7)	5.4 (2.3)	5.9 (2.2)	$F(3, 35) = 5.0$**
Number of different diagnoses considered	1.5 (0.7)	3.8 (1.7)	6.1 (1.3)	5.0 (1.4)	$F(3, 35) = 9.1$***

[a] Adapted from Rudolph (2003).
Note: Means are given with standard deviation in parentheses. ** $p < .01$. *** $p < .001$.

Although previous studies of fixation error (also known as premature closure or tunnel vision) generally concluded that broadening the range of alternatives is the needed antidote to fixation (Gaba, 1989; Johnson, Hassenbrock, Duran, & Moller, 1982), Rudolph found that broadening could also be a problem; her data indicated a third mode, diagnostic *vagabonding*.[2] "Diagnostic vagabonds" generated a wide range of plausible diagnoses and jumped from one to another without utilizing multiple action steps (such as giving a medication) of the treatment algorithms for addressing and ruling out these diagnoses (they pursued just 1.5 steps on average).

Finally, the *adaptive* sensemaking mode, which looks very much like canonical models of effective clinical reasoning (Elstein, Shulman, & Sprafka, 1978), was characterized by generation of one or more plausible diagnoses and exploitation of multiple steps of known treatment algorithms. This process allowed those in the adaptive mode to rule out some diagnoses, take effective action, and, unlike those in any other problem-solving mode, resolve the breathing problem.

This context has several features that are particularly promising for launching a modeling effort: there is a diagnostic challenge that requires sensemaking, but a limited set of plausible diagnoses (about seven); diagnoses, once generated, can only be assessed through active treatment and testing; there is only one correct diagnosis; the actions necessary for testing and treatment are relatively clear; and the consequences for misdiagnosis are dire. Further, although Rudolph's study provides rich behavioral data on a variety of problem-solving modes, there are very limited data on the dynamic processes underlying these modes and only a single scenario from which to generalize. Hence, the motivating data are a "theory fragment" ripe for elaboration with formal modeling (Davis, Eisenhardt, & Bingham, 2007).

METHODS: SIMULATION FOR THEORY DEVELOPMENT

We explored the mechanisms that produce variation in diagnostic problem solving by developing a formal model. Rather than deducing a formal model from general principles, we used a theory development process now well articulated in the system dynamics community to induce a model from empirical findings in one or several studies (Black, Carlile, & Repenning, 2004; Perlow et al., 2002; Repenning & Sterman, 2002; Rudolph & Repenning, 2002; Sastry, 1997). While commonly used to build theory from raw data using qualitative analysis, the grounded theory approach we employed is also useful to develop theory from theory (Glaser & Strauss, 1967; Strauss & Corbin, 1994). There is a growing appreciation in the organization studies community that inducing formal models from empirical data and stylized facts facilitates the identification of structures common to the different narratives and enforces the internal consistency of the emerging theory (Davis et al., 2007).

With the motivating example of Rudolph's (Rudolph, 2003; see also Rudolph & Raemer,

[2] Dörner (1997) identified a similar phenomenon among public officials attempting to identify effective strategies for public policy.

2004) taxonomy of four diagnostic problem-solving modes in a medical emergency as a starting point, we pursued the following five steps to carry out our theory development. First, using the steps of grounded theory building to build theory from theory (Strauss & Corbin, 1994; Suddaby, 2006), we started by translating the constructs and relationships described in Rudolph's narratives of each problem-solving mode into the system dynamics language of stocks, flows, and feedback loops (Forrester, 1961; Sterman, 2000). We also gathered fragments of other relevant theory, proposing constructs and relationships among them and sketching bits of feedback structure (Davis et al., 2007). For example, central processes in the problem-solving literature are described as involving adaptive or reinforcing processes, but the exact feedback relationships that drive these cycles are not clear. The sensemaking literature describes the interplay of acting and interpreting as recursive and adaptive (Weick, Sutcliffe, & Obstfeld, 2005: 409), and fixation is described as self-reinforcing (De Keyser & Woods, 1990), but neither has been explicitly mapped.

Second, using the constructs drawn from Rudolph's study and from the sensemaking literature and decision-making literature, we experimented with different causal loop diagrams and compared our diagrams with constructs and relationships from this literature. It was at this point that we recognized that problem solvers had to make a choice between competing diagnoses. Third, we revised and streamlined our causal loop diagram, converging on just three central processes. Fourth, we translated the links and loops of this diagram into a formal model. For example, each stock is represented by an equation that captures the rate of its inflows and outflows and each connection, such as one between the perceived plausibility of a diagnosis and openness to new cues, as a mathematical function (see Appendixes A and B). Theorizing about dynamic processes with the benefit of a formal model allowed us to identify important logical gaps and inconsistencies in existing theory (Sastry, 1997; Sterman, 1994). The model's robustness and its ease of interpretation were enhanced by using standard system dynamics formulations (familiar fragments of model structure that occur frequently) wherever possible (Sterman, 2000).

Finally, we ran computer simulations of the model, varying the parameters to explore its behavior over a range of conditions. These scenarios or experiments allowed us to examine whether the parts of the model were operating in a psychologically reasonable way and whether the model was reproducing the empirical results. The model and its output provide a theory explicitly mapping the role of both balancing and reinforcing processes in effective and ineffective problem solving. While grounded in previous work, the model also provides new insights.

CONCEPTUAL MODEL

Overview

In Rudolph's findings we observed that doctors need to make sense of a rich and sometimes confusing stream of information *and* they need to choose among the diagnoses they generate in order to make sense of this information. Iterating between Rudolph's problem-solving modes and theories of decision making and sensemaking, our modeling process converged on the idea that action-oriented problem solving involves three central tasks linked by both reinforcing and adaptive (or balancing) feedback: (1) problem solvers take actions and, thus, make information available for their interpretation; (2) they interpret the flow of information around them to continually reassess the plausibility of their diagnoses; and (3) they cultivate alternative diagnoses even as they pursue the leading diagnoses.

The trigger for action-oriented problem solving, whether for doctors or executives, is usually a surprise or divergence from what the problem solver expects or desires, such as a disappointing corporate initiative or an anomalous building fire (Klein et al., 2005; Louis & Sutton, 1991; Mandler, 1984; Weick et al., 2005). In our motivating example the anesthesiology resident expects the patient to breathe normally but instead observes and seeks to address a serious problem with the patient's ventilation. To succeed, any problem solver (e.g., an executive, a fire commander) must construct an organizing story about what is wrong (Weick et al., 2005). In our clinical example the resident scans the clinical signs and symptoms, the patient history, and the timing of the problem, and then developes a plausible story; this then takes the form of a diagnosis (Elstein, 2001; Elstein et al., 1978; Klein, Phillips, Rall, & Peluso, 2006; Rudolph & Raemer, 2004).

We pick up the story from here, presenting the key constructs in our model of action-oriented problem solving—acting, interpreting, and cultivating new diagnoses—and show how they are linked. While these three processes unfold and interact recursively, we start our model description with the acting process because it is easiest to grasp.

Action Generates New Information

Catalyzed by a deviation from expectations and guided by an initial organizing diagnosis, the problem solver launches into action. In the domain of time-pressured acute medical care, conducting tests and providing treatments often involve following a standardized algorithm—a set of steps combining therapeutic treatment with diagnostic tests (Cook & Woods, 1994; Elstein et al., 1978). This is a rule-based behavior that requires executing a known procedure (Rasmussen, Pejtersen, & Goodstein, 1994). Many other professional domains, from nuclear power plant operations to computer chip manufacturing quality control, have standard diagnostic or operating procedures to address problems (Carroll, Rudolph, Hatakenaka, Wiederhold, & Boldrini, 2001; Cyert & March, 1963; Edmondson, 2002; Gersick & Hackman, 1990; Levitt & March, 1988; Repenning & Sterman, 2002; Winter, 1971). By moving through the steps of a standard operating procedure, problem solvers generate cues that become available for them to notice, bracket off from the stream of cues, and interpret. Having advanced the steps further, the problem solver has access to a larger pool of cues for making meaning in an ambiguous situation.

Since the consequences of action accumulate, we model this progress as a stock variable, "action steps completed," in Figure 1. Stocks, like water in a bathtub, are the accumulations of flows into and out of them and are represented by rectangles (Forrester, 1961; Sterman, 2000). Stocks have a key role in creating dynamics: they create delays, give systems inertia, and provide systems with memory. Accomplishing action steps takes time (delays), sets the problem solver on a particular course of action (inertia), and yields results that remain available for interpretation (memory). The stock is increased by "taking action," a flow variable, represented by a pipe and valve icon. The rate of taking action is determined by "time needed to take steps," which represents the time needed for mentally organizing to execute a step, physically preparing for the step, executing the step, awaiting a response from the system, and noticing the results as cues in the stream of ongoing experience. The rate of taking action also depends on how many of the action steps remain to be executed. Causal arrows are used to depict dependencies, such as the arrow in Figure 1 from "action steps completed" to "taking action."

FIGURE 1
Action Steps Make Cues Available

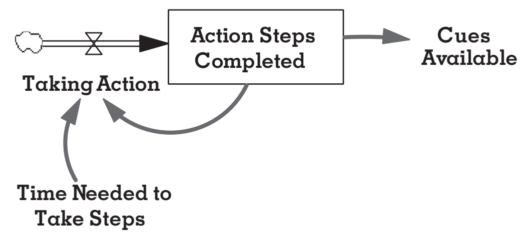

Note: A stock is a reservoir or accumulation (like water in a bathtub) and is represented by a rectangle; flows, like the spigot and drain on a bathtub, fill or drain the stock and are depicted as "pipes" with "valves." The "action steps completed" (represented by the rectangular stock) increases (or stops increasing) as people "take action" (or stop taking action) to complete the standard operating procedure or other algorithm (represented by the inflow pipe). The "time needed to take steps" reflects delays inherent in preparing to take action, taking action, and waiting for a response. "Cues available" is a variable that gives the problem solver confirming or disconfirming data about his or her current leading diagnosis.

Another feature of action is that taking action steps makes "cues available" (see Figure 1) for interpretation. For example, as a doctor accomplishes steps in a clinical treatment algorithm, he or she generates new diagnostic information that can then be considered in the context of his or her diagnosis. Analogously, as an organization carries out a new strategic initiative, executives can observe changes in the competitive environment that become inputs to their assessment of the situation and the merits of their strategies.

Action-Oriented Problem Solving Requires Interpreting Cues

Facing a complex and ambiguous situation where quick action is needed, a problem solver has to create meaning in order to act. Unlike decision-making experiments in a laboratory, where "meaning already exists and is waiting to be found," in complex and ambiguous settings, meaning "awaits construction that might not happen or might go awry" (Weick, 1995: 15). Generating a plausible story or explanation about ambiguous cues helps organize and launch action (Neisser, 1976; Snook, 2000; Weick et al., 2005).

Studies of strategic action (Sutcliffe & Weber, 2003), enactment of organizational structures (Weick, 1995; Weick et al., 2005), naturalistic medical problem solving (Elstein, 2001; Elstein et al., 1978; Johnson et al., 1982), tactical decision making under real-world stress (Cannon-Bowers & Salas, 1998), problem detection (Klein et al., 2005; Mandler, 1982), and problem solving in other naturalistic environments (Carroll et al., 2001; Klein, Orasanu, Calderwood, & Zsambok, 1993; Zsambok & Klein, 1997) all indicate that a plausible explanation, diagnosis, or "story" is the engine of problem solving. Problem solvers such as corporate executives, chess players, or firefighters, for example, use an initial diagnosis or assessment of the situation to develop a plan of action and decide what further information is needed (Dreyfus, 1997; Klein, 1998; Sutcliffe & Weber, 2003). Once an initial diagnosis is established, how plausible people consider their own explanations or diagnoses waxes or wanes as they make sense of new or existing cues (Koehler, 1991; Smith & Blankenship, 1991). It takes time for people to notice, bracket as important, and label new cues and then to change their mental models (e.g., diagnoses) accordingly (Bartunek, 1984; Kleinmuntz, 1985; Marcus & Nichols, 1999; Roberts, 1990).

To capture this process of perceived plausibility increasing or decreasing with the interpretation of cues, we show "plausibility of leading diagnosis" as a stock variable in Figure 2. This stock depicts the problem solver's current subjective assessment of plausibility. By "leading" we mean the problem solver's favorite or current most plausible explanation. "Updating" is the process or flow variable by which the current view of the plausibility of the leading diagnosis

FIGURE 2
Updating Subjective Assessments of Plausibility

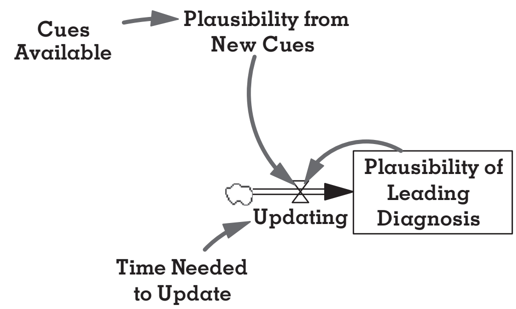

Note: "Plausibility of leading diagnosis" (represented by the rectangular stock) increases or decreases as people incorporate interpretations of "plausibility from new cues" through "updating" (represented by the inflow pipe). "Time needed to update" reflects delays inherent in noticing and processing information.

(the stock) is adjusted to equal "plausibility from new cues." Updating is an ongoing process of incorporating interpretations based on new information with current beliefs. Since changes in perceived plausibility require time, the rate of updating is influenced by a parameter we call "time needed to update." "Plausibility from new cues" describes interpretations of information generated by acting, which, in turn, depends in part on cues available.

Making meaning from a stream of cues available through noticing, bracketing, filtering, and labeling occurs while the problem solver holds a belief in the leading diagnosis, which may influence these interpretive processes. New assessments of plausibility ("plausibility from new cues") are dependent not only on cues available but also on how open the problem solver is to these external cues, which we call "weight on cues" and introduce in Figure 3. Studies of confirmation bias show that once an explanation is set, people prefer supporting to disconfirming information, and this effect is stronger still when cues are presented serially, as in our model (Jonas, Schulz-Hardt, Frey, & Thelen, 2001). Studies of fixation show that as the plausibility of the current diagnosis rises, openness to external cues, especially ones that defy the current view, decreases (De Keyser & Woods, 1990; Johnson, Moen, & Thompson, 1988; Staw, 1976; Xiao & MacKenzie, 1995). In other words, "weight on cues" is a downward-sloping function of "plausibility of leading diagnosis."

FIGURE 3
Feedback Structure of Self-Fulfilling Interpretation

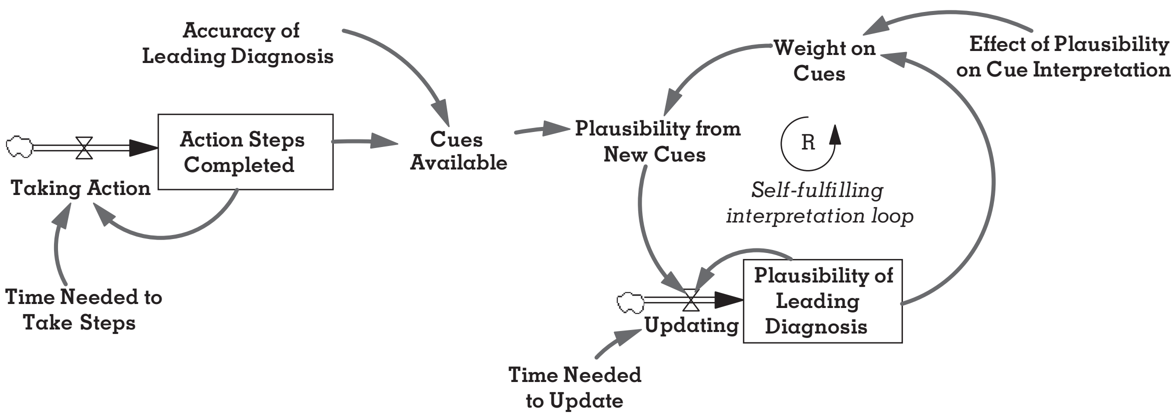

Note: As "plausibility of leading diagnosis" increases, the "weight on cues" decreases, so, all else being equal, the "plausibility from new cues" continues to grow, leading to "updating" that increases "plausibility of leading diagnosis" still further. The process forms a reinforcing feedback loop, "R," the "self-fulfilling interpretation loop."

However, prior research is surprisingly silent regarding the exact form of the relationship between plausibility and weight given to external cues. We postulate that for bold problem solvers a small increase in plausibility will lead to a disproportionately large decrease in weight on cues. For cautious problem solvers a small increase in plausibility will lead either to no or to only a small decrease in the weight on external cues. We use a parameter, "effect of plausibility on cue interpretation," to model the variation, from boldness to caution, in how plausibility influences the weight placed on cues. Appendix A depicts our representations of this relationship for different values of the effect of plausibility on cue interpretation.

To facilitate our mathematical description of the stream of action-generated cues, we assume the information made available for interpretation is either confirming or disconfirming, depending on which leading diagnosis is under consideration. We define the variable "accuracy of leading diagnosis" such that when the problem solver's leading diagnosis is correct, action-generated cues confirm the diagnosis, and when it is incorrect, they disconfirm the diagnosis. This feedback regarding environmental cues mimics the problem-solving situation in a range of domains, such as manufacturing defect elimination or computer debugging, while the simplifying assumption of correct or incorrect diagnosis allows us to place the mechanisms of interpretation in stark relief to clarify how they operate.

The recursive interactions between the interpreting and updating processes form a feedback loop. If new cues arrive that increase the plausibility of the leading diagnosis, then the weight on cues decreases slightly, leading to a small increase in the plausibility from new cues, which, in turn, causes updating to further increase the plausibility of the leading diagnosis, and the cycle continues. This interpretation process amplifies a change through a reinforcing feedback process (the loop identifier "R," for reinforcing). We name this feedback process the "self-fulfilling interpretation loop." In the absence of any offsetting influences, this loop pushes the plausibility of an early-generated diagnosis toward ever-greater plausibility. If the loop is driving toward greater plausibility of an erroneous diagnosis, it will generate the well-known self-confirming pattern of fixation, in which an initially plausible diagnosis and the filtering of external cues recursively influence each other so that the problem solver sticks to the diagnosis, despite discrepant cues. If the loop is driving toward greater plausibility of a correct diagnosis, this is salutary. As we demonstrate later, the interplay between this interpretation process and the processes of acting, interpreting cues, and cultivating alternative diagnoses gives rise to the distinctive patterns of problem solving in Rudolph's (2003; see also Rudolph & Raemer, 2004) study.

Action-Oriented Problem Solving Requires Cultivating New Diagnoses

Problem solvers not only assess the plausibility of their leading diagnosis but also consider alternative diagnoses that they identify through search (Gupta, Smith, & Shalley, 2006; March, 1991), conversations (Weick et al., 2005), explanations (Hirt & Markman, 1995), or imagination (Amabile, 1982). This is a knowledge-based activity relying on expertise (Gonzales, Lerch, & Lebiere, 2003; Rasmussen et al., 1994). The process of inferring the most likely explanation for the available information is known as *abduction* (Josephson & Josephson, 1994; Peirce, 1958) and is complementary to *deduction*—forecasting data that would be the consequences of a presumed hypothesis—and *induction*—drawing conclusions from data. Together, abduction, deduction, and induction form a problem-solving cycle of learning- or sensemaking-through-action.

The problem-solving literature suggests a variety of ways for problem solvers to identify the most plausible hypothesis. At one extreme all possible hypotheses are held simultaneously and updated through Bayesian inference as each new piece of information is received (Fischhoff & Beyth-Marom, 1983). In well-structured clinical situations, for example, there are formal protocols that prioritize lists of diagnoses and clinical tests and treatments. Such decision rules are particularly useful for atypical problems or less experienced problem solvers (Gonzalez et al., 2003). The other extreme is represented by Klein's (1998) naturalistic decision-making model, which proposes that for ill-structured problems representative of many real-life situations, problem solvers consider only one hypothesis at a time, mentally simulate the implications of the hypothesis given the available information, and take action if the mental simulation confirms the plausibility of the hypothesis. Only if the mental simulation fails to confirm the hypothesis do the problem solvers imagine a new diagnosis and check it for validity. In the middle, behavioral decision theory recognizes the constructive processes at work in decision behavior (Payne, Bettman, & Johnson, 1992). The problem-solving literature includes models of exemplar-based memory retrieval (Gonzalez et al., 2003, Logan, 1988) and the concept of a race between exemplars and general heuristics to produce a solution or response (Logan, 1988).

For our simulation model we rejected both the extreme Bayesian rationality model and the serial diagnosis model. The behavioral decision-making literature strongly challenges the realism of Bayesian updating, arguing that limited cognitive capacity makes such omniscience humanly impossible, even for well-structured problems (Fischhoff & Beyth-Marom, 1983). Klein's model seems well suited to situations in which the costs of information and the costs of erroneous diagnosis are very high—for example, his fire commanders "test" their diagnosis by going into a building that may collapse around them. On the other hand, if tests are easily run, erroneous diagnoses are easily replaced, and decision makers are not sufficiently expert to see a pattern (Elstein & Schwarz, 2002; Gonzales et al., 2003), it makes sense that problem solvers will hold multiple diagnoses in mind and seek further evidence to distinguish among them. Indeed, rather than being trained not to have hypotheses or to have all possible hypotheses, clinicians are taught to develop a "differential diagnosis," which encourages them to hold more than one (but not "all") diagnoses in mind and perform diagnostic tests that distinguish among the active diagnoses (Barondess & Carpenter, 1994). In Rudolph's study almost all the subjects at some point compared simultaneous diagnoses. We therefore modeled a process that was psychologically reasonable and computationally simple, involving comparison of two potential diagnoses.

We assume that at any one point in time the problem solver has a preferred or leading diagnosis in mind and is seeking to validate or discredit that diagnosis (Elstein et al., 1978; Klein et al., 2006). Drawing on this research and Rudolph's data, we noted that doctors run tests (albeit with high variance in the quality) that will confirm a correct diagnosis or disconfirm an incorrect one. This process of running tests is not Bayesian updating in a mathematical sense but, rather, some combination of logical rule following ("If I can pass something through the breathing tube, then it isn't blocked") and an intuitive totaling up of the supportive and countervailing evidence (Elstein et al., 1978; March, 1994).

We further assume that there is always at least one other diagnosis that the problem solver is imagining (or taking from a well-learned list of candidates) and that the next best alternative among these others is gathering plausibility as the patient's status worsens despite the problem solver's efforts. We call this process "cultivating" and combine it with the processes of acting and interpreting in Figure 4. The "plausibility of alternative diagnosis" is a stock that is increased by cultivating, the pace of which is determined by the parameter "time needed to cultivate." We also show that the pace of cultivating is reduced when the plausibility of the leading diagnosis is high by including a variable labeled "effect of current plausibility on cultivating."

If the alternative diagnosis catches up to the leading diagnosis, we assume that the leading diagnosis is rejected. Three things happen when the leading diagnosis is rejected: (1) the alternative diagnosis becomes the new leading diagnosis, (2) the problem solver switches to action steps appropriate for the new leading diagnosis (so "action steps completed" starts over), and (3) yet another diagnosis becomes the second-place alternative diagnosis (so "plausibility of alternative diagnosis" starts over). The model variable "change trigger" signals the rejection of the leading diagnosis, and we use dotted lines in Figure 5 to signal these changes.

As we translated this causal structure into a formal mathematical model, we initially set parameters to reasonable values in the context of our motivating clinical example and refined them as we compared simulation output to our source data. We also conducted extensive sensitivity analysis to test for model robustness under extreme conditions. Complete documentation of the model equations appears in Appendix B.

SIMULATING THE DYNAMICS OF PROBLEM SOLVING

We begin with a set of experiments[3] showing how the interplay of acting, interpreting, and cultivating new diagnoses produces the four modes of diagnostic problem solving observed in Rudolph's study (2003). For clarity of exposition, we have chosen a scenario that controls for the effects of random search by assuming that all problem solvers generate alternative diagnoses in

[3] "Experiment" is commonly used in the modeling community to refer to manipulations of the model parameters (Carley, 2001).

FIGURE 4
Core Model Structure Showing the Interaction of Acting, Interpreting, and Cultivating Alternatives

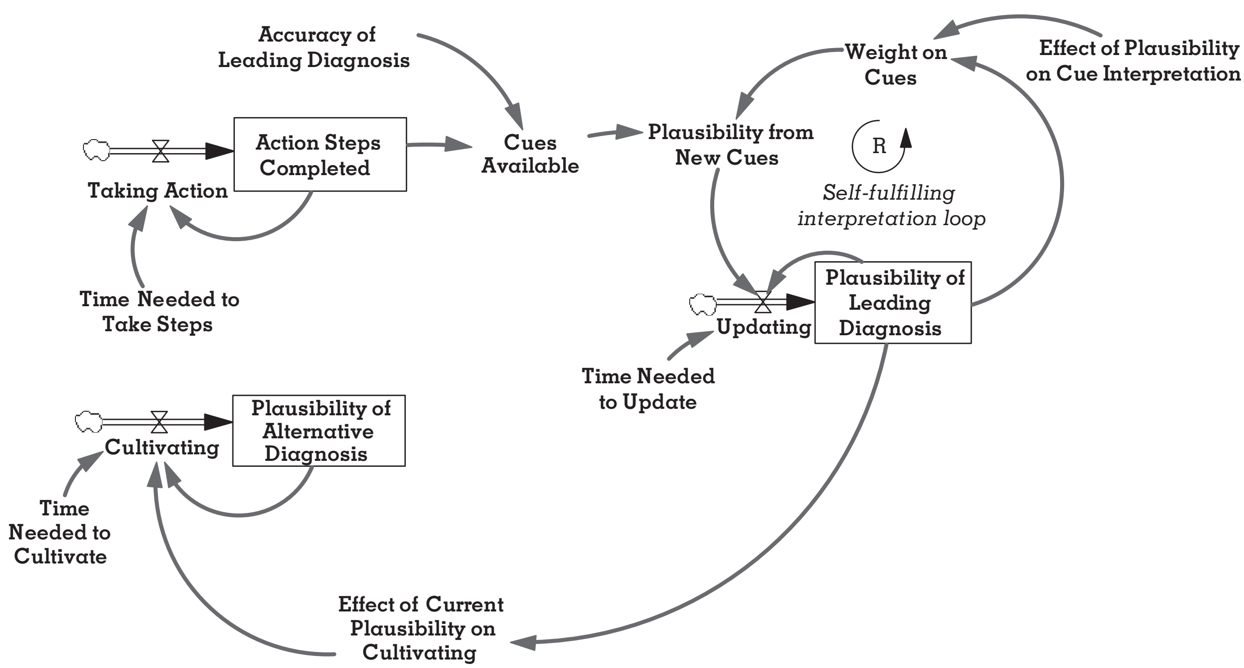

FIGURE 5
Changes to Reset the Model When the Alternative Diagnosis Becomes the Leading Diagnosis

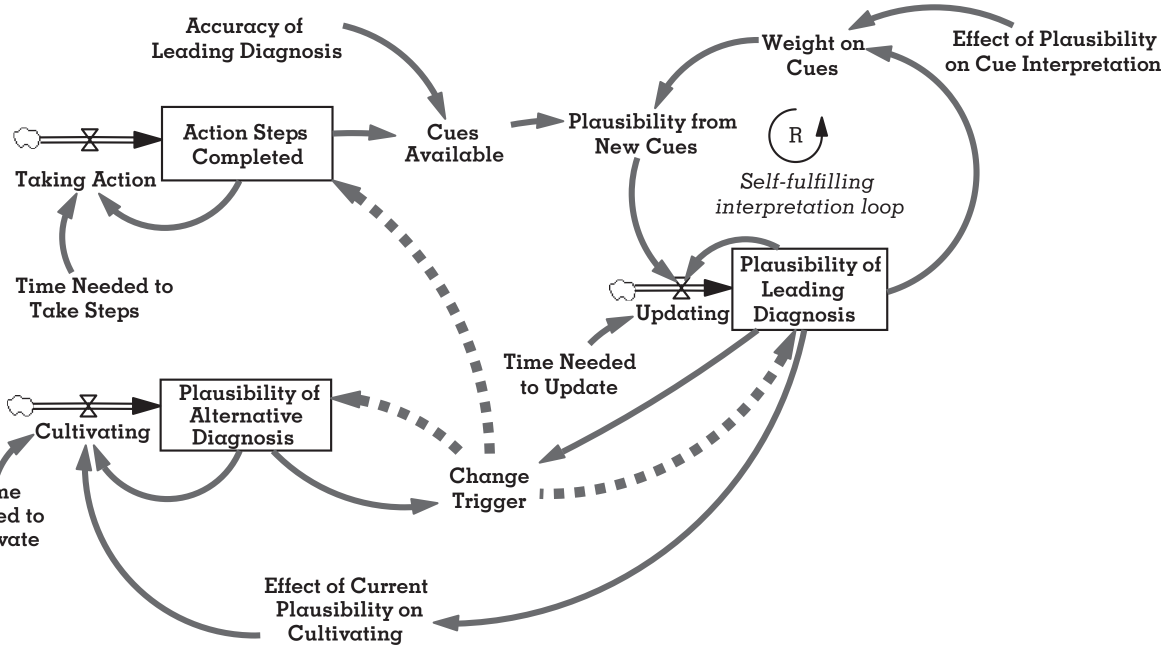

the same modal sequence in Rudolph's data, in which only the fourth diagnosis is correct.[4]

Four Modes of Action-Oriented Problem Solving

To highlight differences among the four problem-solving modes, Figure 6 displays the behavior over time of the plausibility of the leading diagnosis. The top panel is an illustration of the adaptive mode. The problem solver's sense of the plausibility of the first diagnosis begins at its initial value of 0.5 (out of 1.0), and the dynamics of the three component problem-solving processes begin to unfold simultaneously.

First, the problem solver begins taking action steps associated with the first diagnosis, increasing the stock of action steps completed, which results in more cues available. Second, armed with some confidence in his or her diagnosis, the problem solver's interpretations begin to increase plausibility, and the weight on cues falls slowly as the self-fulfilling interpretation loop acts to reinforce the leading diagnosis. In the first few moments the action steps of the diagnostic algorithm have not progressed much, so the limited cues have little effect on plausibility. After a short time the accumulated cues (which are "objectively" disconfirming information because the first diagnosis is incorrect) begin to show their effect on plausibility, and we see a slow decline in the plausibility of the leading diagnosis. Third, the plausibility of the alternative diagnosis builds as the cultivating process unfolds in the face of cues unfavorable to the leading diagnosis. Eventually, the plausibility of the alternative diagnosis overtakes the plausibility of the leading diagnosis. At this moment the first diagnosis is rejected and the second diagnosis becomes the leading diagnosis.

The problem solver begins pursuing the algorithm associated with the new leading diagnosis. The pattern repeats for the second and third diagnosis. When the problem solver begins to consider diagnosis number four, the correct one, plausibility begins to grow as before. However, the new cues available now offer confirmation and are interpreted to build even more plausibility of the leading diagnosis. Moreover, the

[4] Simulation analyses not shown here replicate the main results for scenarios in which the correct diagnosis enters earlier or later than fourth, in which there are two correct diagnoses, and even in which the poor diagnoses have a modest degree of correctness. A summary of several hundred additional simulations demonstrating these results is available from the second author on request.

FIGURE 6
Modes of Action-Oriented Problem Solving

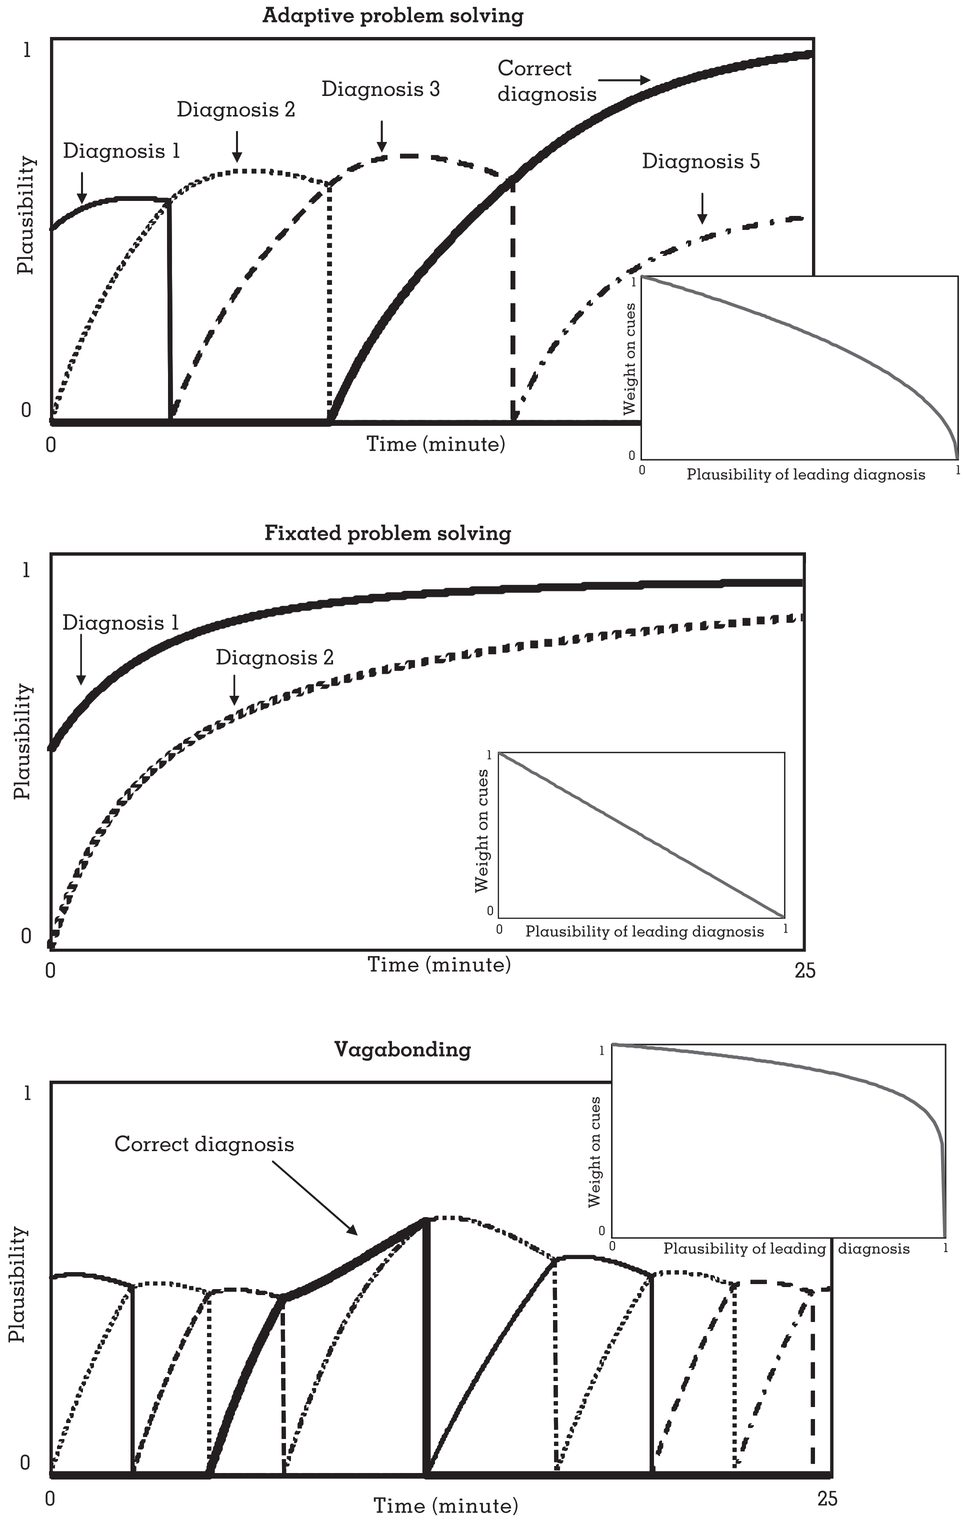

Note: Simulation conditions are identical except for the strength of the self-fulfilling interpretation feedback loop, as determined by the relationship between "plausibility of leading diagnosis" and "weight on cues," shown in the insets.

self-fulfilling interpretation loop reinforces the increases in plausibility, reducing the weight on cues and thus boosting plausibility still further. The problem solver pursues the action steps to completion and converges on a steady-state choice of the correct diagnosis.

The top panel in Figure 6 shows two important features of problem-solving dynamics. First, the consideration of each diagnosis enjoys a honeymoon period during the time it takes for an alternative diagnosis to emerge as a viable contender as the basis for action—a temporal interplay between the leading and alternative diagnoses. In the adaptive mode this temporal interplay is "well balanced" in that the honeymoon period is long enough for the problem solver both to take action and to interpret the results stirred up by that action. Second, there is a dynamic interchange in the roles of acting and interpreting because the cues available accumulate slowly relative to the ongoing process of interpreting experience. So we see in Figure 6 that plausibility increases at first (even for incorrect diagnoses), but it later (for each incorrect diagnosis) decreases. Plausibility increases at first because updating driven by the confirmation-biased interpretation process occurs quickly relative to the accumulation of available cues. Meanwhile, the problem solver continues to interact with the physical world by taking action steps that generate more cues. As the disconfirming evidence mounts, it eventually overcomes the effects of the self-fulfilling interpretation: plausibility reaches a peak and then begins to decline.

The middle panel of Figure 6 shows simulation results that replicate the fixating mode. The difference between this experiment and the one in the top panel is only that the effect of plausibility on cue interpretation, depicted in the inset panel, is stronger.[5] The simulation begins as before with an initial plausibility that starts to rise at first, but the lower weight on cues in this scenario allows the self-fulfilling process to gain momentum. The problem solver acts, interpreting cues and creating meaning that supports the current diagnosis, and the weight on cues falls even more. The self-fulfilling interpretation loop reinforces the current diagnosis, and because the loop is so strong, the first diagnosis is always preferred. The diagnostician does not move on to any other diagnosis. The strong reinforcing effects of the self-fulfilling interpretation loop result in a pattern of problem solving in which the problem solver is completely confident in the incorrect diagnosis. Self-fulfilling interpretations discount some disconfirming evidence, so the current diagnosis locks in prematurely, squeezing out the cultivation of alternatives, and the problem solver never has a chance to find the correct diagnosis.

The bottom panel in Figure 6 shows simulation results that replicate the vagabonding mode. In this experiment the effect of plausibility on cue interpretation is weaker.[6] The first three diagnoses are rejected, but more quickly than in the adaptive case, implying that the problem solver does not take as many action steps, consistent with field data showing that diagnostic vagabonds generated diagnoses but performed few or no steps of the treatment/test algorithm (1.5 steps on average). When the fourth diagnosis, the correct one, enters as the leading diagnosis, plausibility increases, but not as rapidly as in the adaptive case. The problem solver places a higher weight on cues (owing to a weaker effect of plausibility on cue interpretation), but the cues to confirm the diagnosis accumulate somewhat slowly because they must be made available by advancing through action steps. Meanwhile, an alternative diagnosis gains favor and eventually overtakes the correct diagnosis, and the problem solver also rejects the correct diagnosis. Once this diagnosis is rejected, the problem solver continues identifying alternatives, choosing them as the leading diagnosis, and rejecting them in favor of the next emerging alternative.

The stylized problem solver in the vagabonding experiment is quite capable of cultivating alternatives and attending to cues but, lacking more confident beliefs about the plausibility of a diagnosis, does not hold onto a diagnosis long enough to adequately advance forward with action steps. The result is vagabonding, a pattern of diagnostic problem solving in which the problem solver jumps from one plausible diagnosis to the next without treating the patient. The dynamic interplay among acting, interpreting, and cultivating alternatives is out of balance: the

[5] Specifically, Effect of Plausibility on Cue Interpretation = 1.

[6] Specifically, Effect of Plausibility on Cue Interpretation = 0.15.

pace of generating new cues associated with the leading diagnosis (acting) is too slow relative to the pace of cultivating alternatives. The problem solver gets stuck in a mode of generating new alternatives but not discovering enough about them to reach an effective conclusion. This mode fails because the effect of plausibility on interpretation is so weak that even the correct diagnosis is rejected.

The model can also generate a mode in which the problem solver is stalled—unable to move forward to take any action steps (e.g., when both taking action steps and cultivating diagnoses are extremely slow). Rudolph's analysis classified only two out of thirty-nine doctors as stalled. Both exhibited behaviors of advancing treatment steps little or not at all and establishing working diagnoses very slowly, consistent with this example. However, with so few examples and so little action to learn from, we omit this mode from subsequent analysis.

The Interplay of Acting, Interpreting, and Cultivating Diagnoses: Sensitivity Analysis

To examine the critical dynamic interactions among acting, interpreting, and cultivating diagnoses more closely, we conducted experiments in which we varied the pace of these processes. The first set of simulations held all parameters the same as in the vagabonding case of Figure 6, except that we varied the pace of acting.[7] We discovered that the eleven simulation runs generated, shown in Figure 7, separate into two distinct patterns. The set corresponding to faster acting is adaptive: the plausibility of diagnosis number four climbs smoothly toward one. The other set, based on slower acting, displays vagabonding: diagnosis number four is rejected and new alternatives continue overtaking the lead.

This experiment points to two important results. First, different rates of taking action generate qualitatively different dynamics. A bias for action (taking action steps faster) offsets the effects of less self-fulfilling interpretation and protects the problem solver from incorrectly rejecting the correct diagnosis. Second, small differences in the rate of acting can mean the difference between adaptive sensemaking and vagabonding. This result raises the question of just what pace of taking action is needed to escape from the perils of vagabonding.

To shed light on this question, we conducted an extensive set of experiments to test the relationship among the pace of acting, the pace of

[7] We did this by setting the time needed to take steps to values ranging from very fast (1 minute) to very slow (16 minutes).

FIGURE 7
Sensitivity Analysis Showing System Behavior for Various Rates of Taking Action Steps

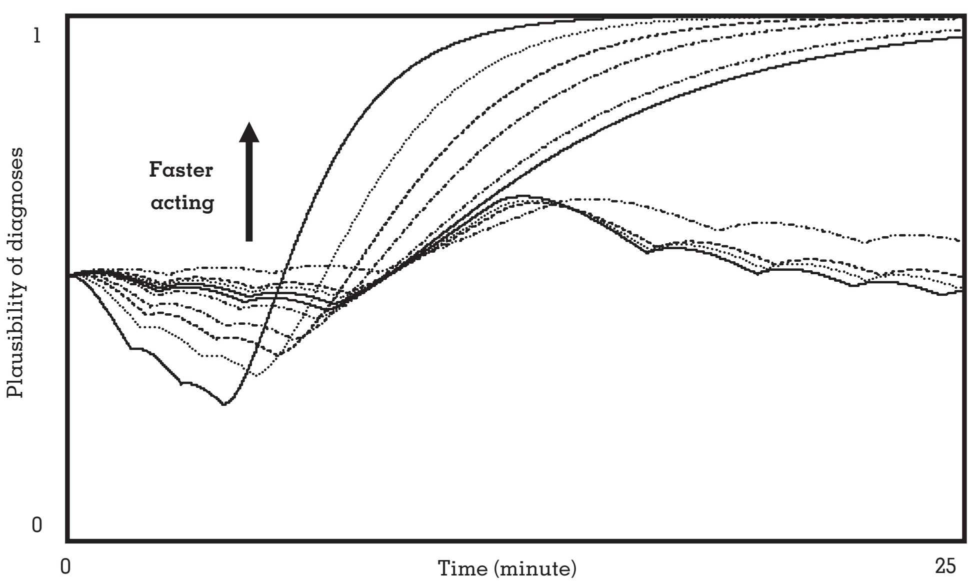

cultivating alternatives, and the strength of the effect of plausibility on interpretation. We performed many sets of simulations like those in Figure 7 for various values of the effect of plausibility on cue interpretation and the time needed to cultivate. For each combination of these two parameters, we found the threshold pace of time needed to take steps that separates vagabonding from adaptive sensemaking and the threshold that separates adaptive sensemaking from fixating.[8] Each of the four panels in Figure 8 shows the results for one value of the time needed to cultivate. Looking first from left to right in any panel, as the strength of the interpretation effect increases from very weak to moderate, the threshold pace of acting needed to avoid vagabonding gets slower. For interpretation effects of moderate strength, even a very slow pace of acting will lead to adaptive problem solving. But, as the interpretation effect gets very strong, the threshold pace of acting needed to avoid fixation gets faster.

The panel views in Figure 8 highlight the interactions between the pace of action and the strength of the interpretation effect. At the left of each panel, the weaker effect of the plausibility on cue interpretation requires faster action to avoid vagabonding. A weak interpretation effect describes a problem solver who wants more cues, so the pace of acting must be faster to lead to adaptive sensemaking. When the appetite for cues is high (weak effect of plausibility on cue interpretation), slow action induces vagabonding. But faster action may not always be possible or desirable. In such cases a modest degree of confidence in the leading diagnosis (i.e., a stronger effect of plausibility on interpretation) thwarts the lurking threat of vagabonding. At the right of each panel, where there is a stronger

[8] To find the thresholds, we classified a simulation as vagabonding if diagnosis number four was rejected, as fixated if diagnosis number one was not rejected by minute 30, and as adaptive otherwise.

FIGURE 8
Threshold Values for the Pace of Taking Action

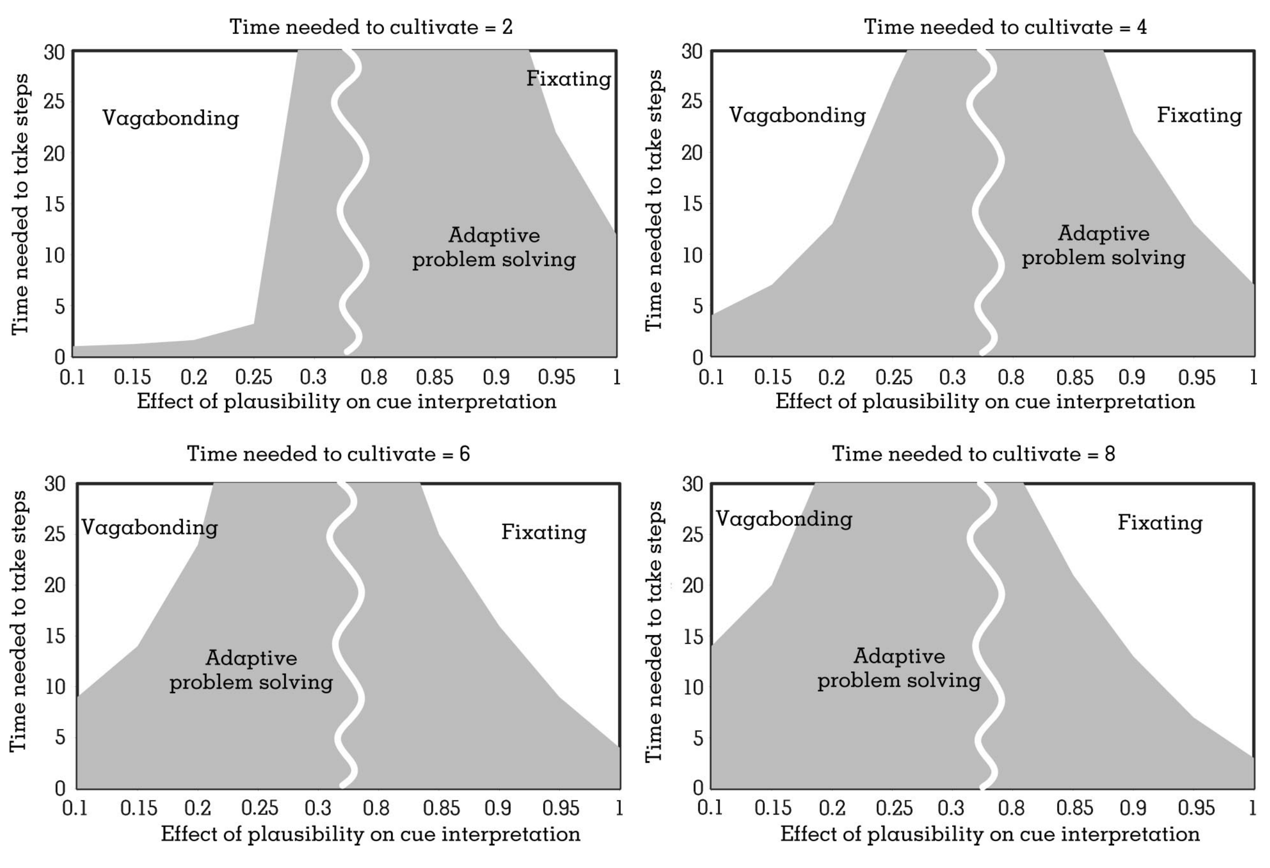

effect of plausibility on cue interpretation, faster action helps deter fixation. A strong interpretation effect describes a problem solver who pays little attention to cues, so the pace of acting must be faster to generate the overwhelming amount of disconfirming evidence needed to pierce the high sense of plausibility. When the attention to cues is low (a strong effect of plausibility on cue interpretation), slow action allows premature convergence on an incorrect diagnosis. Alternatively, a modest degree of skepticism in the leading diagnosis (i.e., weaker effect of plausibility) prompts an openness to cues that can prevent fixation.

Looking across the family of curves in the four panels in Figure 8, we see how the thresholds for the pace of taking action depend on both the effect of plausibility on cue interpretation and the time needed to cultivate alternate diagnoses. When the pace of cultivating alternatives is very fast, the risk of vagabonding is quite high and not mitigated much by stronger interpretation effects. Very rapid action is still needed. For a slower pace of cultivating alternatives, small increases in the strength of the interpretation effect yield stronger buffers against the threat of vagabonding, so slower paces of action are still adequate to avoid vagabonding. However, when the pace of cultivating alternatives is very slow, the risk of fixating is high and not mitigated much by weaker interpretation effects. Moreover, in contrast to the interaction around the vagabonding threshold, for a *faster* pace of cultivating alternatives, small decreases in the strength of the interpretation effect yield larger improvements in avoiding fixation.

DISCUSSION

Our theory of action-oriented problem solving is intended to be broadly applicable yet not universal. Three boundary conditions are reflected in the structure of the model and present important caveats for interpreting our findings and the propositions we develop next. They are (1) action-based inquiry—information cues become available only by taking action; (2) temporal dynamism—doing nothing or holding an erroneous diagnosis means the situation is deteriorating over time, so there is pressure on the problem solver to act on the leading diagnosis or to cultivate a new one; and (3) action endogeneity—moving through the steps of a course of action changes the characteristics of the environment and the stream of cues that become available.

The modeling exercise and theory development process contribute three new insights to understanding the mechanisms that link interpretation and choice in action-oriented problem solving. First, while current theories of sensemaking and decision making specialize in examining interpretation and choice, respectively, we find that they are inextricably linked through the interplay of acting, interpreting, and cultivating diagnoses. Second, we show how the interactions and pacing among acting, interpreting, and cultivating diagnoses tip people between adaptive problem solving on the one hand and fixation or vagabonding on the other. Third, we highlight the surprising result that reinforcing feedback, usually seen as the driver of maladaptive patterns such as fixation, can actually play a beneficial role in adaptive problem solving.

Linking Sensemaking and Decision-Making Perspectives

Our motivating example of anesthesiology residents, like many real-world managerial problem-solving situations, includes processes familiar from both sensemaking and decision-making theories. Consistent with the sensemaking literature, interpretation drives action, and the cues generated by such action then serve as inputs to further sensemaking. Such cues are not immediately meaningful inputs but, rather, have to be interpreted or made sense of, analogous to "editing" in prospect theory (Kahneman & Tversky, 1979). New diagnoses are generated through a process that mimics "insight" (Langley et al., 1995), and they are then updated as a function of incoming information. Consistent with the decision-making literature, choices are being made between the leading and the second diagnoses, based on their relative plausibility. While "choosing" is not a named variable in our model,[9] choosing between the leading and the second diagnoses is a continual option,

[9] Our model does not include a choice among actions; instead, it makes the simplifying assumption that problem solvers are following a protocol that specifies the proper test or treatment given the leading diagnosis; in situations without such a protocol, an expanded model would include a choice among actions as well as a choice among diagnoses.

and the closeness of the contest between the two choices varies as the plausibility of each diagnosis waxes and wanes.

In action-oriented problem solving, choice follows sensemaking. Indeed, there can be no choice without a previous sensemaking process; action and interpretation are needed to make available the information used for choice. And sensemaking follows choice as the selection of a new guiding hypothesis sets the problem solver on a new course of action. In contrast, classic rational choice approaches posit that preconstituted information is given and perfectly understood, and the decision process is decoupled from the sensemaking process that created the preconstituted information.

This leads us to propose our first proposition.

> *Proposition 1: In the context of action-oriented problem solving, the outputs of sensemaking and decision making are inputs to each other.*

Failure Modes and Success: The Dynamics of Acting, Interpreting, and Cultivating Diagnoses

The pace of acting, interpreting, and cultivating new diagnoses and the strength of the self-fulfilling interpretation loop interact to produce both the successful adaptive problem-solving mode and the three failure modes of fixation, vagabonding, and stalling. For example, although an open mind (in our model, a high weight on cues) coupled with deliberation is often crucial for effective action, in action-oriented problem solving, acting quickly, even on an uncertain diagnosis, can produce information needed to improve the diagnosis. In many situations characterized by our boundary conditions, rapid action is essential to generate diagnostic information and avert impending deterioration of a strategic position, a patient, or a wildfire. But how much confirmation bias or self-fulfilling interpretation is enough or too much? When is acting quickly an asset? When is it a liability? Our model clearly shows that there is no one answer to such questions; each problem-solving mode is a kind of "syndrome" produced by a unique combination of the pace and power of different processes. Our results, however, suggest some patterns.

Adaptive problem solving calibrates and balances self-confirmation, speed of acting, and speed of cultivating alternatives. The window for adaptive problem solving gradually closes as the competition between the leading and the alternative diagnosis gets tighter. Just at the moment when the problem solver could disconfirm the leading diagnosis and switch to another, the momentum of a strong effect of plausibility on cue interpretation sweeps him or her into fixation—or just at the moment when sticking with the leading diagnosis would help advance the appropriate action steps, rapid generation of alternative diagnoses combined with slow accumulation of cues sweeps him or her into vagabonding. Analogously, organizations can fall into dysfunctional modes by over- or underemphasizing the speed of decision making (Eisenhardt, 1989; Perlow et al., 2002). Hence, rather than emphasizing beneficial processes over dysfunctional biases, our model shows that the same processes that drive adaptive problem solving, when out of balance, lead to dysfunctional problem solving in the form of fixation or vagabonding.

The simulation results suggest that faster action is beneficial in creating a larger region of adaptive problem solving (Figure 8). Our results show no detrimental effects of rapid action, but in many cases rapid action may be problematic, such as when *not* acting on short-term trends is beneficial (cf. March 1991; Sastry, 1997), or when fast action itself transforms the expected pace of adaptation in an entrepreneurial environment in a way that is harmful to all players (Perlow et al., 2002). However, when these "side effects" are of little concern, a faster test or faster information feedback is better, all else being equal. Indeed, in the extreme case where action and information feedback are instantaneous, the problem approximates a static decision-making choice in which information is preconstituted and given.

> *Proposition 2: In action-oriented problem solving, faster action decreases the risk of failure from fixating or vagabonding.*

The results also show how problem solvers should adapt to different paces of change in the environment. Consider the challenge of how quickly to cultivate alternative explanations. A slower pace of cultivating alternative hypotheses is useful to reduce the threat of vagabonding, but too slow a pace may increase the risk of fixation. Cultivating new diagnoses exerts a constant pressure on the problem solver to improve the situation, either by advancing a current course of action or switching to a different course. The extent to which this pressure is needed depends on the pace at which the problem situation is worsening—that is, the degree of temporal dynamism. Slower cultivation keeps open the window of opportunity to more fully explore a leading diagnosis. An environment that is changing rather slowly offers the luxury of slow cultivation of alternatives. But when the status of the environment is rapidly deteriorating, even short stints of holding onto an erroneous leading diagnosis can lead to disaster.

Proposition 3: In action-oriented problem solving, the pace of cultivating alternatives should correspond to the speed of undesirable change in the environment.

Other features of the environment—particularly, how quickly or slowly cues become available and how time consuming it is to generate alternatives—can exacerbate or inhibit problem-solving failures. Our analysis shows that a rapid pace of acting, even when interpretation effects are weak, prevents vagabonding, and a rapid pace of cultivating alternatives fosters vagabonding.

Proposition 4: Vagabonding will be more likely when action is constrained or information feedback is delayed.

Proposition 5: Vagabonding will be more likely when it is relatively easy to generate plausible alternatives.

Similarly, our analysis shows that either a rapid pace of acting when interpretation effects are strong or a rapid pace of cultivating alternatives reduces the risk of fixation. People are more susceptible to fixation when the rate at which cues emerge and the rate at which new diagnoses are cultivated are both relatively slow in relation to the pace at which plausibility builds. Working through action steps more quickly can provide the necessary cues to weaken the self-fulfilling interpretation loop and also allow new data to speed up the pace of cultivating alternatives.

Proposition 6: Fixation will be more likely when action is constrained or information feedback is delayed.

Proposition 7: Fixation will be more likely when it is relatively difficult to generate plausible alternatives.

Overreliance on a well-learned combination of action pacing, self-fulfilling interpretation, and rate of cultivating new diagnoses can derail problem solving when environmental demands change. The strength of the effect of plausibility on cue interpretation moderates the balance between acting and cultivating. Problem solvers prone to either vagabonding or fixation have three possible solutions. To avoid vagabonding, they may speed up action, hold diagnoses more confidently, or slow down cultivating. To avoid fixation, they may speed up action, hold diagnoses less confidently, or speed up cultivating. Some problem solvers will use a narrow range of combinations of these tactics in their action-oriented problem solving, whereas others, perhaps more expert, will develop a wider range of solution sets and the ability to better calibrate them to the situation at hand.

Proposition 8: Effectiveness in action-oriented problem solving is a function of people's ability to vary the pace of action, intensity of self-fulfilling interpretation, and pace of cultivating new diagnoses in relation to environmental demands.

Reinforcing Loops and Confirmation Bias Can Be Beneficial

Our modeling results also highlight the counterintuitive benefits of reinforcing loops and confirmation bias, which both tend to be associated with vicious cycles, as in theories of fixation and escalation of commitment (Arkes & Blumer, 1985; Staw, 1976). Confirmation bias is generally seen as either present or not and as undesirable when present, but our model shows that this bias varies in degree and sometimes helps. Although an overly strong reinforcing process may lead to fixation, diagnostic vagabonds can benefit from a moderately strong self-fulfilling process that allows them to sustain attention to the correct diagnosis long enough to find the right corrective action. Appendix A provides four hypothetical examples of confirmation bias that can produce sensemaking modes ranging from adaptive problem solving to fixation to vagabonding.

The unexpected benefit of a self-fulfilling interpretation loop emerges in situations in which it is faster to generate an alternative idea than to take the action steps needed to evaluate the merits of a particular course of action, especially in the face of incomplete, delayed, and ambiguous information. Some inertia is necessary to launch and sustain action. Self-fulfilling interpretation allows the problem solver to "preserve the frame" (Klein et al., 2006) while the leading diagnosis gathers support, shielded from the pressure to jump to a new diagnosis. Like hill climbing in a rugged solution terrain, some self-fulfilling interpretation allows problem solvers to stay the course beyond a couple of short-term failures to ascertain if they are indeed real failures (Fleming & Sorenson, 2004). We therefore suggest the following proposition.

> *Proposition 9: When accessing cues from the environment is a lengthy process and generating alternatives is quick and easy, a moderate degree of self-fulfilling interpretation is beneficial.*

Limitations

Given the interacting constructs and relationships in the model, there is more than one way to generate the problem-solving behaviors observed by Rudolph (2003). Our argument is not that our model is an exact or complete representation of psychological processes in action-oriented problem solving but, rather, that the variations in model behavior have been sufficiently realistic and informative to provide insights about theory. The motivating study that sparked this theorizing examined doctors at an intermediate level of training working on one clinical problem in one specialty. It is likely that results might be different in different clinical specialties, different professions, and at different levels of expertise. However, by articulating mechanisms in a stylized way, we have suggested constructs and relationships that could guide further empirical exploration as well as further theory development.

Improving Problem Solving in Practice

Whereas researchers and practitioners have previously identified several ways to avoid fixation, some of these may increase vagabonding. We suggest three general strategies for individuals and organizations to improve problem solving in practice: awareness, training, and task design. Managers, trainers, and individual professionals need to enhance their awareness of problem-solving patterns so that they can assess their tendencies toward fixation, vagabonding, stalling, and adaptive problem solving. Even in the heat of the moment, professionals can be more self-aware, recognizing repeat treatments as a signal of fixation and multiple disconnected actions as a signal of vagabonding.

Training is a classic and effective approach to knowledge-based work. By recognizing the problem-solving patterns, their symptoms, and causes, training can be enhanced. For example, professionals such as managers, executives, disaster first responders, and clinicians could use full-field or computer simulations to experience the "feel" of acting too slowly or too quickly while "keeping the bubble" of understanding the situation and its possibilities. Problem solvers could learn to recognize the moment at which two diagnoses have similar plausibility and when it is critical for them to observe their own problem-solving process to calibrate the strength of the self-fulfilling interpretation and the balance between rates of generating new diagnoses and accumulating information.

Problem-solving environments are often presumed to be given, but task design and coordination processes strongly affect outcomes. In Rudolph's study residents were expected to use protocols that simplified their diagnostic challenges; the protocols were based on accumulated clinical experience and possibly controlled experimentation. In hospitals, manufacturing plants, and computer programming labs, the flow of information from action depends at least as much on the quality of organizational processes, coordination, and support (the speed and quality of lab tests, prototypes, and beta tests) as on the mental activity of the manager or clinician (Gittell, 2009). Managers can rethink their policies and investment priorities to relieve constraints that limit the pace of action or the pace of generating alternatives.

APPENDIX A

Figure A-1 describes different relationships between the plausibility ascribed to the leading diagnosis and how heavily the problem solver weights cues. Looking first at the axes of the figure, when the plausibility of the leading diagnosis is at its extreme value of 0 (the problem solver perceives the leading diagnosis as completely implausible), the weight on cues is at its extreme value of 1 (the problem solver pays full attention to cues). Conversely, when plausibility equals 1, the weight on cues is equal to 0 (except when the effect of plausibility on cue interpretation is equal to 0, in which case no matter how plausible or implausible the problem solver deems his or her current diagnosis, he or she always gives full weight to cues).

When the effect of plausibility on cue interpretation is equal to 0.15, which characterizes a cautious problem solver, even large increases in the plausibility of the leading diagnosis do not much diminish the weight the problem solver places on cues. It is not until he or she is almost completely certain the diagnosis is plausible that the weight on cues diminishes. When the effect of plausibility on cue interpretation is 0.5, describing a moderately bold problem solver, decreases in weight on cues are greater for a given increase in the plausibility of the leading diagnosis. When the effect is equal to 1, depicting a slightly bolder problem solver, a decrease in plausibility brings about a proportional decrease in weight on cues.

Weight on Cues (t) = (1 − Plausibility of Leading Diagnosis (t)) ^ Effect of Plausibility on Cue Interpretation,

where the exponent "effect of plausibility on cue interpretation" is a parameter chosen to represent possible individual and/or situational differences.

FIGURE A-1
Weight on Cues As a Function of Plausibility of Leading Diagnosis for Various Settings of Effect of Plausibility on Cue Interpretation

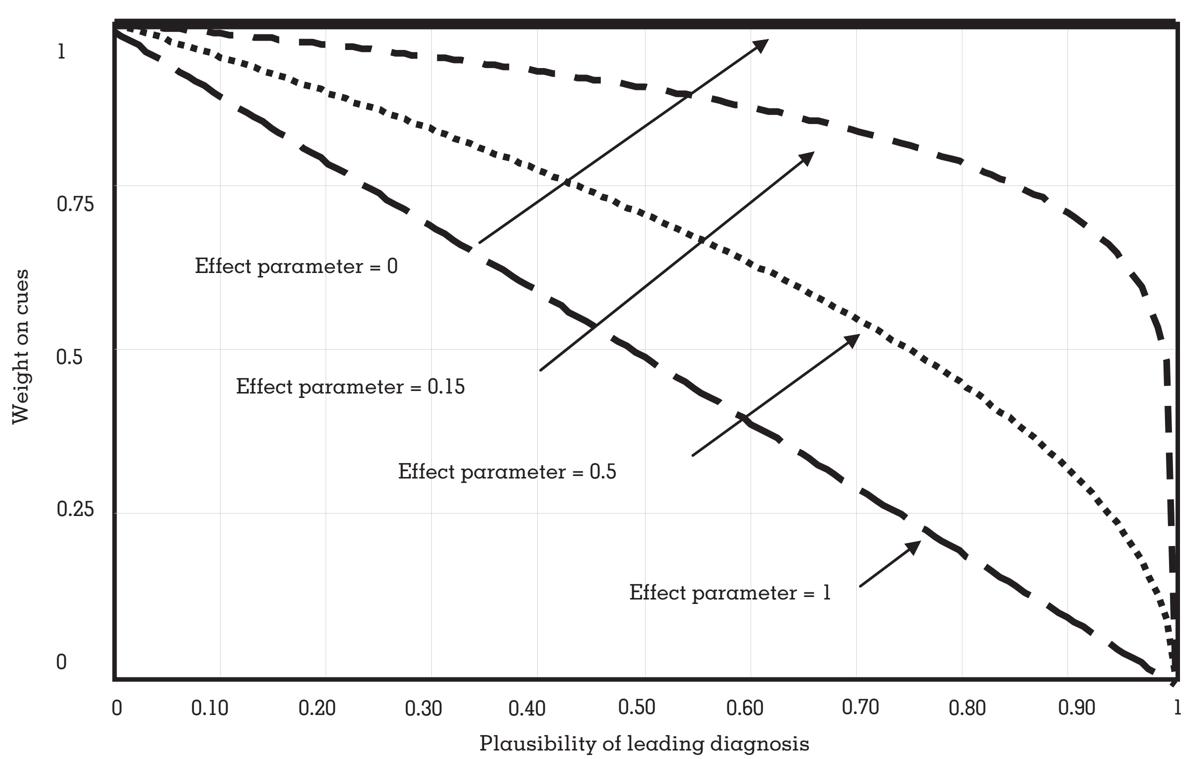

APPENDIX B

Integral equations are written in this appendix using the following notation: Stock = INTEG (Inflow − Outflow, Initial Value of Stock), where the INTEG function means the integral from time 0 to time *t* (the current time) of the inflow less the outflow plus the initial value of the stock. The model is simulated using Vensim DSS software (available from www.vensim.com).

Equations for the Acting Subsection (Figure 1)

Action Steps Completed = INTEG (Taking Action − Resetting Action Steps, 0)
Units: Dimensionless

Taking Action = (1 − Action Steps Completed)/Time Needed to Take Steps
Units: Dimensionless/Minute

Time Needed to Take Steps = 8
Units: Minute

Cues Available = (Starting Plausibility of Leading Diagnosis + Action Steps Completed * (Accuracy of Leading Diagnosis − Starting Plausibility of Leading Diagnosis))
Units: Dimensionless

Accuracy of Leading Diagnosis =
IF THEN ELSE (Current Diagnosis = True Diagnosis, 1, 0)
Units: Dimensionless

True Diagnosis = 4
Units: Dimensionless

Equations for the Interpreting Subsection (Figures 2 and 3)

Plausibility of Leading Diagnosis = INTEG (Updating + Carry Over to Leading − Resetting Leading, Initial Plausibility)
Units: Dimensionless

Updating = (Plausibility from New Cues − Plausibility of Leading Diagnosis)/Time Needed to Update
Units: Dimensionless /Minute

Plausibility from New Cues =
Cues Available * Weight on Cues + (1 − Weight on Cues)
Units: Dimensionless

Time to Needed Update = 2
Units: Minute

Weight on Cues = (1 − Plausibility of Leading Diagnosis) ^ Effect of Plausibility on Cue Interpretation
Units: Dimensionless

Effect of Plausibility on Cue Interpretation = 0.5
Units: Dimensionless

Equations for the Cultivating Alternatives Subsection (Figure 4)

Plausibility of Alternative Diagnosis =
INTEG (Cultivating − Resetting Alternative, 0)
Units: Dimensionless

Cultivating = Effect of Current Plausibility on Cultivating * **(1 −** Plausibility of Alternative Diagnosis)/Time Needed to Cultivate
Units: Dimensionless/Minute

Effect of Plausibility on Alternative = min (1, 2 − 2 * Plausibility of Leading Diagnosis)
Units: Dimensionless

Time Needed to Cultivate = 4
Units: Minute

Equations for Switching Diagnoses (Figure 5)

Change Trigger = IF THEN ELSE (Plausibility of Leading Diagnosis < Plausibility of Alternative Diagnosis, 1, 0)/ TIME STEP
Units: Dimensionless/Minute

Resetting Action Steps = Action Steps Completed * Change Trigger
Units: Dimensionless/Minute

Resetting Leading = Plausibility of Leading Diagnosis * Change Trigger
Units: Dimensionless/Minute

Carry Over to Leading = Resetting Alternative
Units: Dimensionless/Minute

Resetting Alternative = Plausibility of Alternative Diagnosis * Change Trigger
Units: Dimensionless/Minute

Starting Plausibility of Leading Diagnosis = INTEG (New Plausibility − Resetting Starting Plausibility, Initial Plausibility)
Units: Dimensionless

New Plausibility = Resetting Alternative
Units: Dimensionless/Minute

Resetting Starting Plausibility = Change Trigger * Starting Plausibility of Leading Diagnosis
Units: Dimensionless/Minute

Initial Plausibility = 0.5
Units: Dimensionless

Current Diagnosis = INTEG (Diagnosis Counter, 1)
Units: Dimensionless

Diagnosis Counter = Change Trigger
Units: Dimensionless/Minute

REFERENCES

Amabile, T. M. 1982. Social psychology of creativity: A consensual assessment technique. ***Journal of Personality and Social Psychology,*** 43: 997–1013.

Arkes, H. R., & Blumer, C. 1985. The psychology of sunk cost. ***Organizational Behavior and Human Decision Processes,*** 35: 124–140.

Arthur, W. B. 1990. Positive feedbacks in the economy. ***Scientific American,*** 268: 92–99.

Barondess, J., & Carpenter, C. (Eds.). 1994. ***Differential diagnosis.*** Philadelphia: Lea & Febiger.

Bartunek, J. M. 1984. Changing interpretive schemes and organizational restructuring: The example of a religious order. ***Administrative Science Quarterly,*** 29: 355–372.

Black, L., Carlile, P., & Repenning, N. P. 2004. A dynamic theory of expertise and occupational boundaries in new technology implementation: Building on Barley's study of CT scanning. ***Administrative Science Quarterly,*** 49: 572–607.

Cannon-Bowers, J. A., & Salas, E. 1998. ***Making decisions under stress: Implications for individual and team training.*** Washington, DC: American Psychological Association.

Carley, K. 2001. Computational approaches to sociological theorizing. In J. Turner (Ed.), ***Handbook of sociological theory:*** 69–84. New York: Kluwer Academic/Plenum.

Carroll, J. S., Rudolph, J. W., & Hatakenaka, S. 2002. Learning from experience in high-hazard industries. ***Research in Organizational Behavior,*** 24: 87–137.

Carroll, J. S., Rudolph, J. W., Hatakenaka, S., Wiederhold, T. L., & Boldrini, M. 2001. Learning in the context of incident investigation team diagnoses and organizational decisions at four nuclear power plants. In E. Salas & G. Klein (Eds.), ***Linking expertise and naturalistic decision making:*** 349–365. Mahwah, NJ: Lawrence Erlbaum Associates.

Cook, R. I., & Woods, D. 1994. Operating at the sharp end: The complexity of human error. In B. S. Bogner (Ed.), ***Human error in medicine:*** 255–310. Hillsdale, NJ: Lawrence Earlbaum Associates.

Cyert, R. M., & March, J. G. 1963. ***A behavioral theory of the firm.*** Englewood Cliffs, NJ: Prentice-Hall.

Davis, J. P., Eisenhardt, K. M., & Bingham, C. B. 2007. Developing theory through simulation methods. ***Academy of Management Review,*** 32: 480–499.

De Keyser, V., & Woods, D. D. 1990. Fixation errors: Failures to revise situation assessment in dynamic and risky systems. In A. G. Colombo & A. Saiz de Bustamante (Eds.), ***Systems reliability assessment:*** 231–251. Amsterdam: Kluwer.

Dörner, D. 1997. ***The logic of failure: Recognizing and avoiding error in complex situations.*** New York: Perseus.

Dreyfus, H. L. 1997. Intuitive, deliberative, and calculative models of expert performance. In C. E. Zsambok & G. Klein (Eds.), ***Naturalistic decision making:*** 17–28. Mahwah, NJ: Lawrence Erlbaum Associates.

Edmondson, A. 2002. Disrupted routines: Team learning and new technology implementation in hospitals. ***Administrative Science Quarterly,*** 46: 685–716.

Eisenhardt, K. M. 1988. Agency theory: An assessment and review. ***Academy of Management Review,*** 14: 57–74.

Eisenhardt, K. M. 1989. Making fast strategic decisions in high-velocity environments. ***Academy of Management Journal,*** 32: 543–576.

Elstein, A. S. 2001. Naturalistic decision making and clinical judgment. ***Journal of Behavioral Decision Making,*** 14: 363–365.

Elstein, A. S., & Schwarz, A. 2002. Evidence base of clinical diagnosis: Clinical problem solving and diagnostic decision making: Selective review of the cognitive literature. ***British Medical Journal,*** 324: 729–732.

Elstein, A. S., Shulman, L. S., & Sprafka, S. A. 1978. ***Medical problem solving: An analysis of clinical reasoning.*** Cambridge, MA: Harvard University Press.

Fischhoff, B., & Beyth-Marom, R. 1983. Hypothesis evaluation from a Bayesian perspective. ***Psychological Review,*** 90: 239–260.

Fleming, L., & Sorenson, O. 2004. Science as a map in technological search. ***Strategic Management Journal,*** 25: 909–928.

Forrester, J. W. 1961. ***Industrial dynamics.*** Portland, OR: Productivity.

Gaba, D. M. 1989. Human error in anesthetic mishaps. ***International Anesthesiology Clinics,*** 27(3): 137–147.

Gersick, C. J. G., & Hackman, J. R. 1990. Habitual routines in task-performing groups. ***Organizational Behavior and Human Decision Processes,*** 47: 65–97.

Gittell, J. H. 2009. ***High performance healthcare: Using the power of relationships ot achieve quality, efficiency and resilience.*** New York: McGraw-Hill.

Glaser, B. G., & Strauss, A. L. 1967. ***The discovery of grounded theory: Strategies for qualitative research.*** New York: Aldine.

Gonzales, C., Lerch, J. F., & Lebiere, C. 2003. Instance-based learning in dynamic decision making. ***Cognitive Science,*** 27: 591–635.

Gupta, A. K., Smith, K. G., & Shalley, C. E. 2006. The interplay between exploration and exploitation. ***Academy of Management Journal,*** 49: 683–706.

Hirt, E. R., & Markman, K. 1995. Multiple explanation: A consider-an-alternative method for debiasing judgments. ***Journal of Personality and Social Psychology,*** 69: 1069–1086.

Hodgkinson, G. P., & Healey, M. P. 2008. Cognition in organizations. ***Annual Review of Psychology,*** 59: 387–417.

Johnson, P. E., Hassenbrock, F., Duran, A. S., & Moller, J. H. 1982. Multimethod study of clinical judgment. ***Organizational Behavior and Human Performance,*** 30: 201–230.

Johnson, P. E., Moen, J. B., & Thompson, W. B. 1988. Garden path errors in diagnostic reasoning. In L. Bolc & M. J. Coombs (Eds.), ***Expert system applications:*** 395–427. Berlin: Springer-Verlag.

Jonas, E., Schulz-Hardt, S., Frey, D., & Thelen, N. 2001. Confirmation bias in sequential information search after preliminary decisions: An expansion of dissonance theoretical research on selective exposure to information. ***Journal of Personality and Social Psychology,*** 80: 557–571.

Josephson, J. R., & Josephson, S. G. (Eds.). 1994. ***Abductive inference: Computation, philosophy, technology.*** New York: Cambridge University Press.

Kahneman, D., & Tversky, A. 1979. Prospect theory: An analysis of decision under risk. ***Econometrica,*** 47: 263–291.

Klein, G. 1998. ***Sources of power.*** Cambridge, MA: MIT Press.

Klein, G., Phillips, J. K., Rall, E., & Peluso, D. A. 2006. A data/frame theory of sensemaking. In R. Hoffman (Ed.), ***Expertise out of context.*** Mahwah, NJ: Lawrence Erlbaum Associates.

Klein, G., Pliske, R., Crandall, B., & Woods, D. 2005. Problem detection. ***Cognition, Technology and Work,*** 7: 14–28.

Klein, G. A., Orasanu, J., Calderwood, R., & Zsambok, C. E. 1993. ***Decision making in action.*** Norwood, NJ: Ablex.

Kleinmuntz, D. N. 1985. Cognitive heuristics and feedback in a dynamic decision environment. ***Management Science,*** 31: 680–702.

Koehler, D. J. 1991. Explanation, imagination, and confidence in judgment. ***Psychological Bulletin,*** 110: 499–519.

Langley, A., Mintzberg, H., Pitcher, P., Posada, E., & Saint-Macary, J. 1995. Opening up decision making: The view from the black stool. ***Organization Science,*** 6: 260–279.

Levitt, B., & March, J. G. 1988. Organizational learning. ***Annual Review of Sociology,*** 14: 319–340.

Logan, G. D. 1988. Toward an instance theory of automatization. ***Psychological Review,*** 95: 492–527.

Louis, M. R., & Sutton, R. I. 1991. Switching cognitive gears: From habits of mind to active thinking. ***Human Relations,*** 44: 55–76.

Mandler, G. 1982. Stress and thought processes. In S. Goldberger & S. Breznitz (Eds.), ***Handbook of stress:*** 88–164. New York: Free Press.

Mandler, G. 1984. ***Mind and body.*** New York: Norton.

March, J. G. 1991. Exploration and exploitation in organizational learning. ***Organization Science,*** 2: 71–87.

March, J. G. 1994. ***A primer on decision making: How decisions happen.*** New York: Free Press.

March, J. G. 1997. Understanding how decisions happen in organizations. In Z. Shapira (Ed.), ***Organizational decision making:*** 9–32. Cambridge: Cambridge University Press.

Marcus, A. A., & Nichols, M. 1999. On the edge: Heeding the warning of unusual events. ***Organization Science,*** 10: 482–499.

Neisser, U. 1976. ***Cognition and reality: Principles and implications of cognitive psychology.*** San Francisco: Freeman.

Payne, J. W., Bettman, J. R., & Johnson, E. J. 1992. Behavioral decision research: A constructive processing perspective. ***Annual Review of Psychology,*** 43: 87–131.

Peirce, C. S. 1958. ***Collected papers.*** Cambridge, MA: Harvard University Press.

Perlow, L., Okhuysen, G., & Repenning, N. P. 2002. The speed trap: Exploring the relationship between decision making and temporal context. ***Academy of Management Journal,*** 5: 931–955.

Rasmussen, J., Pejtersen, A. M., & Goodstein, L. P. 1994. ***Cognitive systems engineering.*** New York: Wiley.

Repenning, N. P., & Sterman, J. D. 2002. Capability traps and self-confirming attribution errors in the dynamics of process improvement. ***Administrative Science Quarterly,*** 47: 265–295.

Roberts, K. H. 1990. Some characteristics of one type of high reliability organization. ***Organization Science,*** 1: 160–176.

Rudolph, J. W. 2003. ***Into the big muddy and out again: Error persistence and crisis management in the operating room.*** Chestnut Hill, MA: Boston College.

Rudolph, J. W., & Raemer, D. B. 2004. Diagnostic problem solving during simulated crises in the OR. ***Anesthesia and Analgesia,*** 98(5S): S34.

Rudolph, J. W., & Repenning, N. P. 2002. Disaster dynamics: Understanding the role of quantity in organizational collapse. ***Administrative Science Quarterly,*** 47: 1–30.

Sastry, M. A. 1997. Problems and paradoxes in a model of punctuated organizational change. ***Administrative Science Quarterly,*** 42: 237–275.

Simon, H. 1957. ***Models of man: Social and rational.*** New York: Wiley.

Smith, S. M., & Blankenship, S. E. 1991. Incubation and the persistence of fixation in problem solving. ***American Journal of Psychology,*** 104: 61–87.

Snook, S. A. 2000. ***Friendly fire: The accidental shootdown of U.S. Black Hawks over Northern Iraq.*** Princeton, NJ: Princeton University Press.

Starbuck, W. H., Greve, A., & Hedberg, B. L. T. 1978. Responding to crises. ***Journal of Business Administration,*** 9: 111–137.

Staw, B. M. 1976. Knee-deep in the big muddy: A study of escalating commitment to a chosen course of action. ***Organizational Behavior and Human Performance,*** 16: 27–44.

Sterman, J. 2000. ***Business dynamics.*** Chicago: Irwin-McGraw Hill.

Sterman, J. D. 1994. Learning in and about complex systems. ***System Dynamics Review,*** 10: 291–330.

Strauss, A., & Corbin, J. 1994. Grounded theory methodology: An overview. In N. K. Denzin & Y. S. Lincoln (Eds.), ***Handbook of qualitative research:*** 273–285. Thousand Oaks, CA: Sage.

Suddaby, R. 2006. From the editors: What grounded theory is not. ***Academy of Management Journal,*** 49: 633–642.

Sutcliffe, K., & Weber, K. 2003. The high cost of accurate knowledge. ***Harvard Business Review,*** 81(5): 74–87.

Torbert, W. R., & Taylor, S. S. 2008. Action inquiry: Interweaving multiple qualities of attention for timely action. In P. Reason & H. Bradbury (Eds.), ***The Sage handbook of action research:*** 239–251. London: Sage.

Tversky, A., & Kahneman, D. 1974. Judgement under uncertainty: Heuristics and biases. ***Science,*** 185: 1124–1131.

Weick, K. E. 1993. The collapse of sensemaking in organizations: The Mann Gulch disaster. ***Administrative Science Quarterly,*** 38: 628–652.

Weick, K. E. 1995. ***Sensemaking in organizations.*** Thousand Oaks, CA: Sage.

Weick, K. E., Sutcliffe, K., & Obstfeld, D. 2005. Organizing and the process of sensemaking. ***Organization Science,*** 16: 409–421.

Weick, K. E., Sutcliffe, K. M., & Obstfeld, D. 1999. Organizing for high reliability: Processes of collective mindfulness. ***Research in Organizational Behavior,*** 21: 81–123.

Winter, S. G. 1971. Satisficing, selection, and the innovating remnant. ***Quarterly Journal of Economics,*** 85: 237–261.

Xiao, Y., & MacKenzie, C. F. 1995. ***Decision making in dynamic environments: Fixation errors and their causes.*** Paper presented at the Human Factors and Ergonomics Society 39th annual meeting, Santa Monica, CA.

Zsambok, C. E., & Klein, G. 1997. ***Naturalistic decision making.*** Mahwah, NJ: Lawrence Erlbaum Associates.

Jenny W. Rudolph (jwrudolph@partners.org) is assistant clinical professor of anesthesia at Harvard Medical School and assistant in research in the department of Anesthesia and Critical Care at Massachusetts General Hospital. She is associate director of the Institute for Medical Simulation at the Center for Medical Simulation. She received her Ph.D. in management from Boston College. Her research focuses on the role of communication and cognition in dynamic problem solving and crisis management, primarily in health care.

J. Bradley Morrison (bmorriso@brandeis.edu) is assistant professor of management at Brandeis International Business School, Brandeis University. He received his Ph.D. from the MIT Sloan School of Management. His research focuses on how people and organizations cope with dynamic complexity. Current interests include problems in business process implementation, learning, and organizational change.

John S. Carroll (jcarroll@mit.edu) is the Morris A. Adelman Professor of Management at the MIT Sloan School of Management and the MIT Engineering Systems Division. He received his Ph.D. from Harvard University. His research focuses on decision making and learning in organizational contexts, especially in high-hazard industries such as nuclear power and health care.