

CHAPTER 3

Evidence in Practice: A Framing of the Terrain

JAMES P. SPILLANE AND DAVID B. MIELE

While much of the recent educational literature has been devoted to explaining how investigators can produce high quality, practical research evidence (e.g., Cook, 2002; Feuer, Towne, & Shavelson, 2002; Shavelson & Towne, 2002; Slavin, 2002; Towne, Wise, & Winters, 2005), little attention has been paid to how evidence can and should be used by teachers and school leaders. Our goal is *not* to review the empirical literature on teachers' and school leaders' use of evidence, but rather to identify the conceptual tools that frame our thinking about this work. Policymakers often work on the assumption that evidence-based practice should be a simple and straightforward process for school practitioners; that is, practitioners need only follow the guidance offered by evidence—typically equated with qualitative research findings and trends in student achievement data—when deciding what they should do and how they should do it. However, this belief is based on several questionable assumptions.

First, it is sometimes assumed that the chief source of evidence practitioners should consider when making decisions is social science research (Hammersley, 2001). But there are a variety of formal (e.g., standardized test scores) and informal (e.g., personal experience) sources of information that also contribute to the decision making process. Second, it is sometimes assumed that judgments about the most appropriate course of action in a given educational context can be inferred from relevant research findings. However, research findings merely inform practitioners about what the *general* outcomes are of different kinds of decisions. They *do not* answer questions about what the social value of these outcomes is or should be. Such questions are philosophical and political in nature, and thus should be addressed as

James P. Spillane is the Spencer T. and Ann W. Olin Professor in Learning and Organizational Change at Northwestern University's School of Education and Social Policy, Kellogg School of Management and Institute for Policy Research. David B. Miele is a doctoral student in social psychology at Northwestern University.

part of a separate (although related) process of deliberation (Hammond, Harvey, & Hastie, 1992; Phillips, 2007). Nor do such research findings address what the *specific* outcomes will be when these decisions are made within a particular context. Thus, research findings must also be considered *in situ*; that is, practitioners should interpret new research in light of local beliefs, knowledge, values, and problems. Third, it is assumed that the various research findings that are relevant to a particular educational decision will generally point to a single and clear conclusion and, therefore, it should be relatively easy to determine the effectiveness of a particular intervention or educational approach. However, social science findings are usually contested by other scholars in the field and thus are open to multiple interpretations. Consequently, the job of weighing competing interpretations, combining multiple sources of evidence (both formal and informal), and determining the most likely outcome of a particular intervention or approach is left to school practitioners themselves.

Some readers might argue that the state of social science research in education is such that it offers very weak evidence with respect to what actually works when it comes to educating children. The idea is that an increase in the production of quality educational research (i.e., research that established causal relationships between practices and outcomes) will lead practitioners to use social science research more often and to better effect. Although we agree that improving the quality of educational research and working to make it more accessible among local practitioners would likely increase its use, we urge caution because these proposals are often based on rather simplistic and hyper-rational notions about relations between social science research and decision making in schools and school districts. If the goal is for educational practice to make better use of social science research, we believe that it is essential to understand the processes by which practitioners select, evaluate, and utilize information from a variety of sources (including from their own situation) to *construct* evidence.

In our view, evidence use is first and foremost an issue of evidence construction. But research does not construct evidence—people do. As Phillips (2007) points out, evidence is not a synonym for information, facts, or even data for that matter—evidence “is information selected from the available stock and introduced at a specific point in the argument in order to persuade a particular audience of the truth or falsity of a statement. . . .” (Majone, 1989, p. 11). In other words, the implications that a set of facts holds for a particular problem are never self-evident. It is only when we *interpret* these facts as confirming or

contradicting a proposed definition or solution to the problem that they exist as evidence. Thus, in terms of evidence construction, the problem comes first. That is, we begin by observing or noticing something askew in our environment—a puzzle or problem of some sort (e.g., declining mathematics achievement in the middle grades). We then work at defining or clarifying the problem, either by drawing from our previous experience (e.g., as teachers) or by collecting preliminary data. It is only once the problem has been firmly established and potential solutions (or hypotheses)¹ have been proposed that we make more focused attempts to collect data. Finally, this data gets interpreted in terms of the support (or opposition) it provides for a particular solution.

School leaders and teachers use many different sorts of information to construct evidence, including published research reports, word of mouth accounts of what works, student achievement scores, demographic data, and personal experience. These sources are central to the work practices of teachers and administrators to the extent that they are used (alone or in combination) to justify or criticize different ways of doing things in the classroom, department, or school. That being said, what we intend to focus on in this chapter is how information (regardless of source) can be interpreted as evidence for different and, in some cases, competing conclusions about how to resolve a particular problem. Contrary to many people's intuitions, a single piece of information (e.g., "test scores are down") does not entail a single conclusion (e.g., "classroom instruction is poor"). This is in part because information is always interpreted with respect to a person or organization's existing beliefs, values, and norms (Cronbach et al., 1980). These constructs serve as the lens through which new information is understood and thus influence how information gets shaped into evidence (Coburn, 2006; Phillips, 2007).

From a sense-making perspective, *what* is noticed in a school environment, *whether* this information is understood as evidence pertaining to some problem, and *how* it is eventually used in practice (perhaps to formulate a solution to the problem) depends on the cognitions of individuals operating within that school (e.g., teachers and administrators). New information—such as empirical research on how to teach reading to first graders—cannot be deposited directly into the minds of those we trust with educating our children. Practitioners must notice, frame, and interpret new information before they can put it into practice as evidence.

But sense making is not entirely an individual or solo affair. What we notice and what sense we make out of this information depends

on our situation. Of course, “situation” or “context” have been used as catchall terms for everything from the class background of a student population to the teaching experience of a school’s staff. Moreover, researchers working in different academic traditions treat context differently. Some focus on how individuals make sense of new information in their environment. Others complicate things considerably by adopting a situated or distributed cognition perspective, arguing that “situation” is not simply a backdrop for sense making but a defining or constituting element of human practice. Based on this latter perspective, we argue for attention to the *practice* of sense making, viewing it as *distributed* across an interactive web of actors and key aspects of their situation—including tools and organizational routines (Greeno, 1998).

We begin the chapter by considering sense making from an individual perspective. In the second section, we move beyond an individual level of analysis and consider the role of the situation in the construction of evidence. Specifically, we argue for attention to work practices as a way of thinking about how the situation might influence what gets constructed as evidence.

Constructing Evidence: The Individual as Sense Maker

Although the terms “interpretation” and “sense making” are often used interchangeably, sense making actually refers to a set of cognitive processes that encompass interpretation. From a sense-making perspective, stimuli must be selected from the environment *before* they can be interpreted (Weick, 1995). Thus, we begin by considering how people attend to objects and events in their environment and then go on to examine how they interpret this newly acquired information as evidence for or against a particular set of beliefs (see also Spillane, Reiser, & Reimer, 2002). Next, we take a closer look at the kinds of knowledge representations (i.e., schemas and mental models) that play a role in both processes. We also examine the various outcomes (i.e., assimilation versus accommodation) of attention and interpretation as well as how these outcomes are moderated by expertise. We conclude by examining various forms of bias that, in contrast to expertise, lead to undesirable sense-making outcomes.

Attention

Inundated with stimuli from our environment, we tend to notice things that are relevant to our goals and expectations (i.e., the signal) and ignore things that are not (i.e., the noise). Through this process of *selective attention* we are able to maintain a certain level of cognitive

efficiency. Further, what gets noticed and singled out as relevant to our goals depends on the mental representations we have abstracted from our experience. Sense making involves filtering stimuli through the lens of our existing knowledge framework (Starbuck & Milliken, 1988). In terms of evidence use, this means that whether or not a stimulus gets noticed and classified as evidence depends in part on whether the practitioner possesses a set of mental representations that make the stimulus salient. For example, whether or not a teacher notices that a student is confused (especially if the student is attempting to hide his or her confusion) may depend on whether the teacher's representation of a confused student includes nonverbal indicators of uncertainty (Webb, Diana, Luft, Brooks, & Brennan, 1997). It should be noted that the process of selective attention does not usually operate at the level of conscious awareness. We are often unaware of having focused our attention on a particular feature of the environment; as a result, we sometimes mistakenly assume that other people have reached the same conclusion after being exposed to similar stimuli. For example, an expert teacher who assumes that a novice teacher has noticed a student's confused facial expression may find it difficult to understand why the novice does not provide that student with extra help. Conversely, the novice teacher may not understand why the expert is making a fuss over the student, because according to the information she possesses (e.g., an explicit denial of uncertainty by the student), the student does not require extra help. Thus, it is important to keep in mind that what may seem like a disagreement over interpretation may actually have resulted from differences in the information considered by each party (i.e., difference in attention). Of course, there are plenty of instances in which different conclusions are drawn from the same information. To understand how these instances arise, it is necessary to consider the process of interpretation in more detail.

Interpretation

New information (e.g., a widening achievement gap in middle-grade mathematics) is always understood in light of what is already known (Brewer & Nakamura, 1984; Greeno, Collins, & Resnick, 1996). Prior knowledge, including the tacit expectations and beliefs we abstract from our experience, influence not only which stimuli we notice, but also how these stimuli are encoded, organized, and interpreted. Thus, it can be argued that local consumers of information construct interpretations of new stimuli based on their previous interactions with the environment. In a sense, it is the consumer's interpretations of previously

encountered stimuli (and the inferences about the future that these interpretations afford) that allow him or her to construe new information as “evidence” for or against a particular conclusion. Practitioners continually make use of “personal data” (i.e., information drawn from personal experience) to make sense of things.

Schemas

If every detail of our past experience was stored in memory and used to make sense of new stimuli, the process of interpretation would be extremely slow and inefficient. For this reason, we selectively encode useful information about the concrete objects and events we regularly encounter in the world (e.g., students, homework, grading, etc.), as well as information about the abstract relationships between these entities (e.g., fairness), into knowledge representations known as schemas (Mandler, 1984; Rumelhart, 1980; Schank & Abelson, 1977). Schemas are not simply collections of loosely associated features; rather, they represent beliefs or “theories” about how these features relate to one another (Markus & Zajonc, 1985; Murphy & Medin, 1985). Also, they are not limited to information about the physical world—in fact, schemas are often used to generate expectations about social relationships (Cantor & Mischel, 1979; Trope, 1986), such as how students and teachers should interact in a classroom context.

Schemas are *automatically* activated as we attend to new stimuli in our environment (Higgins, 1996). Once activated, a schema can guide the sense-making process in at least three ways. First, by providing additional information about the objects and events that constitute our experience, schemas allow individuals to *fill out* their initial understanding of what was said or observed (Brewer & Nakamura, 1984; Rumelhart & Ortony, 1977; Schank & Abelson, 1977). Second, when this initial understanding is vague or ambiguous, schemas provide *clarity* by allowing individuals to infer the existence of objects and events that may not have been directly experienced (i.e., missing details). And, third, by allowing individuals to selectively attend to signals and cues in the environment that are relevant to their processing goals, schemas enable individuals to *focus* their sense-making efforts in a cognitively efficient manner.

Mental Models

Research on sense making has also focused on how dynamic processes (e.g., children’s learning of arithmetic) are encoded into sophisticated knowledge representations known as mental models

(Gentner & Stevens, 1983). Mental models are more elaborate than schemas to the extent that they represent the causal relationships between objects and events. Whereas schemas are used to represent our beliefs about what things are (i.e., their defining features), mental models are used to represent our beliefs about how things work. Once a general model of some process has been constructed, it can be instantiated (or “run”) with information about a specific situation in order to make predictions about what might happen in the future. For example, a teacher’s model of student motivation can be run with specific information about a child’s behavior (“exhibits little intrinsic motivation”) to make predictions about how long the child will persist at a given task (“not very long”).

Mental models can be tacitly constructed from experience, without reference to formal explanations of the phenomenon (Greeno, 1989). That is, people often develop an implicit understanding of how things work simply by interacting with objects in their environment. This is not to say that our mental models are unaffected by other people’s explanations (Vosniadou & Brewer, 1992). However, it is important to keep in mind that a person’s inability to explain how something works does not necessarily mean that he or she lacks a model for understanding it. In some cases, a person may even possess two mental models of the same phenomenon: an “espoused” model, which the person uses to explain the phenomenon to other people, and an “in-use” model, which guides the person’s behavior when responding to the phenomenon directly. Teachers, for example, often activate an espoused model of how children learn when talking with each other, but activate a separate, in-use model when actually teaching (Strauss, 1993).

Often, teachers do not even realize that they possess two models of the same phenomenon. What this suggests is that information is unintentionally construed in different ways depending on the activity or work practice in which the practitioner is engaged—a child’s reluctance to participate in class may be construed as evidence of low self-efficacy when the practitioner is talking about the child with a colleague (espoused model), but as evidence of low intelligence when she is actually interacting with the child in the classroom (in-use model). The challenge then, in terms of evidence use, is to ensure that the sophisticated understanding some practitioners develop at an explicit level (e.g., through in-service training) gets translated into practice at an implicit level. Evidence is of little practical benefit if it is not incorporated into practitioners’ in-use models of how things work within an educational context.

Accommodation Versus Assimilation

The process of interpreting new information is not always as straightforward as applying an existing schema or mental model; in some cases, it may require constructing a new knowledge representation or reorganizing a collection of existing representations. Although “assimilation,” or the encoding of stimuli into existing knowledge representations, often serves as our default mode of interpretation, early theories of cognitive development (Piaget, 1972) stressed the importance of “accommodation,” or the restructuring of existing knowledge to account for new concepts and ideas. Whereas assimilation is a conserving process that strives to “make the unfamiliar familiar, to reduce the new to the old” (Flavell, 1963, p. 50), accommodation is a modifying process that strives to maintain the novelty of the unfamiliar in order to *reconcile* the new with the old.

More recent theories of cognitive development have focused on the difficulty we face when attempting to restructure deeply entrenched representations (Carey, 1985). According to these theories, successful accommodation requires that we repeatedly engage and reconsider information that is difficult to reconcile with what we already know. Failure to do so increases our risk of interpreting new pieces of information as minor variations on an old theme, rather than as critically important contributions to our understanding of an issue.

To a large extent, whether new information is assimilated or accommodated depends on its similarity to our existing knowledge representations (Gentner, Rattermann, & Forbus, 1993; Ross, 1987). When the information is only superficially similar to the schemas and mental models that are activated, it tends to get assimilated. But when the information is deeply analogous to these representations (i.e., in terms of both its content *and* its organization), it is more likely to be accommodated.

In many cases, superficially similar representations are activated because they are all that are available (i.e., because deeper alternatives have not yet been constructed). For example, Hill (2001) found that teachers working on a district committee to adopt materials that would support the state’s mathematics policy understood the state reform ideas in different ways than policymakers had intended. These teachers, perceiving little difference between their own position and the state’s, assumed that a traditional curriculum was sufficient for implementing great chunks of a state policy that was designed to initiate fundamental change in mathematics education. What is striking is that these

committee members devoted substantial time to interpreting the state's mathematics policy, which suggests that their "misunderstandings" cannot be explained in terms of limited attention to the policy's directives. Presumably, teachers interpreted this "new" way of teaching mathematics to be similar (in some superficial sense) to what they already understood about how to teach math and, thus, assimilated the content of the policy to their existing knowledge framework. One could imagine a similar scenario in which the teachers struggle to accommodate new research findings on mathematics learning because they are distracted by the superficial similarities between the findings and their knowledge representations.

Expertise

One way that practitioners can avoid the pitfalls of assimilation is to develop elaborate knowledge representations that enable them to move beyond superficial similarities. Research shows that the complexity of our representations in a particular domain corresponds to our level of expertise. With increasing expertise, we construct knowledge representations that encompass a broader range of experiences and that exhibit deeper levels of organization. These representations then allow us to perceive sophisticated patterns in information that may not be apparent to novices (Chase & Simon, 1973; Larkin, McDermott, Simon, & Simon, 1980; VanLehn, 1989). As experts, we are less likely than novices to be distracted by the superficial similarities that initially capture our attention; that is, we are less likely to "lose the forest for the trees" (Novick, 1988).

What this suggests is that practitioners and policymakers with varying levels of expertise in a particular domain (e.g., mathematics pedagogy) are likely to develop different (and possibly conflicting) interpretations of the evidence offered in support of a proposed reform. For example, only an expert math teacher may be able to distinguish between a reform scenario in which manipulatives serve as a basis for exploration and a more traditional scenario in which manipulatives serve as a procedural training tool (Cohen, 1990). If a novice teacher does not see differences where the expert does, how can we expect practitioners as a group to develop a coherent interpretation of the evidence for or against proposed reforms? From a sense-making perspective, the problem leading to disagreement is not that each party brings its own set of evidence to the table, but that each party interprets the same evidence in fundamentally different ways. Therefore, an obvious but important step that can be taken to prevent irresolvable

disagreements about how new or existing evidence should be interpreted is to encourage the parties involved to reflect on and share with each other the beliefs and assumptions that inform their understanding of the issue (Ross & Ward, 1995).

Bias

Although the same information can be construed as evidence in very different ways depending on a person's existing knowledge representations and level of expertise, there are normative standards of evidence which determine whether a particular interpretation should be accepted or dismissed as biased within a given context (MacCoun, 1998). Such standards include: (1) whether the interpretation references information that is irrelevant, inappropriate, or unjust (e.g., some scholars have argued that IQ is irrelevant to the assessment of learning disabilities because intelligence and achievement are really separate constructs; see Siegel, 1989); (2) whether the interpretation fails to take into account relevant information that may lead to a different conclusion (e.g., a claim that the No Child Left Behind Act has been successful may fail to account for the lack of equivalent testing standards across states; see Linn, Baker, & Betebenner, 2002); and (3) whether the interpretation conflicts with an equally plausible explanation of existing information (e.g., one group may interpret low standardized test scores as indicative of academic failure, while another group may interpret the same scores as indicative of incompatible assessment; see Shepard, 2000).

An additional challenge to understanding how information is constructed as evidence, and how this evidence is then evaluated and used by practitioners, stems from the fact that these standards of evidence often vary from person to person (e.g., some people may endorse all three of the above standards, others may endorse only one). Furthermore, a particular person's standards may not always be applied correctly (or fairly), even when it is that person's intention to be as objective or impartial as possible. This is because the process of weighing and interpreting information is vulnerable to numerous motivational and cognitive biases, many of which operate outside of the practitioner's awareness.

Many readers will be familiar with intentional biases such as fraud (the intentional effort to distort, fabricate, or conceal information for personal gain) and advocacy (the selective use or emphasis of information in support of one's position). Although the "rational" model of policymaking frequently ignores the biasing effects of advocacy and

rhetoric (Majone, 1989), they are part and parcel of decision-making processes from the statehouse to the schoolhouse.

Biases that operate outside of conscious awareness are also prevalent. This is because practitioners do not activate all of their schemas and mental models when attempting to interpret new stimuli. Instead, they selectively activate only those representations that will help them fulfill their current processing goal (Kunda, 1990). In many cases, this goal is to arrive at an accurate conclusion (e.g., to determine whether a new intervention is truly effective). However, in other cases, the goal may be to arrive at a preferred or desired conclusion (e.g., to determine that a new intervention is effective, as hoped). There are a number of fundamental motives that tend to give rise to these “directional” goals, including the need for self-esteem, the need for cognitive consistency, and the need for belonging. Whenever information processing serves one of these needs, sense making is considered to be “hot” and, thus, susceptible to “motivated directional biases” (Kunda, 1990).²

Processes of interpretation are particularly susceptible to a form of hot bias known as *biased assimilation*, which occurs when evidence that is offered in support of a desired conclusion is automatically perceived to be of higher quality than evidence that supposedly challenges this conclusion. For example, people are likely to rate research studies that corroborate their own position to be superior in quality and persuasiveness to studies that contradict their position, *irrespective of the research methodologies that the studies supposedly employed* (Lord, Ross, & Lepper, 1979). This effect results from our tendency to scrutinize contradictory evidence more than we scrutinize supporting evidence (Ditto & Lopez, 1992). While we are often willing to accept preference-consistent information “at face value,” we are more likely to scrutinize information that does not fit with our expectations or desires.

Even when people are *not* motivated to reach a particular conclusion, they are still capable of interpreting information in a biased manner—i.e., they are susceptible to “cold” biases (Aarkes, 1991; Wilson & Brekke, 1994). One type of cold bias that deserves special consideration is *confirmation bias*, which occurs when a hypothesis is more likely to be confirmed than disconfirmed, regardless of its validity. Confirmation bias results from the use of a positive test strategy to evaluate hypotheses (Klayman & Ha, 1987). With a positive test strategy, we “tend to test those cases that have the best chance of verifying current beliefs rather than those that have the best chance of falsifying them” (p. 211). In other words, we are more likely to confirm a hypothesis when we search for supporting evidence without making an effort to determine whether

opposing evidence exists.³ Research suggests that local administrators and teachers frequently employ a positive test strategy when evaluating evidence (Birkeland, Murphy-Graham, & Weiss, 2005; Coburn, Toure, & Yamashita, 2006; Coburn & Talbert, 2006), although not always to detrimental effect.

In sum, individual sense making involves attention and interpretation. Both *what we notice* and *how we interpret what we notice* depends, to a large extent, on a person's existing schemas and mental models. When new information is made to conform to existing representations, assimilation occurs; but, when new information is used to restructure existing representations, accommodation occurs. Because experts are more likely than novices to perceive complex patterns in new information, they are less likely to be distracted by superficial similarities that exist between the information and their knowledge representations. But even experts are susceptible to bias. While some biases are intentional (e.g., fraud and advocacy), others operate outside of conscious awareness.

Constructing Evidence: Sense Making and Situation

Sense making is not a solo affair—how evidence gets constructed depends on the context or “situation” in which people notice and interpret information. Studying the role of the situation in the construction and use of evidence can be quite challenging, for two reasons. First, the situation of an organization (e.g., a school) is broad and multifaceted. Numerous situational factors—including the organization's size, structure, and social capital, the diversity of its staff in terms of race, career stage, and life stage, and the degree to which its core work is defined—are believed to influence how people work within an organization (Bryk & Schneider, 2002; Galbraith, 1973). Thus, it is often difficult to decide which factors matter most when it comes to evidence use.

Second, the causal links between situation, behavior, and mind are far from clear. Some social scientists, especially those operating from a cognitive science perspective, tend to focus on how mental processes influence our behavior (including our social interactions). Others, especially those operating from a sociological perspective, focus on how the structure and dynamics of our social situation shape what we do. For over half a century, these competing foci of the social sciences—individual cognition versus the social situation—have been at the center of a fundamental debate concerning the nature of human agency. While

some argue that human agency is ultimately determined by the structure of the situation (Althusser, 1971), others point out that the situation is itself constructed from the beliefs, intentions, and actions of individuals (Berger & Luckmann, 1966). Thus, when studying the role of the situation in evidence use, deciding whether to treat the situation as a cause or as an effect of individual sense-making efforts can be difficult.

In this section, we argue for an approach to the situation that attempts to address both of these challenges. By focusing on the work practices that exist within schools and school systems, our approach reveals the reciprocal influence of individual agency and social structure and provides a criterion for deciding which situational variables matter when it comes to studying evidence use.

Critiquing the structure–agency dualism evident in many studies of the situation, some scholars have proposed alternative conceptions of agency and structure (Giddens, 1979, 1984). For example, Giddens argues that structure is both the medium for human activity and its outcome. Structure constitutes human action, providing the rules and resources that guide our behavior; but, at the same time, structure is also created, reproduced, and potentially transformed by human action—it is our behavior that defines the rules. The challenge still remains, however, as to which aspects of the situation or social structure are important to study. One way of dealing with this is to argue that situational factors only matter to the extent that they pertain to actual work practices; the situational factors that matter are those that enable (or constrain) human activity, that are “instantiated in activity,” or that serve as rules of conduct or “rights to resources” (Whittington, 1992, p. 696). Following Whittington, we believe that a productive means of exploring relations between the situation and human agency in evidence use involves attending to work practice.

Work Practice

Not surprisingly, anchoring investigations of evidence use in work practice involves examining day-to-day practice in schools. Such practice, familiar to most, includes monitoring instruction, attending grade-level meetings, setting instructional priorities, and so on. It is in this kind of practice that new information is encountered (or overlooked), interpreted as evidence for one thing or another (or dismissed as irrelevant), and eventually put to use (or ignored).

Practice is one of those frequently used words that is believed to represent a fundamental ingredient or lever for improving schools;

however, more often than not, the word practice is glossed over and its meaning taken for granted. For instance, some researchers talk generally about “best practices”—the strategies or activities that school staff should engage in to be successful. However, practice actually refers to the complicated pattern of behavior that emerges from people’s interactions with each other and with their social situation *over time* (Bourdieu, 1981). That is, the study of work practice involves more than identifying easily followed strategies that can be implemented at any point in time. To ignore the dynamics of the situation is to disconnect behavior from the “urgency of practice” (Bourdieu, 1990). Something happens, people must react, but they do so in relation to others—and it is in these interactions that practice takes form. Once practice actually unfolds, the best laid plans and well-honed strategies often turn out very different from what originally was expected.

Although practice unfolds in the here and now, it is irrevocably tied to the past. People who are engaged in practice draw on a logic that is informed by their past interactions—a logic that “is able to organize all thoughts, perceptions, and actions by means of a few generative principles” (Bourdieu, 1990, p. 86). Consider, for example, practitioners at Costen School (Hallett, in press). A new principal worked to comply with the school district’s mandate for improved student achievement and increased accountability. In doing so, she implemented leadership routines (such as actively monitoring classroom instruction) that were drawn from her prior experience and from her principal preparation program. These routines were not implemented in a vacuum but in relation to other staff members, and it is in these interactions that a new pattern of work practice eventually began to unfold at Costen. It is important to note that the resulting interactions, especially responses to the principal’s efforts by some veteran teachers, were also guided by past experience. At Costen, teachers were accustomed to having substantial autonomy in their classrooms—the previous school leaders had worked to buffer them from external influences and had allowed them to determine their own instructional approaches. As a result, the principal was not able to institute a completely new way of doing things; she was instead forced to negotiate practice that extended (but still fit with) what the teachers were accustomed to doing.

What this example illustrates is that simply extracting actions or strategies from their place and time is insufficient for understanding work practice. Only by examining the conflicting expectations of the principal and teachers is it possible to understand how this particular pattern of work practice eventually emerged. The key to understanding

practice is to understand how it arises from people's ongoing attempts to negotiate their relationship with their situation—social, material, cultural, and historical.

A distributed perspective on work practice. Frameworks for studying work practice are scarce, and those that do exist tend to dwell on either individual agency or the determining influence of social structure. This has led some scholars to argue that investigations of work practice require the development of new conceptual frameworks, “frameworks built out of concepts that speak directly to practice” (Pickering, 1992, p. 7). Drawing from work in the fields of organizational theory (Weick, 1979), distributed cognition (Greeno, 1998; Hutchins, 1995; Latour, 1987; Leont’ev, 1981), and activity theory (Cole & Engeström, 1993; Engeström & Middleton, 1998), we argue for taking a distributed perspective to studying work practice (Spillane, 2006).

Some psychologists have argued that more attention should be paid to the situation in which sense making occurs, especially to the interactive web of actors that comprise the situation. Consequently, the defining characteristic of research on distributed cognition is its “focus on interactive systems that are larger than the behavior and cognitive processes of an individual agent” (Greeno, 1998, pp. 5–6). These interactive systems (what we call work practice) are broader than individual cognition to the extent that they are *distributed* across a web of actors, artifacts, and situations. Thus, work practice cannot be grasped simply by looking at the actions or strategies of particular individuals (e.g., school leaders, teachers, district office administrators); instead, close attention must be paid to the *interactions* among individuals, particularly with respect to how these interactions are facilitated or constrained by specific aspects of the situation.

The complexity of social practice is perhaps better appreciated with an analogy to a simple dance like the two-step. While the actions of partner one and partner two are important in their own right, the practice of the two-step is really about the interplay *between* partners. That is, a simple description of each partner's individual actions fails to capture the practice of the two-step since it does not account for the partners' interactions. We might go so far as to argue that the practice of the two-step is about the interplay between the partners and the music. This is not meant to undermine the contribution of each individual (having a partner with two left feet certainly influences how the dance is performed), but merely to highlight how the practice of the two-step unfolds in its interactions. The same holds true for work

practice in school—it unfolds in the interactions among staff that are mediated by aspects of the situation (e.g., plans, policy directives, organizational routines) (Spillane & Orlina, 2005).

Another important implication of our distributed perspective is that cognitive performance cannot be equated with mental capacity; the success of a given work practice depends on more than the cognitive abilities of the people involved (Resnick, 1991). Language (e.g., referring to grades as “feedback”), rules (e.g., “all curriculum changes must be approved by the principal”), norms (e.g., “do not interrupt another teacher’s class”), tools (e.g., textbooks, computer programs, maps), and organizational routines (e.g., taking attendance, grading homework, faculty meetings) serve as “mediational means” that enable and constrain work practice (Leont’ev, 1981; Vygotsky, 1978; Wertsch, 1991). The situation is a defining or constituting element of work practice.

Organizational Routines and Tools in Work Practice

Work practices in schools are not composed solely of interactions among people; aspects of the situation, such as organizational routines and tools, mediate these interactions (Wertsch, 1991). By organizational routine we mean “a repetitive, recognizable pattern of interdependent actions” (Feldman & Pentland, 2003, p. 96), which in schools includes everything from improvement planning to grade-level meetings. Tools, on the other hand, are defined as externalized representations of ideas that are used by people in their practice (Norman, 1988; Wertsch, 1998). In schools, tools include everything from protocols for evaluating teaching practice to curriculum standards. It is important to note that although routines and tools are constitutive of practice, it is also the case that they are created and recreated in and through practice.

Organizational routines. By offering a rhythm to the workday, week, and month, organizational routines structure much of what happens within schools and school districts. While some organizational routines are part of a school’s formal structure (e.g., school improvement planning), other organizational routines are informal (e.g., the fourth-grade teachers’ coffee klatch) and thus are easily overlooked when attempting to understand local work practices. Formal and informal routines serve similar functions—they allow for efficient, coordinated action, reduce conflict (as they represent an agreement about how to do the work of a particular organization), and provide stability over time. Although

organizational routines are often portrayed as the nemesis of change, innovation, and growth (Freeman & Hannan, 1983), this view has been challenged by scholars who believe that routines actually increase the flexibility of organizations (Feldman, 2000; Feldman & Pentland, 2003; Suchman, 1983).

Organizational routines are critical for understanding evidence use in schools. Through routines such as faculty meetings, department meetings, parent–teacher conferences, and the weekly coffee klatch, teachers and school leaders encounter a range of information, such as dips or gains in student achievement and how well new instructional strategies are working in the classrooms. Furthermore, it is through these routines that teachers and school leaders work out which information counts as evidence, and then decide what this evidence suggests for their practice. For example, at a faculty meeting in which the principal describes new research supporting a proposed curriculum reform, veteran teachers may respond that the studies do not count as evidence because they were conducted at schools that differ from theirs in terms of racial composition, size, and funding. In this case, the school has a normative standard of evidence stipulating that new information only counts as evidence if it takes the school’s local context into account—a standard that is enforced by teachers and administrators through the practice of routines such as faculty meetings.

Organizational routines can be thought of as having both an ostensive aspect and a performative aspect (Feldman & Pentland, 2003; Latour, 1986). The *ostensive aspect* refers to “the ideal or schematic form of a routine” (Feldman & Pentland); it lays out how the routine should unfold within the context of a particular situation. Another way to think of the ostensive aspect of a routine is as a broad script. By script we do not mean something that is followed verbatim, but rather something that suggests who is supposed to do what, when they are supposed to do it, and where it should occur. How the script is actually executed at a particular time and place (regardless of fidelity to the ostensive aspect) is what we mean by the performative aspect of a routine—it is how the routine occurs *in practice*. An important part of a routine’s “performative aspect” is improvisation. The way in which actors spontaneously embellish the script of a routine not only determines how work practice unfolds at that moment, it also determines how practice will unfold in the future. In other words, the ostensive aspect of a routine may be modified and revised as the actors spontaneously add and subtract from a script over time. Once again, this illustrates the reciprocal nature of the relationship between situation and practice—the situation (osten-

sive aspect) shapes practice (performative aspect) and in turn the situation is changed through practice. Both the ostensive and performative aspects of routines are critical to understanding work practices in schools and school districts (Sherer, in press; Spillane, Sherer, & Coldren, 2005).

Tools. Tools are not accessories to interactions or devices that merely allow individuals to increase the efficiency of their work practice. Tools *mediate* interactions among people. By virtue of their affordances and constraints, tools help to structure interactions in a manner that is actually constitutive of practice. Although this added structure limits the ways in which practice can unfold, tools never constrain practice completely—the same tool in different hands can be used in different ways. Thus, it is important to study *how* tools are used in practice (Cole, 1996). Of course, tools can also be studied apart from practice (e.g., to understand their general affordances), but it is only *in* practice that the researcher is able to comprehend how tools *constitute* practice.

The tools used in school work practice include student assessment data, teacher evaluation protocols, curricular frameworks, lesson plans, and student work. But schools can employ different tools for what would otherwise be identical organizational routines. Moreover, these different tools can contribute to the ways in which evidence is constructed about instruction and its improvement. Consider teacher evaluations, a common routine in most schools. Many school districts require school principals to use an observation protocol when evaluating teachers' classroom practice. These observation protocols, which identify dimensions of classroom instruction that the principal should focus on, vary across districts. One district's protocol might focus on a generic teaching process, such as the teacher's use of praise and wait time. Another district's protocol might focus more on the cognitive complexity of the academic tasks that students are working on. These two observation protocol tools focus the interaction between the teacher and the school principal on distinctly different aspects of instruction. As a result, these tools help define classroom monitoring practice in different ways. More important, the information generated (as well as the evidence that might be constructed about practice from this information) is likely to differ. While one protocol is more likely to generate information about the quality of academic work, the other is more likely to generate evidence about teaching processes. Although tools do not straightjacket our interactions, by focusing and framing our interactions in particular ways, they do contribute to defining practice.

Now consider student assessment data, a tool used in most schools. Student assessment data, as measured by standardized tests and as disseminated by the district office, was frequently used as a tool by the Chicago schools that participated in the Distributed Leadership Study (Spillane, 2006). The data reported on how students in a particular school were doing relative to district averages (see Diamond & Cooper, 2007). Typically, the interactions and organizational routines in these schools centered on skills that students had not mastered yet; in fact, the decision was often made by staff to designate these skills as priorities for the following year.

At Baxter School, however, the school principal redesigned the student assessment data as it was received from the school district and built a system of integrated organizational routines designed to support discussion about these data as well as generate additional data. Baxter School had no need to worry about student achievement because the school had met national standards in core subject areas year after year. Still, the principal took the assessment data and performed longitudinal analyses of student achievement. In doing so, he identified grade- and cohort-level trends which showed that Baxter was not doing as well as suggested by the district averages. The principal then shared graphs of these trends with his staff; the graphs helped focus and frame their discussions about whether there was a problem at Baxter, what the nature of this problem was, and how it might be addressed (Spillane, 2006). In this way, a common tool (i.e., student achievement data) was transformed by the principal into one capable of uncovering a problem with achievement at Baxter. The tool focused interactions among staff in organizational routines (including literacy committee meetings and school leadership team meetings), such that Baxter ended up revising its curriculum. Thus, as a tool, student assessment data was both a situational outcome of practice at Baxter *and* an aspect of the situation that helped structure practice.

Though federal, state, and local policies often equate student assessment data with evidence, there are other types of information that get constructed as evidence in schools and school systems. In fact, such information is often needed to go beyond the initial understanding of a problem that student achievement data affords. For example, although the reanalyzed student achievement data suggested a problem at Baxter, the school staff did not stop there. Using a range of tools (including teacher surveys and curriculum analyses), Baxter generated additional information that helped explain why student achievement was declining as students moved through the grades. The school was then able to

develop constructive solutions that could readily be put into practice (see Burch, in press).

Again, tools and organization routines, along with other aspects of the situation, contribute to defining school work practice. In this way, tools and routines that we often take for granted can shape what gets constructed as evidence in schools. Take the writing folder review routine that was an outcome of practice at Hillside School (see Coldren, in press; Spillane, 2006). Teachers at Hillside submitted a monthly folder to the principal that contained a writing sample for each student in their class (including the feedback that they gave to each student). Based on a review of these samples, the principal provided guidance to teachers on their writing instruction. This organizational routine and its accompanying tools focused and centered the interactions between the principal and teachers on what students actually wrote. Now imagine a different organizational routine—a writing lesson plan review—which would focus the interactions not on what students actually wrote but on what teachers intended to teach. By virtue of how they structure the interactions among staff, these routines generate qualitatively different information about instruction, information that may or may not be constructed as evidence (Spillane & Diamond, in press; Spillane et al., 2005).

Bias in work practices. In this section we consider bias at the level of social interactions. School work practices yield judgments that can deviate from normative standards of evaluation (i.e., socially approved criteria for accepting or rejecting new evidence). Research from social psychology suggests that social interactions can lead to group-level bias in two ways. First, unless there is an organizational press for detecting and eradicating cognitive and motivational biases, biased interpretations of information at the individual level (e.g., biased assimilation) will lead to biased interpretations of information at the group level. For example, if each group member is motivated to reach a conclusion that supports his or her point of view, it is likely that the group will reach a decision that reflects what the majority of group members want to believe (regardless of whether or not this decision is supported by the evidence constructed in support of it) (Kerr & Tindale, 2004). However, if people are pressed to account for each other's motives (as part of an established organizational routine), it is possible that a more objective discussion of information will occur and that a relatively less biased decision will be reached.

Bias will also occur when interactions in organizational routines limit or distort the manner in which information is shared between group members. Although there are a number of biases that fall into this category, four are particularly worth discussing.

Information pooling refers to the tendency of a group to pool or share the information that its members already have in common, as opposed to the information that is unique to particular individuals (Stasser & Titus, 1985). Pooling information and expertise should, in theory, enable a group to make more informed decisions. That is, it should allow individuals to exchange disparate information, which can lead to new insights, understandings, and perspectives on an issue (Brown & Campione, 1990; Brown, Collins, & Duguid, 1989). Unfortunately, research shows that because group members often fail to discuss the information they do *not* have in common, decisions are made that do not accurately reflect the available evidence (i.e., the outcome is biased; Stasser & Titus, 1985; Winquist & Larson, 1998).

Groupthink refers to a way of thinking “that people engage in when they are deeply involved in a cohesive in-group, when the members’ striving for unanimity overrides their motivation to realistically appraise alternative courses of action” (Janis, 1972, p. 9). In groupthink situations, routines for building consensus prevent people from seeking new evidence and from considering alternative interpretations of existing evidence (Janis & Mann, 1977). For example, if a group of teachers feel that it is important to reach a consensus about a proposed solution by the end of their faculty meeting, it is unlikely that opinions contrary to the initial consensus will be expressed (even if many of the teachers disagree with the solution). As a result, it is common for practitioners in groupthink situations to support solutions that are not well thought out.

Group polarization describes the majority’s tendency to intensify its position on an issue after repeatedly discussing the evidence in support of this position (Isenberg, 1986). For example, if all teachers engaged in performing a particular routine (e.g., a grade-level meeting) believe that declines in student reading scores are a function of changing student demographics, then repeated interactions may only serve to reinforce and intensify this belief.

Finally, *group escalation of commitment* refers to situations in which those responsible for performing a routine stay the course and continue their behavior despite evidence that what they have attempted has failed (Staw, 1976; Whyte, 1993). In an effort to justify past decisions or to make these decisions seem rational, evidence (e.g., early grade reading

scores) that a particular approach (e.g., whole language or phonics-based approach to early reading) does not improve student achievement is ignored. For school or district leaders who are responsible for the success of important organizational routines, the press for consistency and validity (which is associated with leadership in general) may only serve to increase group escalation of commitment.

In sum, a distributed perspective requires us to understand *how* aspects of the situation (such as organizational routines and tools) enable and constrain work practice and, in turn, how they are made and remade through work practice. Routines and tools do more than just enable people to work more efficiently; they also structure practice by mediating the interactions among people. They help define how and what people must be heedful of when they interact with one another. Other aspects of the situation also frame interactions in schools, but they are beyond the scope of this chapter. It is through work practice, as described in this section, that new information gets interpreted as evidence for or against a proposed reform.

Conclusion

Over the past decade, many social scientists and some government agencies have fretted about (and tried to improve) the quality of evidence in education, especially research-based evidence. Moreover, policymakers at every level of the educational system have employed a broad array of policy instruments in an effort to ensure that student achievement data will serve as a primary source of evidence for making school-relevant decisions. Together, these efforts represent a dramatic shift in American education. In this chapter, we have attempted to foreground a different aspect of the debate about evidence—how it gets used in local practice. Our account is premised on the notion that increases in high-quality educational research and improvements in student achievement measures will only matter to the extent that they find their way into the day-to-day practice of schools and school districts. The information these sources provide will have to contend with a cacophony of other data that already inhabit most school environments. Certainly, policymakers and researchers can do much to get on the radar screen of schools. However, policymakers and researchers have rather blunt instruments for influencing what happens in the daily practice of most schoolhouses.

Evidence-based practice is not nearly as simple and straightforward as we are often led to believe. And this is because how evidence is

constructed depends to a large extent on the existing beliefs, values, and norms of local practitioners. Taking a sense-making perspective, we argued that in order to understand evidence use we must examine how people make sense of their environment. We must account for what they notice and how they interpret what they notice (as evidence). A central argument in this chapter is the importance of attending to work practices in schools and school districts. Further, we argue that attention to work practice necessitates moving beyond an exclusive focus on actions to examine the interactions among school staff, as mediated by aspects of the situation, including organizational routines and tools.

For those in the trenches, our take-home message is this: work practice is where the rubber meets the road in the schoolhouse. Understanding how information becomes evidence and how this evidence gets used or goes unused requires attention to work practice. Work practice can be difficult to access and even more difficult to analyze. We often gloss over these difficulties by focusing on simplified strategies for evidence-based decision making. But to understand evidence use, we must attend to practice, which necessitates attention to interactions among people, as well as to how these interactions are mediated by aspects of the situation (such as organizational routines and tools). A simple first step for school leaders might involve taking stock of the formal and informal organizational routines in their schools. A next and somewhat more complex step might involve asking some tough design questions about these organizational routines—what purpose do they serve? How should they work in order to achieve this purpose? How should they not work? A third and substantially more complex step involves analyzing the interactions in the performance of these routines—attending not only to the people involved, but also to how aspects of the situation frame and focus their interactions. What tools are in use, and how do they frame the interactions? How might these tools be redesigned, or new tools developed, to frame and focus the interactions in new ways?

School leaders and teachers, like the rest of us, construct evidence to define problems and craft solutions. Solutions sometimes come before the problems they are mobilized to address; in policymaking, a problem is sometimes defined to fit a particular solution. In constructing evidence, school leaders and teachers attend selectively to information in their environment. What they select and what they make of what they select is in part a function of their knowledge, beliefs, and experiences. But school leaders don't make sense alone. They do so in concert with others, and aspects of the situation structure these interactions; as

a result, understanding evidence use necessitates attention to work practice.

AUTHORS' NOTE

This chapter was made possible by support from the National Science Foundation (Grant No. OSR-9250061 and REC-9873583), the Spencer Foundation, the Consortium for Policy Research in Education (CPRE) (Grant No. OERI-R308A60003, U.S. Department of Education), and Northwestern University's Institute for Policy Research and School of Education and Social Policy. We are grateful to two anonymous reviewers for their thoughtful suggestions on an earlier draft of the chapter. We are especially thankful to Pamela Moss for her detailed comments and insightful questions on multiple drafts of the chapter that helped us clarify our thinking. Opinions expressed in this paper do not necessarily reflect the views of any of the funding agencies.

NOTES

1. A potential solution is really just a particular type of hypothesis—one that takes the form of “if we do X then problem Y will be solved.” Other kinds of hypotheses that are frequently tested in educational settings include hypotheses about the existence of a particular problem and hypotheses about why a particular problem occurs (“when schools do X, problem Y occurs”).

2. It should be noted that “cold” biases, which do not result from directional processing goals, are still motivated in the sense that the individual desires to reach an accurate conclusion and desires are motivational by definition.

3. This is different from biased assimilation, in which opposing evidence is actively discounted.

REFERENCES

- Aarkes, H.R. (1991). Cost and benefits of judgment errors: Implications for debiasing. *Psychological Review*, *110*, 486–498.
- Althusser, L. (1971). Ideology and ideological state apparatus. In B. Brewster (Trans.), *Lenin and philosophy and other essays* (pp. 229–236). New York: Monthly Review Press.
- Berger, P.L., & Luckmann, T. (1966). *The social construction of reality*. New York: Doubleday.
- Birkeland, S., Murphy-Graham, E., & Weiss, C. (2005). Good reasons for ignoring good evaluation: The case of the drug abuse resistance education (D.A.R.E.) program. *Evaluation and Program Planning*, *28*(3), 247–256.
- Bourdieu, P. (1981). Men and machines. In K. Knorr-Cetina & A.V. Cicourel (Eds.), *Advances in social theory and methodology* (pp. 304–317). London: Routledge.
- Bourdieu, P. (1990). *The logic of practice* (R. Nice, Trans.). Stanford, CA: Stanford University Press.
- Brewer, W.F., & Nakamura, G.V. (1984). The nature and function of schemas. In R.S. Wyer & T.K. Srull (Eds.), *Handbook of social cognition* (pp. 119–160). Hillsdale, NJ: Erlbaum.
- Brown, A.L., & Campione, J.C. (1990). Communities of learning and thinking, or a context by any other name. *Human Development*, *21*, 108–125.
- Brown, J.S., Collins, A., & Duguid, P. (1989). Situated cognition and the culture of learning. *Educational Researcher*, *18*, 32–41.
- Bryk, A., & Schneider, B. (2002). *Trust in schools: A core resource for improvement*. New York: Russell Sage Foundation.

- Burch, P. (in press). School leadership practice and the school subject: The Baxter case. In J. Spillane & J. Diamond (Eds.), *Distributed leadership in practice*. New York: Teachers College Press.
- Cantor, N., & Mischel, W. (1979). Prototypes in person perception. *Advances in Experimental Social Psychology*, 12, 3–52.
- Carey, S. (1985). *Conceptual change in childhood*. Cambridge: MIT Press, Bradford Books.
- Chase, W.G., & Simon, H.A. (1973). The mind's eye in chess. In W.G. Chase (Ed.), *Visual information processing* (pp. 215–281). New York: Academic Press.
- Coburn, C.E. (2006). Framing the problem of reading instruction: Using frame analysis to uncover the microprocesses of policy implementation in schools. *American Educational Research Journal*, 43(3), 343–379.
- Coburn, C., Toure, J., & Yamashita, M. (2006). *District evidence use: An analysis of instructional decision-making*. Paper presented at the annual conference of the American Education Research Association, San Francisco.
- Coburn, C.E., & Talbert, J.E. (2006). Conceptions of evidence-based practice in school districts: Mapping the terrain. *American Journal of Education*, 112(4), 469–495.
- Cohen, D.K. (1990). A revolution in one classroom: The case of Mrs. Oublier. *Educational Evaluation and Policy Analysis*, 12, 327–345.
- Coldren, A. (in press). Spanning the boundary between school leadership and classroom instruction at Hillside Elementary School. In J. Spillane & J. Diamond (Eds.), *Distributed leadership in practice*. New York: Teachers College Press.
- Cole, M. (1996). *Cultural psychology: A once and future discipline*. Boston: The Belknap Press of Harvard University Press.
- Cole, M., & Engeström, Y. (1993). A cultural-historical approach to distributed cognition. In G. Salomon (Ed.), *Distributed cognitions: Psychological and educational considerations* (pp. 1–46). New York: Cambridge University Press.
- Cook, T.D. (2002, February). *The three cultures of research on effective practices in schools: Given their diversity, what can (and should) the Campbell Collaboration do?* Paper presented at the Campbell Collaboration Colloquium, Philadelphia.
- Cronbach, L.J., Ambron, S.R., Dornbusch, S.M., Hess, R.D., Hornik, R.C., & Phillips, D.C., et al. (1980). *Toward reform of program evaluation*. San Francisco: Jossey-Bass.
- Diamond, J.B., & Cooper, K. (2007). The uses of testing data in urban elementary schools: Some lessons from Chicago. In P.A. Moss (Ed.), *Evidence and decision making. The 106th yearbook of the National Society for the Study of Education*, Part I (pp. 241–263). Malden, MA: Blackwell Publishing.
- Ditto, P.H., & Lopez, D.F. (1992). Motivated skepticism: Use of differential decision criteria for preferred and nonpreferred conclusions. *Journal of Personality and Social Psychology*, 63, 568–584.
- Engeström, Y., & Middleton, D. (Eds.). (1998). *Cognition and communication at work* (2nd ed.). Cambridge: Cambridge University Press.
- Feldman, J. (2000). Minimization of Boolean complexity in human concept learning. *Nature*, 407, 630–633.
- Feldman, M.S., & Pentland, B.T. (2003). Reconceptualizing organizational routines as a source of flexibility and change. *Administrative Science Quarterly*, 48(1), 96.
- Feuer, M.J., Towne, L., & Shavelson, R.J. (2002). Scientific culture and educational research. *Educational Researcher*, 31(8), 4–14.
- Flavell, J.H. (1963). *The developmental psychology of Jean Piaget*. Princeton, NJ: Van Nostrand.
- Freeman, J., & Hannan, M.T. (1983). Niche width and the dynamics of organizational populations. *American Journal of Sociology*, 88(6), 1116–1145.
- Galbraith, J.R. (1973). *Designing complex organizations*. Reading, MA: Addison-Wesley.
- Gentner, D., Rattermann, M.J., & Forbus, K.D. (1993). The roles of similarity in transfer: Separating retrievability from inferential soundness. *Cognitive Psychology*, 25, 524–575.

- Gentner, D., & Stevens, A. (Eds.). (1983). *Mental models*. Hillsdale, NJ: Erlbaum.
- Giddens, A. (1979). *Central problems in social theory: Action, structure, and contradiction in social analysis*. Berkeley and Los Angeles: University of California Press.
- Giddens, A. (1984). *The constitution of society: Outline of the theory of structuration*. Berkeley and Los Angeles: University of California Press.
- Greeno, J.G. (1989). Situations, mental models, and generative knowledge. In D. Klahr & K. Kotovsky (Eds.), *Complex information processing: The impact of Herbert A. Simon* (pp. 285–318). Hillsdale, NJ: Erlbaum.
- Greeno, J.G. (1998). Where is teaching? *Issues in Education*, 4(1), 110–119.
- Greeno, J.G., Collins, A.M., & Resnick, L.B. (1996). Cognition and learning. In D.C. Berliner & R.C. Calfee (Eds.), *Handbook of Educational Psychology* (pp. 15–46). New York: Macmillan.
- Hallett, T. (in press). The leadership struggle. In J. Spillane & J. Diamond (Eds.), *Distributed leadership in practice*. New York: Teachers College Press.
- Hammersley, M. (2001, September). *Some questions about evidence-based practice in education*. Paper presented at the annual conference of the British Educational Research Association, University of Leeds, England.
- Hammond, K.R., Harvey, L.O., & Hastie, R. (1992). Making better use of scientific knowledge: Separating truth from justice. *Psychological Science*, 3, 80–87.
- Higgins, E.T. (1996). Knowledge activation: Accessibility, applicability, and salience. In E.T. Higgins & A.W. Kruglanski (Eds.), *Social psychology: Handbook of basic principles* (pp. 133–168). New York: Guilford Press.
- Hill, H. (2001). Policy is not enough: Language and the interpretation of state standards. *American Educational Research Journal*, 38(2), 289–318.
- Hutchins, E. (1995). *Cognition in the wild*. Cambridge, MA: MIT Press.
- Isenberg, D.J. (1986). Group polarization: A critical review and meta-analysis. *Journal of Personality and Social Psychology*, 50, 1141–1151.
- Janis, I.L. (1972). *Victims of groupthink*. Boston: Houghton Mifflin.
- Janis, I.L., & Mann, L. (1977). *Decision making: A psychological analysis of conflict, choice and commitment*. New York: Free Press.
- Kerr, N.L., & Tindale, R.S. (2004). Group performance and decision making. *Annual Review of Psychology*, 56, 623–655.
- Klayman, J., & Ha, Y.W. (1987). Confirmation, disconfirmation, and information in hypothesis testing. *Psychological Review*, 94, 211–228.
- Kunda, Z. (1990). The case for motivated reasoning. *Psychological Bulletin*, 108, 480–498.
- Larkin, J.H., McDermott, J., Simon, D.P., & Simon, H.A. (1980). Expert and novice performance in solving physics problems. *Science*, 208, 1335–1342.
- Latour, B. (1986). The powers of association. Power, action and belief. A new sociology of knowledge? In J. Law (Ed.), *Sociological Review Monograph* (Vol. 32, pp. 264–280). London: Routledge & Kegan Paul.
- Latour, B. (1987). *Science in action: How to follow engineers and scientists through society*. Milton Keynes: Open University Press.
- Leont'ev, A.N. (1981). *Problems of the development of the mind*. Moscow: Progress.
- Linn, R.L., Baker, E.L., & Betebenner, D.W. (2002). Accountability systems: Implications of requirements of the No Child Left Behind Act of 2001. *Educational Researcher*, 31(6), 3–16.
- Lord, C.G., Ross, L., & Lepper, M.R. (1979). Biased assimilation and attitude polarization: The effects of prior theories on subsequently considered evidence. *Journal of Personality and Social Psychology*, 37, 2098–2109.
- MacCoun, R.J. (1998). Biases in the interpretation and use of research results. *Annual Review of Psychology*, 49, 264.
- Majone, G. (1989). *Evidence, argument, and persuasion in the policy process*. New Haven, CT: Yale University Press.
- Mandler, G. (1984). *Mind and body*. New York: Norton.

- Markus, H., & Zajonc, R.B. (1985). The cognitive perspective in social psychology. In G. Lindzey & E. Aronson (Eds.), *Handbook of Social Psychology* (3rd ed., pp. 137–230). New York: Random House.
- Murphy, G.L., & Medin, D.L. (1985). The role of theories in conceptual coherence. *Psychological Review*, 92, 289–316.
- Norman, D.A. (1988). *The design of everyday things*. New York: Doubleday Currency.
- Novick, L.R. (1988). Analogical transfer, problem similarity, and expertise. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 14, 510–520.
- Phillips, D.C. (2007). Adding complexity: Philosophical perspectives on the relationship between evidence and policy. In P.A. Moss (Ed.), *Evidence and decision making. The 106th yearbook of the National Society for the Study of Education*, Part I (pp. 376–403). Malden, MA: Blackwell Publishing.
- Piaget, J. (1972). *The psychology of the child*. New York: Basic Books.
- Pickering, A. (Ed.). (1992). *Science as practice and culture*. Chicago: University of Chicago Press.
- Resnick, L.B. (1991). Shared cognition: Thinking as social practice. In L.B. Resnick, J.M. Levine, & S.D. Teasley (Eds.), *Perspectives on socially shared cognition* (pp. 1–20). Washington, DC: American Psychological Association.
- Ross, B. (1987). This is like that: The use of earlier problems and the separation of similarity effects. *Journal of Experimental Psychology; Learning, Memory, and Cognition*, 13, 629–639.
- Ross, L., & Ward, A. (1995). Psychological barriers to dispute resolution. In M. Zanna (Ed.), *Advances in experimental social psychology* (Vol. 27, pp. 255–304). New York: Academic Press.
- Rumelhart, D. (1980). Schemata: The building blocks of cognition. In R. Spiro, B. Bruce, & W. Brewer (Eds.), *Theoretical issues in reading comprehension* (pp. 33–58). Hillsdale, NJ: Erlbaum.
- Rumelhart, D., & Ortony, A. (1977). The representation of knowledge in memory. In R. Anderson, R. Spiro, & W. Montagne (Eds.), *Schooling and the acquisition of knowledge* (pp. 99–135). Hillsdale, NJ: Erlbaum.
- Schank, R.C., & Abelson, R.P. (1977). *Scripts, plans, goals, and understanding*. Hillsdale, NJ: Erlbaum.
- Shavelson, R.J., & Towne, L. (Eds.). (2002). *Scientific research in education*. Committee on Scientific Principles for Education Research. Division on Behavioral and Social Sciences and Education. Washington, DC: National Academy Press.
- Shepard, L.A. (2000). The role of assessment in a learning culture. *Educational Researcher*, 29(7), 4–14.
- Sherer, J. (in press). The practice of leadership in mathematics and language arts: The Adams Case. In J. Spillane & J. Diamond (Eds.), *Distributed leadership in practice*. New York: Teachers College Press.
- Siegel, L.S. (1989). IQ is irrelevant to the definition of learning disabilities. *Journal of Learning Disabilities*, 22(8), 469–478, 486.
- Slavin, R. (2002). Evidence-based education policies: Transforming educational practice and research. *Educational Researcher*, 31(7), 15–21.
- Spillane, J. (2006). *Distributed leadership*. San Francisco: Jossey-Bass.
- Spillane, J., & Diamond, J. (Eds.). (in press). *Distributed leadership in practice*. New York: Teachers College Press.
- Spillane, J., & Orlina, E. (2005). Investigating leadership practice: Exploring the entailments of taking a distributed perspective. *Leadership and Policy in Schools*, 4, 157–176.
- Spillane, J., Reiser, B.J., & Reimer, T. (2002). Policy implementation and cognition: Reframing and refocusing implementation research. *Review of Educational Research*, 72(3), 387–431.
- Spillane, J., Sherer, J., & Coldren, A. (2005). Distributed leadership: Leadership practice and the situation. In W. Hoy & C. Miskel (Eds.), *Educational leadership and reform* (pp. 149–167). Charlotte, NC: IAP Publishing.

- Starbuck, W., & Milliken, F. (1988). Executives' perceptual filters: What they notice and how they make sense. In D. Hambrick (Ed.), *The executive effect: Concepts and methods for studying top managers* (pp. 35–65). Greenwich, CT: JAI.
- Stasser, G., & Titus, W. (1985). Pooling of unshared information in group decision making: Biased information sampling during discussion. *Journal of Personality and Social Psychology*, 48(6), 1467–1478.
- 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.
- Strauss, S. (1993). Teachers' pedagogical content knowledge about children's minds and learning: Implications for teacher education. *Educational Psychologist*, 28, 279–290.
- Suchman, L.A. (1983). Office procedures as practical action: Models of work and system design. *ACM Transactions on Office Information Systems*, 1(4), 320–328.
- Towne, L., Wise, L.L., & Winters, T.M. (Eds.). (2005). *Advancing scientific research in education*. Washington, DC: National Academy Press.
- Trope, Y. (1986). Identification and inferential processes in dispositional attribution. *Psychological Review*, 93, 239–257.
- VanLehn, K. (1989). Problem solving and cognitive skill acquisition. In M. Posner (Ed.), *Foundations of cognitive science* (pp. 526–579). Cambridge, MA: MIT Press.
- Vosniadou, S., & Brewer, W.F. (1992). Mental models of the earth: A study of conceptual change in childhood. *Cognitive Psychology*, 24, 535–585.
- Vygotsky, L.S. (1978). *Mind and society: The development of higher mental processes*. Cambridge, MA: Harvard University Press.
- Webb, J.M., Diana, E.M., Luft, P., Brooks, E.W., & Brennan, E.L. (1997). Influence of pedagogical expertise and feedback on assessing student comprehension from non-verbal behavior. *The Journal of Educational Research*, 91(2), 89.
- Weick, K.E. (1979). *The social psychology of organizing* (2nd ed.). New York: McGraw-Hill.
- Weick, K.E. (1995). *Sensemaking in organizations*. Thousand Oaks, CA: Sage Publications.
- Wertsch, J.V. (1991). *Voices of the mind: A sociocultural approach to mediated action*. Cambridge, MA: Harvard University Press.
- Wertsch, J.V. (1998). *Mind as action*. New York: Oxford University Press.
- Whittington, R. (1992). Putting Giddens into action: Social systems and managerial agency. *Journal of Management Studies*, 29(6), 693–712.
- Whyte, G. (1993). Escalating commitment in individual and group decision making: A prospect theory approach. *Organizational Behavior and Human Decision Processes*, 54, 430–455.
- Wilson, T.D., & Brekke, N. (1994). Mental contamination and mental correction: Unwanted influences on judgments and evaluation. *Psychological Bulletin*, 116(1), 117–142.
- Winquist, J.R., & Larson, J.R., Jr. (1998). Information pooling: When it impacts group decision making. *Journal of Personality and Social Psychology*, 74, 371–377.