Pounds, J., & Fallesen, J. J. (1995). *Familiarity effects on strategy use in tactical problem solving.* Poster presented at Society for Judgment and Decision Making Conference. November 11–13, 1995. Los Angeles, CA.

Pounds, J., & Fallesen, J. J. (1997). *Problem solving of mid-career Army officers: Identification of general and specific strategies* (Research Note 97–21). Alexandria, VA: U.S. Army Research Institute for the Behavioral and Social Sciences.

Schmitt, J. F. (1994). *Mastering tactics: A tactical decision game workbook.* Quantico, VA: Marine Corps Association.

Serfaty, D., MacMillan, J., Entin, E. E., & Entin, E. B. (1997). The decision-making expertise of battle commanders. In C. E. Zsambok & G. Klein (Eds.), *Naturalistic decision making* (pp. 233–246). Mahwah, NJ: Lawrence Erlbaum Associates.

Shanteau, J. (1995, August). *Why do experts disagree?* Paper presented at the 15th Bi-Annual Conference on Subjective Probability, Utility, and Decision Making, Jerusalem, Israel.

Shanteau, J. (1996, April). *The upper limits of expert performance and the search for world records: What is the question?* Invited paper presented at the meeting of the Oklahoma/Kansas Judgment and Decision Making Society, Norman, Oklahoma.

Woodward, B. (1991). *The commanders.* New York: Schuster.

5

Solving the Problem of Employee Resistance to Technology by Reframing the Problem as One of Experts and Their Tools

Lia DiBello
Workplace Technology Research Group
CUNY Graduate School

This chapter addresses the relationship between human cognition and tools as it applies to the problem of rapidly changing information technology, an issue I have been studying for a number of years (e.g., DiBello, 1996a, 1996b, 1997; DiBello & Kindred, 1992; DiBello & Spender, 1996; Scribner, Sacks, DiBello, & Kindred, 1991; Scribner, DiBello, Kindred, & Zazanis, 1992). Although we focus on relatively complex information technologies, we continually find that the fundamental mechanisms of tool mastery and appropriation remain deceptively the same, regardless of tool complexity. We believe that most of what human beings do with complex tools could be predicted by close observation of how people use and modify even the simplest tools.

Since the 1970s large information technologies have been fundamentally changing many industries. Of particular interest to us are those that are changing work by making the manipulation and analyses of information easier and more widespread. These systems are usually large, highly integrated information systems that capture relatively "live" data for purposes of analysis and make possible rapid changes in business strategy. Two examples we have studied in depth are Materials Requirements Planning (MRP) and MRPII and Computerized Maintenance Management Systems (CMMS). A brief description of these two types of systems will help clarify some of our later points.

THE IMPACT OF COMPLEX TOOLS

MRPII

MRP has been characterized as a theory of inventory and material management. It instantiates certain key economic concepts such as *zero inventory* and *just-in-time* production and is based on principles of manufacturing (for example, formulas regulating how future orders are forecast) developed over the last several decades (Harrington, 1974; Timms & Pohlen, 1970). In many ways it is considered a somewhat counterintuitive approach to material planning in that it "begins" in the future and moves backward in time. However, when properly executed, on-hand inventory can be reduced by as much as 70%, freeing a large part of a company's capital.

In general, MRP approaches are in contrast to traditional "aggregate" planning. Aggregate planning methods evolved after WWII and were influenced by material scarcity concerns. The goal of most aggregate planning methods was to accumulate as much raw material as possible to cope with growing demand. With improved distribution and material availability, inventory surplus soon became a significant and costly business problem. So called just-in-time approaches evolved as a result of a better understanding of fixed or known demand.

However, in actual practice MRP systems are often used to automate aggregate planning practices. In these cases, the implementation is not considered successful by industry standards. In fact, it is actually easier for employees to make things worse when they misuse the system.

CMMS

CMMS are a rapidly growing set of systems for managing activity. Industry leaders estimate that there are now about 80 different products on the market. In some ways, they begin where MRP systems leave off. MRP systems are concerned with planning "things" (raw material, components, assemblies) to meet anticipated and very specific demand. The activities associated with the material are implied by MRP in the form of material routings. CMMS, on the other hand, is somewhat of a mirror image of this process: these systems plan activity, and material, components, and assemblies are only locations for activity.

CMMS systems are appropriate when activity and its details are the focus. Just as MRP and MRPII have become important in manufacturing, CMMS has become important in industries where maintenance or monitoring activity with an enduring asset ensures smooth delivery of a product or service. Some examples are railroads, bridges, or power generation plants.

CMMS systems assume that assets and their internal components and assemblies have fixed life cycles based on chronological time, operational time (such as service hours) or mileage. The underlying assumption is that life cycles are a function of equipment-environment interactions and that these interactions must be tracked in order to empirically derive and predict the length and nature of life cycles. Therefore, most of these systems are built to collect data from work orders (to do something to equipment) usually triggered by a "symptom." Over time, these systems link symptom types with defect types, and ultimately link both to a piece of equipment or one of its components. The systems then use this information to predict the time span of "natural" life cycles and plan when equipment needs to be replaced or serviced due to their completion. The ideal assumption underlying these systems is that the entire out of service time of a piece of equipment due to component failure can be predicted and prevented by these systems. Failure trends are then used to plan pre-failure "change outs," ensuring virtually seamless operation.

The approach represented in CMMS is very much opposed to traditional methods for asset maintenance, which are—at least formally—highly reactive. As with aggregate material planning methods, reactive methods of maintenance are also an outgrowth of post-WWII scarcity. They also assume that life cycles of equipment are unpredictable and that the most cost-effective approach to maintenance is to milk an asset for all it is worth by running it to failure. Now that components can be procured easily, however, reactive methods of maintenance are considered to be unnecessarily costly. In fact, recent data provided by the Society of Automotive Engineers indicate that in the 1960s the parts/labor ratio was 2/1. In the 1990s the ratio completed a full reversal (parts/labor = 1/2). Maintenance practices that emphasize reactive repairs also require redundant systems, large spare factors, and significant loss of revenue opportunity when equipment is down for repair. Further, when large fleets of buses or trains support the economic functioning of large metropolitan areas, and reliable service is expected, running to failure introduces unacceptable uncertainty.

Important as they are, technologies such as those described above have enjoyed mixed success in workplaces. Current literature on their failure (e.g., Boldt, 1994, 2000) indicates that large information technologies are hard to implement. Typical implementation times for the introduction of complex information technologies such as MRPII and CMMS are on the order of 12 to 18 months (per site) and success rates have been low, especially for systems that offer the opportunity for company-wide resource management through planning. For example, studies of technologies of this kind by the Gartner Group and others (for the transportation industry) have shown that as many as 50% of new systems are abandoned in the first year and possibly 90% never reach their full potential. In the manufacturing arena, success rates are as low as 20%. I have been told by senior managers for vendors of MRP systems that these methods have a return rate as high as 76% for the software.

Most businesses now recognize that the failure of these kinds of technologies is essentially a "user" problem rather than a technology problem. However, we think that businesses, especially, grossly misunderstand the nature of this user problem. The most striking misunderstanding involves the idea of "resistance."

For example, one commonly given reason for the failure of these technologies is workers who fear and are resistant to change. In general, most firms now acknowledge that technologies such as CMMS and MRP require 99% data accuracy and are sensitive to level-of-detail issues. Therefore, many attempts have been made to deploy technologies among frontline workers (such as mechanics and assemblers) who are in contact with the details of day-to-day operations or have detailed knowledge of equipment. However, this group has not responded well. Often, they are seen as not having adequate computer skills. When training has been attempted, frontline workers have not learned much from the (usually vendor provided) classroom instruction. They may not learn what is needed, do not transfer what they learn to practice, or resist the training experience itself.

I propose that a cognitive analysis of behavior that looks like resistance yields a deeper understanding of how people change and learn and offers greater opportunities for productive technology deployment. For example, frontline workers' resistance to technology instruction is actually partly rooted in their success as craftspersons, which both selects and develops a learning style that is based on problem solving, experimentation, and hands-on contact. As will become clear below, when pedagogy is modified to fit their learning style, frontline workers show themselves to be superior learners with resistance acting as a catalyst.

THE THEORETICAL PERSPECTIVE

The focus of my work concerns the cognitive impact of the introduction of technology into the workplace. Specifically, I am interested in exploring how workers' ways of thinking and understanding are affected by changes in the nature of work and workplace organization. Many of my questions have been addressed using a number of different models, such as "novice-expert shift" (e.g., Chi, Glaser, & Farr, 1988), "situated cognition" (e.g., Rogoff & Lave, 1984), or "naturalistic decision making" (e.g., Orasanu & Connolly, 1993) and my work has been influenced by the methods and theoretical models from all of these various approaches. However, since the focus of my inquiry concerns the *development* of different ways of thinking in different domains, my research has been most influenced by the theories and methods of developmental psychology and particularly by the developmental theories of Vygotsky (1987) and Scribner's application of them to workplaces and workplace cognition (e.g., as summarized in Scribner et al., 1996).

Cognition and skills develop in the service and support of activities at work (DiBello, 1996a, 1996b). This is the principle difference between school learning and ongoing learning at work. As one participates in a particular industry or occupation, specific strategies and ways of understanding the business at hand are selected and reinforced as they prove over time to have a direct bearing or accomplishing important goals (DiBello & Kindred, 1992; Scribner et al., 1992). Workers typically understand they are learning the right things when they are more fully able to participate in meaningful problem solving and are recognized for their value by their coworkers (Scribner et al., 1991; Scribner et al., 1996). This set of skills and these ways of understanding work comprise the culture of any workplace. Over time, the culture of practice takes on a life of its own, being passed on to new workers as they "learn the ropes."

When looking at workplace culture as really being about skills developed collectively (and over time) in service of accomplishing goals, it becomes clear why the culture of a workplace becomes the main impediment when a widespread rapid change in business practice is being introduced. This is even more the case when the sudden changes are represented in an information technology that affects every job. A long-standing culture of practice can become suddenly obsolete, at least in part.

When a change is being introduced, change agents (i.e., new management, consultants or a process improvement team) will often disregard any usefulness that previous strategies may add in the process of change. They are often unaware of the important role that prior knowledge can play in the "new" vision (Chamberlain & DiBello, 1997). Their strategy is often to replace all legacy practices, by either "selling" the change or eliminating key resisters. This does not acknowledge the importance of content knowledge employees accumulate over the years. The process of integrating useful aspects of legacy skills with practices that support new and changing business goals is required for any positive change (DiBello, 1997a, 1997b).

The Role of Constructive Activity in Learning

In my first studies of workplaces undergoing technology changes my colleagues and I made a small discovery at a plant north of New York City that influenced a great deal of my subsequent work. In a study of workers using MRP (Scribner, Di Bello, Kindred, & Zazanis, 1992) in two different factories—one with a successful implementation and one with an unsuccessful implementation—classroom instruction was shown to be a poor way of preparing workers to use MRP effectively at either plant (DiBello & Glick, 1993; Scribner et al., 1991; Scribner et al., 1992). Despite this, at one plant many individuals managed to master MRP and reduce their inventory by 72%. It turned out that on-the-job activity proved to be critical to developing the necessary skills. An analysis of day-to-day job activity by people in three comparable titles and levels of responsibility revealed two distinct patterns of activity: constructive and procedural. Briefly, *constructive* activities are those that have clearly defined goals and poorly defined means. The employee is compelled to "construct" a procedure, form, tool, or artifact that accomplishes some meaningful goal in an iterative fashion. In contrast, *procedural* activities are those that have clearly specified means and order of execution but goals that are either clearly conveyed or not. Important to note, constructive activities were associated with an in-depth understanding of MRP's underlying logic while procedural activities were not. In fact, when several variables—job title, years of experience, level of formal education, and number of opportunities (weekly) for constructive activities—were correlated with measures of in-depth grasp of MRP principles, only number of opportunities for constructive engagement was found to be significantly associated with mastery ($r = .69$, $p = <.01$; see Di Bello and Glick, 1993, for discussion). However, this study also showed that opportunities for constructive activities are usually fortuitous and ill structured. For example, they often occur because the person who knows what to do has left the job without documenting procedures for others to use.

After this study ended, my aim was to better understand the role of constructive activity when it comes to technology and to find the means to increase opportunities for it in the workplace. In the effort described below my colleagues and I tried to systematize opportunities for constructive technological activity through specially designed exercises. These exercises were developed to help employees better understand that they were participating in a particular set of practices that may have become obsolete, and to help them construct a new set of practices more relevant to their company's goals. As becomes clear, my colleagues and I found we had to develop an in-depth understanding of the company's legacy domain of practice in order to design these exercises.

MRPII and Transit Workers

The initial attempt to bring the benefits of accidental on-the-job constructive activity into an intentional intervention involved transit mechanics learning MRP. As I have already detailed in other publications (e.g., DiBello 1996, 1996b, 1997a), we provided real workers with an opportunity for accelerated constructive activity in manipulative simulations in which they gained an understanding of MRPII by having to construct a manual version, act as MRPII using manual means, and then implement whatever plan they had made. One small study was sponsored by the Spencer Foundation and conducted in 1993 among mechanics in the compressor shop at the New York City Transit Authority subway department. The outcome was that the typical long learning curve for MRP systems was greatly shortened among personnel who normally would not be targeted as users. This occurred because designing a very simple exercise simulating the mechanics' workplace and inventory concerns worked well and seemed to bypass the need for prerequisite knowledge of computers. In fact, the exercises did not focus on the "computerness" of MRPII, but rather on the conceptual differences between aggregate planning and MRP methods of planning. In previously published articles about the experiment I have explained how this training activity is different from other simulation training in two specific ways: 1) the hands-on nature of the simulations provides more avenues of engagement; rather than being a virtual workplace on a computer screen, physical production and manufacturing are required and 2) participants first go through the simulation with very little guidance, in other words, although constraints and goals were made clear, procedures for accomplishing them are not. Participants in the exercise were constrained to select the development of MRP methods according to the goals and with the resources (or tools) provided; all manner of methods were available (including MRP-like and traditional planning sheets and the information needed to use them) but the miniature business could only operate within budget by using MRP methods of material planning.

As elaborated elsewhere, in-depth knowledge of conceptual domains is constructed. However, part of this construction is a kind of "deconstruction" of existing expertise. This leads to the third difference with the hands-on simulations: in order for construction to happen properly, it seems vitally important that the simulation engage the participants' implicit ways of thinking (by introducing time pressure) and allow them to systematically fail.

The idea here is that learning via construction is actually about reorganizing existing knowledge, and existing intuitive expertise could not be reorganized for a new purpose without significant engagement and, ultimately, an activity-based challenge. In a sense, the challenge weakens the a priori nature of expert knowledge (as the learner notices the failure and begins to reassess the situation), and, therefore, fundamental reorganization in one's ways of thinking may have to involve failure.

THE MAINTENANCE INFORMATION DIAGNOSTIC ANALYSIS SYSTEM EXPERIMENT

For many years, NYCTA management wanted to implement a centralized life cycle-based maintenance system, or CMMS, as described in this chapter. Transit professionals had long known that CMMS could help them reduce costs and increase service, but few successful applications existed, and none of them were in the public sector. NYCTA made a number of heroic attempts to bring this approach to its maintenance divisions. Manual systems proved unwieldy, however, given the size of the fleet (over 8,000 buses and subway cars), and it was widely acknowledged that early information technologies failed for many of the same reasons cited for MRPII failure. The information needed to make the system work had to be extremely accurate at just the right level of detail. Ideally, the information needed to be inputted by the mechanic him- or herself. However, efforts to train mechanics on computers were not successful. At NYCTA, many of the workers neither spoke English as a first language (about 80%) nor knew how to use a computer keyboard. Also, frontline workers, in general, usually threatened by information systems on the shop floor, can cause widespread system sabotage or damage to expensive computer equipment.

Eventually, with the availability of powerful relational databases, interest in using new technology for maintenance purposes was renewed. In the early 1990s, NYCTA began planning its version of CMMS and called it a Maintenance Information Diagnostic Analysis System (MIDAS). At the onset of the design process, senior management in the Department of Buses decided to have the hourly Bus Maintainers, Class B (BMBs), enter their own repair data into the system without clerical assistance. There were two reasons for this decision. First, budget cuts forced a reexamination of redundant work; asking mechanics to record information in longhand and then have clerks type the same information into a computer represented a particularly costly practice. Second, considerable evidence showed that original handwritten records were much more accurate than what clerks (or supervisors) eventually entered. Therefore, senior management moved all data entry responsibility directly to the shop floor. The decision was occasion for considerable nervousness in middle-management. In general, this approach had never succeeded anywhere except in a few private transportation companies (such as UPS), where workers are carefully screened before hiring. Since making floor workers responsible for a management technology had never been tried before on such a large scale, the MIDAS team believed that any traditional education approach based on this very different kind of user would be inappropriate at best.

Our relatively minor success with MRPII and the "compressor gang" in 1993 were the impetus for making large-scale, frontline computer systems a reality in the NYCTA Bus Division. Specifically, senior management saw our project as successful in achieving mechanic acceptance and for mitigating system sabotage. During our first conversations, management did not recognize that the mechanics accepted the system because they had learned its business purpose and were using it as a tool for their work. Nor did they agree that user knowledge of the buses might be critical to the successful use of the system from a management perspective. That is, they did not recognize that understanding the reason for the system might affect the quality and nature of the data entered by mechanics, and that this level of quality would, in turn, affect the analytic results of the system's pattern analysis capability. In short, the frontline mechanic was not seen as a person with knowledge of the buses that could be critical to cycle identification.

Rather than attempt to convince management that these factors were important, my colleagues and I proposed a "training pilot" at one location, ostensibly to increase user acceptance and prevent system sabotage. The project we designed and eventually rolled out to 19 locations actually addressed user understanding of the reasons for the system. In fact, our exercises engendered the theory behind cycle-based preventive replacement and techniques for trend analysis

and planning. That is, rather than emphasize procedures for using the system, we emphasized conceptual context. We believed the mechanics' mental model of maintenance was the root cause of their resistance, or poor learning, in the first place; we thought that once they understood life cycle-based maintenance theory and the system's assumptions, the computer would seem like any other tool in the shop. Further, to add to our own research, we designed measurements to examine the relationship between user knowledge, data quality, and financial impact.

Learning About the Culture of Fixing Buses

As indicated earlier, the entire study rests on the assumption that efforts to change a workplace culture most often fail because there is an already-functioning, cohesive culture that is actively competing with the change. In order to effect change, one must know as much as possible about the competing culture of practice.

On the surface, it seemed, in the case of the NYCTA study, that our "competition" was reactive maintenance and the attendant belief that parts do not have natural life cycles. However, this still did not tell me what actual practices instantiate these beliefs. From my experience, I knew that on the frontline of the business (usually the shop floor level) the picture is more complicated. Often there is some tension or inefficiency at the front line of the business that has led decision makers to consider alternatives. In these situations, legacy methods are already failing to meet challenges and new things are being tried. This informal domain of practice is usually the real source of competition and the real source of culture change failure.

Many ask how to "get at" the legacy domain of practice. Very few people in a given workplace are explicitly aware of the dominant domain of practice, but most are aware of when they operate effectively within its parameters. They know who is effective, who knows what is going on, and they are able to assess the meaning and significance of situations that are baffling to outsiders. The trick is to tap into the ways that these workers understand their workplace and its business. There is ample reason to believe that people who have implicit expertise in a given area are not the best at narrating their processes of working and making decisions (Dreyfus, 1997), especially in dynamic settings (Klein, 1999) such as vehicle maintenance.

Therefore, in order to understand more about how mechanics actually think about the business of fleet maintenance I felt we needed to begin by observing them on the job, but in such a way that led to understand from their point of view what it is like to do the job. In order to make this a natural and comfortable observation while still allowing us to ask questions as they worked, my colleagues and I each did our fieldwork as a "quasi-apprentice." In this role, it is normal to ask questions, want explanations for decisions, and be curious about the underlying reasons for doing things. Also, it put the experienced worker in the position of "master" or "teacher," which are roles they have to assume many times when breaking in new workers.

Following is one sample of the recorded dialogue taken during an observation of a periodic inspection.

Ed: Well, there is not much to this. We just go down the checklist. Nothing to it really.

Lia: So we start at the top and just go down . . .

Ed: No, I don't do that. I mean, I skip around the list.

Lia: Why is that?

Ed: Well, the order doesn't make sense. See that guy back there (points to rear of bus), I'll be in his way if I start back there. And if I follow the list exactly, I'll be running around the bus all day, literally. So I begin with the things in front. And since I have it up on the lift, I begin with the things underneath first.

Lia: Okay.

Ed: (Looking at steering arm bushing under bus.) Here, hold this flashlight for me. (Picks at dirt and rust around bushing.)

Lia: What's that?

Ed: that's the bushing. What's bothering me here is that it looks like some rust here. That's not good. Shows me there's a problem. Let's look and see when this is due back in. (Looks at schedule of inspections and picks more at the dirt and rust around bushing.)

Lia: What's up?

Ed: Well, see this bushing over here. Shine the light right here. This is good. See, no rust mixed in with the dirt. Now look at this one. There is some rust in here. But not too much. Not very red. See that? (Researcher sees no difference.) That bushing really needs to be changed. But given that this is coming in in 3000 miles for an A inspection, we can take care of it then. It's got at least that much time on it left. And they need this bus this afternoon. It's gotta wait. So we will make a note of it.

Lia: How do you know it has another 3000 miles left on it.

Ed: Well, it's obvious. By the color of the dirt. The amount of rust in there.

As can be seen from this transcript, even though the mechanic reported on an earlier occasion that he doesn't "think" but rather does what he is told to do, there are a significant amount of situation assessment, analysis, and information coordination (and life cycle-based maintenance) being done here. What this and other observations reveal is that experienced mechanics have an intuitive understanding of the life cycles and the coordination of life cycles among components within one piece of equipment. In other words, there is already an informal culture of preventive, coordinated maintenance operating when the formal practices of reactive maintenance threaten the depot's ability to make service requirements. However, it not yet explicit or consistent and has a "plan B" status as a practice.

I also was able to observe how the mechanics learn during the course of doing their work. None of the workers we observed considered themselves to be strong classroom learners or "read and write" types. Most contributed to their own ongoing learning by "puzzle solving," and when stumped, drew on the opinions and observations of peers to help them understand the equipment through systematic group experimentation.

These two observations—the existence of an implicit scheduled maintenance domain of practice and the mechanics' evolved method of learning—greatly influenced the next design decisions of the research project, that is, the cognitive probes designed to tap into individuals' ways of thinking about maintenance and the training exercise to move people into a new way of thinking.

The Construction of the Cognitive Probes

The next task was to understand the mental models being used in the daily business of doing work. I had found that the best way to do this is at the level of the individual using cognitive probes. These probes are very similar to those originally used by Klein et al., 1989, and resemble in spirit his critical decision making interview method. However, there are some important differences. Klein's method is an attempt to get at an expert's implicit knowledge and situation assessment skills by asking him or her to tell a work history story and explore the methods by which he or she reasons it through. Our approach was to constrain the problem-solving context and see how our interviewees view and handle the constraints we have defined. This involved setting up the problem and the tools available for solving it in a uniform way, while at the same time having a situation that invited the interviewee's implicit skills and situation assessment biases. The method involved the following steps:

1. Identify the strategies and practices associated with each domain that make sense only within the worldview of that domain. For example, in vehicle maintenance catalogue all the strategies associated with proactive life cycle maintenance (as one domain) and all the strategies associated with reactive run-to-failure practices (as a contrasting domain).
2. Identify behaviors associated with the strategies in the workplace in which I am doing the research.
3. Design a meaningful problem situation that can be solved using the strategies and behaviors from either domain, or from a mix of both.
4. Design one or more additional problem situations that are similar to those in step 3 but which are more abstract and generic than the site-specific versions.

5. Develop a scoring form that permits a coder to easily check off the strategies/behaviors and calculate the proportion of the strategies used from each domain to produce a solution.

For NYCTA, my colleagues and I constructed two basic tasks, each of which had three variations. The first task was an "active" task: given a pile of work orders, we asked the interviewee to look them over and then make five piles for each day of the work week. In other words, schedule the work. Below is a small sample of the strategy/behavior pairs for solving this task arranged according to domain of practice.

The same task was also given in two other forms: using another piece of equipment that is commonly known (bicycles) and using "Machines A-N," which were purely made up items with meaningless codes as defect or component indicators (such as defect Mu8).

TABLE 5.1
Strategies for Scheduling Work

Cycle-Based Scheduled Maintenance	*Traditional Reactive Maintenance*
Strategy: Coordinating work within asset **Behavior:** Interviewee sorts work-order cards first by equipment ID number. Or asks: can the same asset be taken out of service only once to satisfy multiple problems?	**Strategy:** Coordinating work by type or craft **Behavior:** Interview does initial sort of work-order cards by type of job or type of trade needed to do job, regardless of equipment ID number.
Strategy: Maximizes inservice time **Behavior:** Interviewee compares number of assets coming into shop with number needed for service. Assigns work accordingly. Brings in asset twice on two different days only as necessary to make service.	**Strategy:** Coordinating work within shop capacity **Behavior:** Interviewee distributes work-orders evenly among the days, regardless of the type of work needed to be done
Strategy: Identifying component life cycles based on empirical data **Behavior:** Asks if there are any historic records available that might help differentiate "normal" wear from "abnormal" failure.	**Strategy:** Attempts to reduce maintenance costs by identifying defective components or warranties **Behavior:** Asks if there are historic records in order to determine if a component was recently replaced and is therefore defective or under warranty.
Strategy: Coordinating cycles with each other into clusters **Behavior:** Looks at "what else has been failing" on work histories and speculates about clustering preventive replacement for components with similar life cycles.	**Strategy:** Assumes no life cycles but recognizes "infant mortality" **Behavior:** Looks at the symptom remarks on the work-orders to ascertain if problem is repeater and if part will need replacing.

A second group of tasks also took three forms that were more passive and required the interviewee to interpret information. The objects varied in the same way as the first set: interviewees interpreted bus repair histories, bicycle repair histories, and those for "machines A-N." A similar set of strategies (for interpreting data) to those shown above was used to code the protocol. Photographs were taken of the interviewee's piles and of any drawings or writing and all talking and "thinking aloud" were audiotaped.

Within each task, about 70% of each interviewee's strategies were reactive and about 30% were proactive, indicating that some proactive planning skills had developed in the workplace. There was also striking homogeneity among interviewees in the pilot depot, suggesting a strong workplace cultural effect.

An Education Process That Would Enable the BMBs to Do Data Entry

Based on the fieldwork and the cognitive battery results, my colleagues and I decided to construct a three-part manipulative simulation of a miniature depot, constraining the goals and resources in such a way that, in order to "win" (i.e., make service requirements and stay within budget), the participants had to use proactive strategies that were logically consistent with CMMS (and MIDAS in particular). Our previous work suggested that constructive activities in real workplaces lead to learning because they elicit the implicit knowledge that the worker has to bring to the problem and at the same time select against nonworkable strategies (through experiences of failure). Therefore, the first part of the exercise was designed to "engage the default" within the context of new business goals.

Specifically, teams of 8 participants were asked to run a depot of 40 plastic buses with relatively complex interior components. The goals were to maintain 32 buses in service at all times (limiting the number out of service to 8), order all the materials (within a budget) needed for doing so, and evaluate daily operator reports (each "day" being 20 minutes) that might indicate potential problems (e.g., noisy engine). The activity was rigged so that the only way to meet these goals was to predict what was due to break next. The breakdown patterns of all components followed time/mileage cycle rules and were precalculated using a computer. The toys were then actually "broken" according this pattern. The participants were given adequate tools to predict and calculate this breakdown (including printouts of every bus's repair history, among other items), but were given other tools as well, including those similar to those used to do "reactive" maintenance.

The trainers also played a role. One acted as dispatcher, regularly demanding buses to satisfy routes, while the other acted as a parts vendor and a Federal Transit Administration inspector, looking for safety violations or abuses of public funding, such as overspending or cannibalizing.

5. Develop a scoring form that permits a coder to easily check off the strategies/behaviors and calculate the proportion of the strategies used from each domain to produce a solution.

For NYCTA, my colleagues and I constructed two basic tasks, each of which had three variations. The first task was an "active" task: given a pile of work orders, we asked the interviewee to look them over and then make five piles for each day of the work week. In other words, schedule the work. Below is a small sample of the strategy/behavior pairs for solving this task arranged according to domain of practice.

The same task was also given in two other forms: using another piece of equipment that is commonly known (bicycles) and using "Machines A-N," which were purely made up items with meaningless codes as defect or component indicators (such as defect Mu8).

TABLE 5.1
Strategies for Scheduling Work

Cycle-Based Scheduled Maintenance	*Traditional Reactive Maintenance*
Strategy: Coordinating work within asset **Behavior:** Interviewee sorts work-order cards first by equipment ID number. Or asks: can the same asset be taken out of service only once to satisfy multiple problems?	**Strategy:** Coordinating work by type or craft **Behavior:** Interview does initial sort of work-order cards by type of job or type of trade needed to do job, regardless of equipment ID number.
Strategy: Maximizes inservice time **Behavior:** Interviewee compares number of assets coming into shop with number needed for service. Assigns work accordingly. Brings in asset twice on two different days only as necessary to make service.	**Strategy:** Coordinating work within shop capacity **Behavior:** Interviewee distributes work-orders evenly among the days, regardless of the type of work needed to be done
Strategy: Identifying component life cycles based on empirical data **Behavior:** Asks if there are any historic records available that might help differentiate "normal" wear from "abnormal" failure.	**Strategy:** Attempts to reduce maintenance costs by identifying defective components or warranties **Behavior:** Asks if there are historic records in order to determine if a component was recently replaced and is therefore defective or under warranty.
Strategy: Coordinating cycles with each other into clusters **Behavior:** Looks at "what else has been failing" on work histories and speculates about clustering preventive replacement for components with similar life cycles.	**Strategy:** Assumes no life cycles but recognizes "infant mortality" **Behavior:** Looks at the symptom remarks on the work-orders to ascertain if problem is repeater and if part will need replacing.

A second group of tasks also took three forms that were more passive and required the interviewee to interpret information. The objects varied in the same way as the first set: interviewees interpreted bus repair histories, bicycle repair histories, and those for "machines A-N." A similar set of strategies (for interpreting data) to those shown above was used to code the protocol. Photographs were taken of the interviewee's piles and of any drawings or writing and all talking and "thinking aloud" were audiotaped.

Within each task, about 70% of each interviewee's strategies were reactive and about 30% were proactive, indicating that some proactive planning skills had developed in the workplace. There was also striking homogeneity among interviewees in the pilot depot, suggesting a strong workplace cultural effect.

An Education Process That Would Enable the BMBs to Do Data Entry

Based on the fieldwork and the cognitive battery results, my colleagues and I decided to construct a three-part manipulative simulation of a miniature depot, constraining the goals and resources in such a way that, in order to "win" (i.e., make service requirements and stay within budget), the participants had to use proactive strategies that were logically consistent with CMMS (and MIDAS in particular). Our previous work suggested that constructive activities in real workplaces lead to learning because they elicit the implicit knowledge that the worker has to bring to the problem and at the same time select against nonworkable strategies (through experiences of failure). Therefore, the first part of the exercise was designed to "engage the default" within the context of new business goals.

Specifically, teams of 8 participants were asked to run a depot of 40 plastic buses with relatively complex interior components. The goals were to maintain 32 buses in service at all times (limiting the number out of service to 8), order all the materials (within a budget) needed for doing so, and evaluate daily operator reports (each "day" being 20 minutes) that might indicate potential problems (e.g., noisy engine). The activity was rigged so that the only way to meet these goals was to predict what was due to break next. The breakdown patterns of all components followed time/mileage cycle rules and were precalculated using a computer. The toys were then actually "broken" according this pattern. The participants were given adequate tools to predict and calculate this breakdown (including printouts of every bus's repair history, among other items), but were given other tools as well, including those similar to those used to do "reactive" maintenance.

The trainers also played a role. One acted as dispatcher, regularly demanding buses to satisfy routes, while the other acted as a parts vendor and a Federal Transit Administration inspector, looking for safety violations or abuses of public funding, such as overspending or cannibalizing.

Despite loud disclaimers, people tended to construct solutions to even novel problems that fit with their experience, even when explicitly instructed to avoid doing so. In fact, the participants were rarely aware they were replicating their normal methods.

Rather than interfere with this tendency, the trainers allowed the participants to wing it, while carefully documenting the cash flow, labor flow, inventory acquisitions, and the number and type of on-the-road failures that resulted from failing to predict problems. Meanwhile, heavy fines were levied for expensive reactive problem-solving strategies, such as cannibalizing an entire bus for a few cheap parts in order to get other buses back on the road. As the activity progressed, participants were continually shown the financial consequences of their decision making patterns and asked what they were thinking by the vendors/inspectors and dispatchers. In general, by the end of the first day of these sessions, the depot was in crisis and the participants realized their budget was being expended to react to mounting problems. At that point, the activities were stopped and the team was sent back to work or to lunch.

On the second day of the exercise, participants reflected on what they did, as recorded by the trainers. The participants were asked to discuss among themselves what thinking led to various decisions and to begin to identify practices that lead to bad outcomes versus practices that are preventative. Only at this point were participants truly open to new ideas about how to solve the problems of vehicle maintenance. They also began to understand in detail the ways that their "gut feel" decisions revealed how they have actually misunderstood preventative maintenance.

In the last hour of the second day, the trainers facilitated the participants in building a manual scheduled maintenance system. The participants identified cyclical patterns from histories (which were available from the first but which now took on new meaning) and set up predictive data structures, identifying true cycles and—most importantly—coordinating cycles, so that their system brings in a bus only once to satisfy several cycles at the same time. For example, the participants quickly realized that a 15,000-mile cycle and a 30,000-mile cycle can be coordinated so that at least half the time the 15,000-mile cycle co-occurs with a 30,000-mile component cycle. The participants were then given materials to construct a maintenance allocation chart for the whole fleet over a number of months and evaluate the stress this would put on the shop. After doing this, they entered these data on an actual test region in MIDAS and created and assigned the preventive work orders according to this schedule.

During the second day, the participants completed their data entry and printed out their work assignment sheets and work orders. They ran their miniature depot again using MIDAS and saw the difference in profits and ease of workflow. Usually after only 5 "days," the team could afford to buy an additional bus to add to the fleet and thereby increase their fare-box revenue.

The last activity of the workshop involved entering the data on work orders (paying attention to detailing the components, defects, and symptoms involved) and closing out both work orders and work assignment sheets. At this point, participants also learned how to get various reports they now realized they wanted, such as a 30-day history on a bus. Most participants were no longer thinking of MIDAS as a computer per se, but rather a tool for doing what they had been developing manually over a number of days. After operating as MIDAS and then with MIDAS, participants navigated through the actual system more easily, knew what to look for, and asked informed questions. Even computer illiterate individuals showed little hesitation when exploring the system on the third day.

As indicated above, about 80% of the mechanics were not native English speakers and fewer than 20% were computer literate. Many midcareer individuals had not completed high school. None wanted to attend the training and most were resistant to the idea of having to do their own data entry on the first day of training. Despite these features, the trainees mastered the system at record speed. Rather than requiring the expected 12 months for implementation, the hourly staff reached independence with the system in 2 weeks and line supervisors (who do more) took 6 weeks. The one exception occurred at a location that received classroom training but no simulation exercises (this last site was not included in our original contract scope of work). After 8 months, the implementation was declared a failure.

After examining the success of the project and comparing it to the one failure, my colleagues and I believe classroom approaches (which involve explanation, simplification, and instruction) have not worked because individuals have different prior perspectives that must each be taken into account. It seems that learning through "constructive activity" actually involves building on, or reorganizing, the way that one already understands something. Therefore, it is critical to engage prior knowledge, if only to make sure it is eventually changed or reorganized. In this context, resistance to learning may actually be best understood as the assertion of existing expertise and may actually be necessary to learning. Whenever a teacher simplifies material for his or her students, he or she is really anticipating the entry point of the learners. This method often fails with experienced workers because the entry point is not always predictable or universal (for example, *simplified* is often not helpful for those experienced in thinking through vast amounts of detail). The learners—when allowed—actually do better at breaking it down for themselves in a way that is useful to them. What looks like resistance is actually an attempt to construct an entry into a new way of understanding something by beginning with what one has. In a sense, a "wrong" idea is used as raw material for a new idea, with the challenges of the exercise acting to remold the operating knowledge of the learner.

EVALUATION

Because the participants in the exercises were depending on us to orient them to the system, we could not arrange for a true control group. Rather, those who were out of work on the first day of their scheduled exercise (due to illness, personal days, or other reasons) were rescheduled to go through the exercise after everyone else and were measured as a kind of control group until they were trained. When the system went into use, they did receive classroom-based vendor training that differed from the others' only in that they did not through our exercises. There were 12 of these individuals out of about 150 total participants at the pilot site (which is where the cognitive probes were conducted).

After six months, we conducted cognitive postprobes on samples of the participants from both the pilot site (the same people who were preprobed) and the controls. Figure 5.1 shows the cognitive battery results of the pilot trainees. While the controls produced the same profile as they (and all mechanics) had before the training, those who went through our exercises produced a mirror image result compared with their former approach. Rather than solving the scheduling and data interpretation tasks with a reactive dominant approach, they exhibited about 70% of their strategies in the proactive domain. Further, when asked about how this compared to their prior performance, most did not remember doing it another way, and several could not replicate their former solution to the problem.

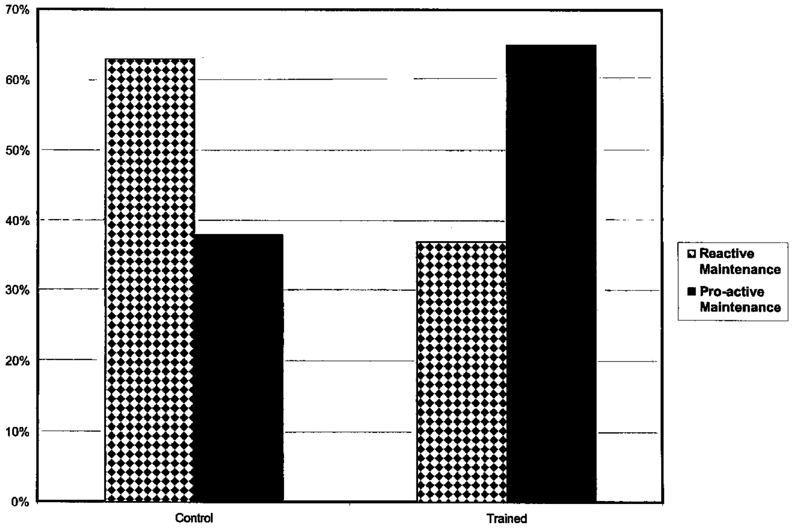

FIG. 5.1. Cognitive battery after 6 months.

Data Entry Patterns

Two measures used to determine the quality of the system use concerned the data inputted by the mechanics. I know from industry standards that coding the location of an equipment defect at a subsystem level is required for trend analysis. When users do not understand the level of detail required, they code at too general a level for the data to be useful. Figure 5.2 shows that our trainees were coding at the fourth level of the equipment template most of the time, which indicates that they are identifying a root cause component rather than a more general assembly. This is an unprecedented result in the transit industry. It indicates that the users were knowledgeable of the uses the system makes of the data and were coding appropriately. As can be seen, the control group's codes reflect a flatter, more general pattern that does not support root cause failure analysis. This indicates that the control group users do not understand the purpose of the data entry (although it was explained to them during classroom kinds of training).

Code Variation Measures

Downloads of workers' navigation through the system and data entry practices were also analyzed for component code variation and homogeneity. For the first, I measured the frequency with which any component code was chosen from a finite universe of about 2,000. In general, the data from systems such as MIDAS have been considered poor in quality, or inaccurate when the same symptom, defect, and component codes are chosen over and over because they are both general and

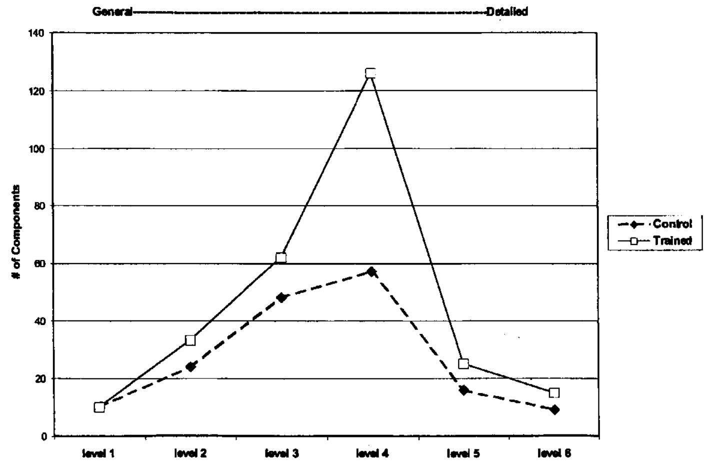

FIG. 5.2. Level of component detail.

easy to remember. An example of inputting a code would be "malfunctioning" (symptom), "broken" (defect), and "fuel system" (component). When analyzed by the system for patterns that indicate life cycles, these kinds of data are basically useless. When users understand this, they tend to code with more specific kinds information. A more detailed example of coding would be: "sporadic power surges" (symptom), "cracked" (defect), and "Injector valve-aft" (component). Simply put, a low average component code frequency indicates greater variability. Therefore, my colleagues and I looked for low hit rates per code, per user.

Figure 5.3 below shows the average frequencies per code, per person in two groups: the trainees and the 12 controls. The trainee group maintained the lowest frequency while the control group scored higher in both frequency and variance.

The low standard deviation among the trainees suggested a homogeneity effect. As a test of homogeneity, we conducted a Scheffé test of the standard deviation. The significant Scheffé indicates there was considerable within-group consistency in the type of detail entered even though the input was much more complex.

These data were initially collected on the 200 maintenance personnel at the pilot site. The striking success of the exercise was used to make a compelling case for deploying MIDAS among all 3400 frontline maintenance workers. For the full rollout, the training exercise was conducted with over 3600 people in a period of about a year. The evaluation process continued as well, with monthly downloads of data entry patterns being analyzed for each mechanic at each site as each location began using the system. The data were analyzed for two years after the rollout began and no degradation in quality was seen. Further, once the full scale implementation was underway, we also evaluated MIDAS success by measuring Mean Distance Between Failure (MDBF).

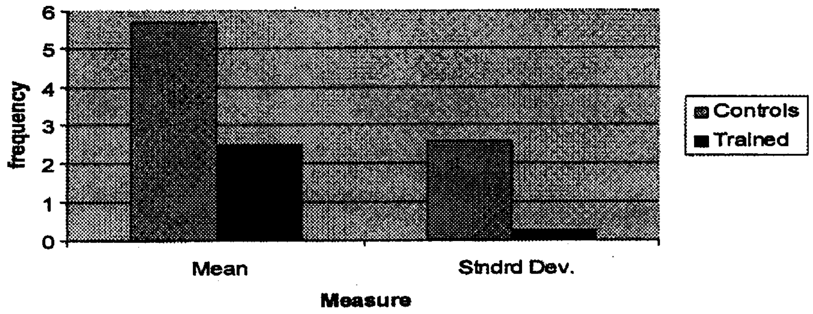

FIG. 5.3. Component frequency.

MDBF

MDBF is obtained by taking the number of in-service vehicle failures divided by the in-service distance traveled.

Rises in MDBF mean that the revenue earning asset is out earning money and is not incurring maintenance cost from repair labor. Simply put, therefore, the higher, the better. As can be seen from Figure 5.4, the MDBF rose system wide at the same rate that the MIDAS mechanics participated in the exercise. The savings from the increased MDBF are estimated to be about $40 Million. The savings in field supervisor time (handling the return of broken-down buses) is estimated to be 208,000 hours times a fully loaded hourly rate of $70, or $14,560,000. These numbers represent the financial benefits that occurred in the first year, before there were enough data collected to support the trend analysis needed for true preventive replacement based on life cycles. That analysis is just beginning.

DISCUSSION

The main point is simple: the way that workers understand their work and their role in the workplace acts as a kind of operating theory that affects how they do their job and what actions they choose at various decision points. Workplace

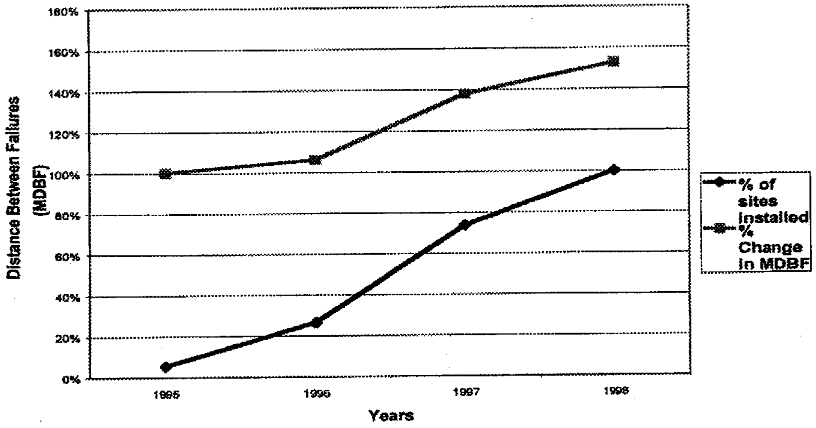

FIG. 5.4. Performance Increase with computerized maintenance management.

culture manifests as a set of activities, practices, and procedures that have evolved historically in response to having to accomplish important goals with specific resources. Ways of thinking also develop in response to goals and existing practices but in different ways depending upon experience. For new workers, the goal is to become a participant in a set of practices already in place for others; in a sense, their struggle is to understand the workplace in the same way as everyone else. For experienced workers, the goal is to get the work done and to contribute to maintaining, evolving, and refining the practices that accomplish this.

Only experienced workers can recognize the new opportunities that new resources provide, because they already intuitively contribute to practice and see the relationship between specific practices and specific outcomes. The barrier, of course, is that they have a goal-practice framework already in place, and struggle with the initial "gestalt switch" required to use their rich experiential knowledge in a new way made possible by a complex technology. Normally, their contribution is more gradual than complex technologies will tolerate. However, as the MIDAS project has shown, there may be more than one way to enlist the considerable knowledge capital of an organization for an aggressive change in goals.

As already indicated, after the MIDAS system was put into place, MDBF rose dramatically and, in truth, I do not really know why. Furthermore, system notes appeared that were increasingly in the private language of mechanics. In a sense, it might be said that the language of mechanics itself developed somewhat as a result of a wider reading and writing audience sharing the same conversation. In any case, as the mechanics grew more comfortable with the system, it became harder for us to know what they were doing with it. In other words, they grew beyond us in their understanding of what the data were saying.

There is an excerpt from the "notes" section of a work order at the New York City Transit Authority:

> Worked on 7016, which came from ENY minus the following items: one entrance door partition, one station upwright and grabrail, one dome light partition cover and front dest sign lock. Remove dest compart locks from bus 7033—which is waiting for other parts—to meet req. All other items listed were obtained from spare buses at yard. Tap-out damage Riv-nuts installed new ones on same. Interior close tbc.

There are two striking features in this passage. The first is the admission of cannibalism, stealing parts from one bus to get another into service, a practice that could have led to dismissal before MIDAS was implemented. Using MIDAS, mechanics soon realized that indicating parts shortages in the components fields helped MIDAS correct parts ordering forecasts, making cannibalizing unnecessary. Telling other mechanics where the stolen parts came from helped them address missing parts problems in the cannibalized buses later. Other notes helped the mechanic on the next shift begin where the previous one had left off. The second striking feature is that I cannot decipher very clearly what is going on. In other words, the notes are not useful to nonmechanics. This trend became stronger as the system produced increasing financial benefits. What happened here is that the system has become a tool for the mechanics, and perhaps the absence of this practice had been the problem all along with failed technologies.

Possbily, the focus of design should be making new goals more visible to workers, with technologies acting as tools for examining how performance compares with or impacts on target goals. I suspect that MDBF is increasing at NYCTA because the mechanics have found a way to use repair histories (which can be accessed instantly by any user on any bus or system of a bus from any terminal) as feedback on their own diagnostic and repair decisions. Mechanics can now address a symptom a particular way and then follow the performance of a vehicle afterward to evaluate the result. The most important thing about this shift may be that it is entirely user controlled and initiated.

As with most works in progress, this issue will involve further study. Just as training turned out not be a straightforward issue, the feedback that supports learning has, and will continue to have, unforeseen subtleties as well.

REFERENCES

Boldt, R. (2000). *TCRP Synthesis of Transit Practice #35, Information Technology Update.*

Boldt, R. (1994). *TCRP Synthesis #5, Management Information Systems, A Synthesis of Transit Practice.*

Capelli, P., Bassie, L., Katz, H., Knoke, D., Osterman, P., & Useem, M. (1997). *Change at work: How American industry and workers are coping with corporate restructuring and what workers must do to take charge of their own careers.* New York: Oxford University Press.

Chamberlain, E. S. III, & DiBello, L. (1997). Iterative design and implementation: A model of successful scheduled maintenance technology deployment. *Transportation Research Record.* August 1997, no. 1571. pp. 42-49.

Chi, M., Glaser, R., & Farr, M. (1988). *The nature of expertise.* Hillsdale, NJ: Lawrence Erlbaum Associates.

Cole, M., Engestrom, Y., & Vasquez, O. (1977). *Mind, culture and activity: Seminal papers from the laboratory of comparative human cognition.* New York: Cambridge University Press.

Davydov, V. V. (1988). Problems of developmental teaching: The experience of theoretical and experimental psychological research. *Soviet Education 30.*

DiBello, L. (1997a). Measuring success in non-trivial ways: How we can know that a DSS implementation has *really* worked. In *Proceedings of the 1997 IEEE International Conference on Systems, Man and Cybernetics* 3(5) pp. 2204-2209.

DiBello, L. (1997b). Exploring the relationship between activity and the development of expertise: Paradigm shifts and decision defaults. In C. Zsambok & G. Klein (Eds.), *Naturalistic decision making* (pp. 163-174). Mahwah, NJ: Lawrence Erlbaum Associates.

DiBello, L., & Spender J-C. (1996). Constructive learning: A new approach to deploying technological systems into the workplace. *International Journal of Technology Management 11*(7/8), pp. 747-758.

DiBello, L. (1996). Providing multiple "ways in" to expertise for learners with different backgrounds: When it works and what it suggests about adult cognitive development. Special Issue of the *Journal of Experimental and Theoretical Artificial Intelligence* 8, pp. 229-257.

DiBello, L., & Glick, J. (1993). *Technology and minds in an uncertain world.* Paper presented at the Annual Conference of the National Society for Performance and Instruction, April 12-16, 1993, Chicago, Il.

DiBello, L., & Kindred, J. (1992). *Understanding MRPII systems: A comparison between two plants* (Tech. Rep. for Laboratory for Cognitive Studies of Work). New York: City University Graduate School.

DiBello, L., Kindred, J., & Zazanis, E. (1992). Third Year Annual Report prepared for the Spencer Foundation by Laboratory for Cognitive Studies of Work. New York: City University Graduate School.

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

Dreyfus, H. L., & Dreyfus, S. E. (1986). *Mind over machine: The power of human intuitive expertise in the era of the computer.* New York: The Free Press.

Harrington, J. (1974). *Computer integrated manufacturing.* New York: Industrial Press.

Harvey, D. (1990). *The condition of postmodernity: An inquiry into the origins of cultural change.* Oxford, England: Blackwell.

Hedegaard, M. (1988). *The zone of proximal development as a basis for instruction.* Aarhus, Denmark: Institute for Psychology.

Heidegger, M. (1977). *The question concerning technology.* New York: Harper & Row.

Hoffman, R. R., Crandall, B. E., & Shadbolt, N. R. (in press). A case study in cognitive task analysis methodology: The critical decision method for the elicitation of expert knowledge. *Human factors.*

Klein, G. (1999). *Sources of power: How people make decisions.* Cambridge, MA: MIT Press.

Klein, G., Calderwood, R., & MacGregor, D. (1989). Critical decision method for eliciting knowledge. *IEEE Transactions on Systems, Man, and Cybernetics 19*(3).

Martin, L. M. W. (forthcoming). Introduction. In V. V. Rubtsov (Ed.), *Learning in children: The organization and development of cooperative actions.* Commack, NY: Nova Science Publishers.

Martin, L. M. W., & Scribner, S. (1991). Laboratory for cognitive studies of work: A case study of the intellectual implications of a new technology. In *Teachers College Record 92.*

Moll, L. C. (Ed.). 1990. *Vygotsky and education: Instructional implications and applications of sociohistorical psychology.* New York: Cambridge University Press.

Murnane, R. J., & Levy, F. (1996). *Teach the new basic skills: Princples for educating children to thrive in a changing economy.* New York: The Free Press.

Nardi, B., & O'Day, V. L. (1999). *Information technologies: Using technology with heart.* Cambridge, MA: MIT Press.

Norman, D. A. (1999). *The invisible computer: Why good products can fail, the personal computer is so complex and information appliances are the solution.* Cambridge, MA: MIT Press.

Orasanu, J., & Connolly T. (1993). The reinvention of decision making. In G. A. Klein, J. Orasanu, R. Caderood, & C. Zsambok (Eds.), *Decision making in action: Models and methods* (pp. 3-20). Norwood, NJ: Alex.

Rogoff, B., & Lave, J. (Eds.). 1984. *Everyday cognition: Its development in social context.* Cambridge, MA: Harvard University Press.

Postman, N. (1993). *Technopoly: The surrender of culture to technology.* New York: Vintage Books.

Scribner, S., DiBello, L., Kindred, J., & Zazanis, E. (1992). *Coordinating knowledge systems: A case study* (Res. Rep. for the Spencer Foundation). New York: Laboratory for Cognitive Studies of Work, City University of New York Graduate School.

Scribner, S., Sachs, P., DiBello, L., & Kindred, J. (1991). *Knowledge Acquisition at Work* (Tech. Rep. No. 22). New York: National Center on Education and Employment, Teacher's College, Columbia University.

Scribner, S., Tobach, E., & Falmagne, R. (Eds.). 1996. *Mind and social practice: Selected writings by Sylvia Scribner (Learning in Doing).* Cambridge, England: Cambridge University Press.

Sternberg, R., & Wagner, R. (1986). *Practical intelligence: The nature and origins of competence in the everyday world.* Cambridge, England: Cambridge University Press.

Timms, H. L., & Pohlen, M. F. (1970). *The production function in business: Decision systems for production and operations management (3rd ed.).* Homewood, IL: Irwin.

Winograd, T., & Flores, F. (1987). *Understanding computers and cognition.* New York: Addison Wesley.

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

Vygotsky, L. S. (1987). Thinking and speech. In R. Rieber & A. S. Carton (Eds.), *The collected works of L. S. Vygotsky: Vol. 1. Problems of general psychology.* New York: Plenum.

Zuboff, S. (1988). *In the age of the smart machine: The future of work and power.* New York: Basic Books.