

GROUP LEARNING

JEANNE M. WILSON
The College of William & Mary

PAUL S. GOODMAN
Carnegie Mellon University

MATTHEW A. CRONIN
George Mason University

We clarify the construct of group learning, encouraging new directions for research. Definitions of group learning vary considerably across studies, making it difficult to systematically accumulate evidence. To reconcile disparate approaches, we first present a set of features for distinguishing group learning from other concepts. We then develop a framework for understanding group learning that focuses on learning's basic processes at the group level of analysis: *sharing*, *storage*, and *retrieval*. By doing so, we define the construct space, identify gaps in current treatments of group learning, and illuminate new possibilities for measurement.

- In an eight-person product development team, one member from Engineering learns a new method for three-dimensional graphing and starts using it in her rough product designs. As a result, the team's development costs decrease.
- Based on its experience with the Love Me cybervirus, a national internet security team agrees that, in the future, the team should wait to send out alerts until it has a tested fix to recommend. Seven weeks later, when the Me Too bug strikes, the team delays sending out an alert for an extra four hours while it develops a patch. This response is roundly criticized in the internet security community.

Does either of these examples represent group learning? The answer is difficult to determine because the literature is so inconsistent about what constitutes group learning. Each of the above examples would be endorsed by some researchers but discounted by others. Because agreement on the definition of a construct is a prerequisite to effectively testing ideas about it (Rosenthal & Rosnow, 1991), a unifying view of group learning would help advance our understanding of this important phenomenon. In this paper we clarify the construct of group learning,

focusing on four objectives. First, we present a set of features differentiating group learning from other concepts. Second, we identify gaps in the research literature on group learning. Third, we present an illustrative set of new propositions from our conceptualization of group learning. Fourth, we outline some implications for research and methods.

Several factors make the need for clarity about group learning increasingly important. First, because groups have become an important building block of organizational effectiveness over the past twenty years, understanding whether and how groups learn is important for predicting organizational performance. Second, group research has shifted from primarily focusing on group effectiveness models to understanding critical group processes, one of which is group learning (Argote & McGrath, 1993). Third, there has been a growing body of theory and empirical research on group learning, but, as in most early stages of research, definitions of the construct have varied considerably across studies, and there are gaps and ambiguities in those conceptualizations (cf. Snyder & Gangestad, 1986). These discrepancies have led others to note that "the group learning literature suffers from the problem of insufficient cohesion. Greater consensus in the development of a theoretical framework would be helpful in generating more empirical research" (Mohammed & Dumville, 2001: 97).

We thank Ed Conlon, Amy Edmondson, Carrie Leana, Michelle Marks, Don Moore, Dick Moreland, John Levine, Brandi Pearce, Denise Rousseau, and Susan Straus for their comments on earlier versions of this paper.

CURRENT STATE OF RESEARCH ON GROUP LEARNING

A review of the current literature on group learning brings this lack of agreement into focus. We examined the existing literature on group learning, including studies that either purported to be about group learning or actually measured group learning, even if done under a different label (cf. Argote, 1996; Argote, Beckman, & Epple, 1990; Blickensdorfer, Cannon-Bowers, & Salas, 1997; Carley, 1992; Edmondson, Bohmer, & Pisano, 2001; Ellis, Hollenbeck, Ilgen, & Porter, 2003; Hollingshead, 2001; Lant, 1992; Moreland, Argote, & Krishnan, 1998; Van der Vegt & Bunderson, 2005; Wong, 2004; Zellmer-Bruhn, 2003). Group learning has been defined as "the activities through which individuals acquire, share and combine knowledge through experience with one another" (Argote, Gruenfeld, & Naquin, 2001: 370). Edmondson has defined group learning as "an ongoing process of reflection and action, characterized by asking questions, seeking feedback, experimenting, reflecting on results and discussing errors" (1999: 353). Others have suggested that learning is primarily a process of error detection and correction (Argyris & Schön, 1995) or that group-level learning is primarily about the processes of interpretation and integration (Crossan, Lane, &

White, 1999). Additional definitions and related constructs are displayed in Table 1.

A review of these definitions reveals little consensus. Some focus on individuals learning in a group, while others focus on a team's level of collective knowledge. Some focus on processes, while others examine outcomes. All of these definitions raise questions about (1) the appropriate level of analysis, (2) critical learning processes, (3) distinguishing learning outcomes from other constructs such as performance, and (4) changes in group learning over time.

Many scholars have made the case that a theory or research model must contain an explicit description of the levels to which generalization is appropriate (Rousseau, 1985) and that critical problems result when the level of the theory is inconsistent with the prevailing level of measurement or statistical analysis (Klein, Danse-reau, & Hall, 1994). Some treatments of group learning confuse levels of analysis by not distinguishing "individual learning in the context of groups" from "group-level learning." What we mean by this is that individuals can learn within the context of a group, and their learning may improve the group's performance, but it still is individual learning unless shared by members of the group. If an individual leaves the group and the group cannot access his or her

TABLE 1
Definitions of Group Learning

Paper	Definition
Edmondson (1999: 129)	An ongoing process of reflection and action, characterized by asking questions, seeking feedback, experimenting, reflecting on results, and discussing errors or unexpected outcomes of actions
Argote, Gruenfeld, & Naquin (2001: 370)	The activities through which individuals acquire, share, and combine knowledge through experience with one another
Edmondson (2002: 129)	A process in which a team takes action, obtains and reflects on feedback, and makes changes to adapt or improve
Sole & Edmondson (2002: S18)	The acquisition and application of knowledge that enables a team to address team tasks and issues for which solutions were not previously obvious
Ellis, Hollenbeck, Ilgen, Porter, & West (2003: 822)	A relatively permanent change in the team's collective level of knowledge and skill produced by the shared experience of team members
Gibson & Vermeulen (2003: 203-204)	The exploration of knowledge through experimentation, the combination of insights through reflective communication, and the explication and specification of what has been learned through codification
Gruenfeld, Martorana, & Fan (2003: 46-47)	The acquisition, persistence, diffusion, and depreciation of group knowledge
London, Polzer, & Omeregic (2005: 114)	The extent to which members seek opportunities to develop new skills and knowledge, welcome challenging assignments, are willing to take risks on new ideas, and work on tasks that require considerable skill and knowledge

learning, the group has failed to learn. As with other group-level constructs, group learning should be an emergent property of the group exerting influence beyond the individual members involved in the original learning process (Morgeson & Hofmann, 1999).

It was apparent from our review that most current conceptualizations of group learning are not explicit about basic learning processes. Existing definitions of group learning contain process verbs, such as "share," "reflect," "feedback," and "interpret," all of which can facilitate learning at the group level. Groups that seek feedback and reflect on errors are more likely to learn (Edmondson, 1999). Nonetheless, research on group-level learning has paid little or no attention to basic learning processes, such as how information is encoded, stored, or retrieved. This is a fundamental problem. The individual learning literature (Anderson, 2000) and recent work on organizational learning (Argote, 1999) show that the basic processes of storing and retrieving new routines are central to learning. We analyzed papers on learning cited in one of the most recent reviews of group research (Ilgen, Hollenbeck, Johnson, & Jundt, 2005). Only 10 percent of the papers in the group learning section of this most recent review were explicit about these processes.

A third problem with current conceptualizations of group learning is a failure to distinguish learning as an outcome from other constructs. Only 20 percent of the empirical papers in the group learning section of the Ilgen et al. review (2005) actually examined group learning rather than other constructs. One of the constructs most commonly confused with group learning is performance, despite a long-standing conceptual distinction in the literature (Tolman, 1932). Still, authors equate learning with performance (Fiol & Lyles, 1985), assuming that no change in performance means that learning did not take place (Cook & Yanow, 1993). We argue that learning may have occurred, even when there was no change in a group's overall performance. For example, the group may have learned something but may not have had an opportunity to apply the learning in a way that would change its performance. Conversely, performance can change without any learning actually taking place—for example, when the environment changes (e.g., when a product development team is able to reduce time to market because a

supplier delivers a key component early). Finally, learning does not always result in positive outcomes. Research on group learning needs to account for the possibility of dysfunctional learning, as in the case of superstitious learning, where a group learns a false connection between its actions and some outcome (Levitt & March, 1988).

Finally, many treatments of group learning do not examine changes over time, even though many fundamental aspects of learning, such as practice or forgetting, occur over time. Learning is a necessarily dynamic construct. Without a change in the repertoire of potential behavior, there is no learning, and in order to assess change, one must consider the role of time. Our coding of the research cited in the Ilgen et al. review (2005) showed that fewer than 30 percent of the studies cited in the section on team learning actually measured or conceptualized the construct of group learning over time. Time is a critical aspect of the definition of group learning, and it helps to distinguish group learning from other constructs, such as group decision making.

OUR APPROACH

Given the current limitations in the literature, we propose four criteria to be considered in a theory of group learning:

1. Level of analysis: Learning must be at the group level of analysis. Hence, we propose as a *definition* that group learning represents a change in the group's repertoire of potential behavior. We are explicitly stating that the theory, measurement, and analysis of group learning should focus on changes in the group's repertoire (cf. Klein et al., 1994). An aggregation of what individual members learn does not constitute group-level learning.
2. Fundamental processes: We propose that the processes inherent in the construct of group learning include *sharing*, *storage*, and *retrieval* of group knowledge, routines, or behavior. The processes of sharing, storage, and retrieval are the basic elements or mechanisms of the learning process (Hinsz, Tindale, & Vollrath, 1997). In general, current research on group learning does not explicitly deal with the processes of sharing, storage, and retrieval or their interrelationships. Attention to processes is important because it helps us understand not only why constructs come about (Whetten, 1989)

but also the systematic reasons for a particular occurrence or nonoccurrence of a phenomenon (Sutton & Staw, 1995)—in this case, of group learning.

3. Learning as an outcome: Our conceptualization treats learning as an outcome—specifically, a change in the range of a group's potential behavior, following Huber (1991). Any change in the group's *range of potential behavior*, whether or not it is manifested in externally observable behavior, constitutes evidence of group learning. Learning as an outcome should be distinguished from other criterion variables, such as performance or decision making.
4. Time: Our definition explicitly incorporates time by requiring a *change* in the group's repertoire of potential behavior over some interval. This feature of our definition allows us to distinguish group learning from other group-level phenomena, such as decision making, shared mental models, and problem solving.

In advancing "a change in the group's repertoire of potential behavior" as a definition of group learning, we implicitly adopt a cognitive approach to this construct. This is not only consistent with general trends in the field of organizational behavior (Ilgen & Klein, 1989) and groups research in particular (Moreland, Hogg, & Hanes, 1994) but is appropriate because group learning is essentially about the internal and external manifestations of information processing.

Applying our criteria to the examples offered at the beginning of the paper clarifies the meaning of group learning. The case of the engineer who learns a new procedure does not meet the criteria for group learning, even though the new procedure she starts using makes the group perform better. When one person in a group learns something that is not shared with other members of the group, as in this example, this constitutes individual learning, not group learning. In the case of the incident response team, however, learning does occur, because the range of the group's behavior changes (i.e., delaying alerts until a fix is available becomes part of the group's repertoire). As this example reveals, however, not all learning results in positive changes in performance. This example helps to illustrate that, taken alone, changes in performance are not accurate indicators of group learning.

To highlight specific points about group learning, we rely on a group situation that we

observed over a period of three years at a national computer emergency response center. The purpose of the center is to respond to threats or attacks on the internet infrastructure (such as a widespread worm or virus). Whenever such attacks occur, an incident response team is formed to deal with the attack. This team works with external experts to identify a fix or patch, keeps the broader community informed about the incident, and generally serves as an unbiased source of information (not affiliated with any software providers). At the time of our observation, the core members of this team included Kyle (the team leader), Aaron, Seth, Alex, Chris, Mitch, and Sam. Depending on the nature of the attack, the team can expand to include members of other incident response teams, experts at vendor sites (when vulnerabilities in their software are being exploited), and government officials. In a serious attack, the team may field hundreds of emails about the incident, work around the clock for several days, and deal with dozens of national media inquiries about the incident.

Building on this case and previous research on group learning, we outline an approach to group learning that specifies the necessary and sufficient conditions for group learning to occur. Our approach to group learning clarifies theoretical issues about level of analysis while disentangling group learning from other related constructs, such as group performance or decision making. We also focus attention on critical but understudied topics, such as negative learning, and neglected processes, such as storage and retrieval. Previous definitions of group learning (Argote et al., 2001; Crossan et al., 1999; Edmondson, 1999) have focused primarily on the process of sharing in group learning; as a result, most of the empirical research to date samples only part of the total construct space. Finally, our approach highlights opportunities for new methods in studying group learning.

BASIC FEATURES OF GROUP LEARNING

Sharing

We define sharing as the process by which new knowledge, routines, or behavior becomes distributed among group members and members understand that others in the group possess that learning. Group learning must be shared, taking on structural properties and ex-

erting influence beyond the individuals who constitute the collective, before it becomes a legitimate group construct (Morgeson & Hofmann, 1999). An example of sharing can be found in Devadas and Argote (1995), who showed that when a group embedded knowledge in its roles and procedures, this learning managed to persist, even in the face of extensive turnover within the group.

There are at least three stages in the development of shared knowledge. In the first stage an individual member's repertoire changes to incorporate some new knowledge, routine, or behavior, x . For example, in the incident response teams we studied, a major responsibility is disseminating accurate information about an attack. And, in this case, an example of a change in knowledge (x) was Alex's realization that "we need to use a set of Frequently Asked Questions (FAQs) because 95 percent of the questions from the media are the same." This does not change the group's repertoire, since only one person possesses the learning. If that person leaves the group, the knowledge is lost.

At the second stage in the acquisition of shared knowledge, imagine that several of the other group members have the same understanding about x , in this case the need for FAQs to reduce redundancy when responding to media inquiries. However, the other members each acquired the same knowledge of how to handle redundant media questions independently, and each person thinks that he or she is the only group member who knows about the redundancy. In such a situation, although each member possesses the learning, there still is no shared understanding at the group level about x , the need for a set of FAQs—no group-level learning has occurred. The learning would not be enacted in a situation that required the other members to share the knowledge, because each person thinks he or she is the only group member who knows about the redundancy.

In the third stage a shared understanding of x allows knowledge to be transferred to new group members and decreases the probability that the learning will be lost over time. We stipulate that group learning occurs when the members possess both the knowledge (in this case, that using FAQs will help them deal with redundant media questions) and an understanding (either explicit or tacit) that others have the same knowledge and it is a property of the

group. This means that a new group repertoire now exists, and it is independent of any particular individual. The process of sharing also serves to legitimate the knowledge for the group. When this happens, the learning becomes a group-level construct that can survive the turnover of any members. More important, people in the group are able to anticipate how other members will respond in certain situations and to act accordingly, a crucial capability when tasks are interdependent.

Shared information is also mutually enhancing for group members—validating members' knowledge and helping group members relate to each other (Wittenbaum, Hubbell, & Zuckerman, 1999). When sharing is complete, these factors are present and, consequently, learning occurs at the level of the group. From this conceptualization of sharing, we propose that the depth and breadth of sharing about any given learning are directly related to the probability of group retrieval of that learning. "Depth of group learning" refers to the level of detail about any particular knowledge, routine, or behavior that is shared by members of the group (such as under what circumstances it is appropriate to apply the new learning). For instance, by focusing on when to use the new knowledge or routine, the retrieval cues become clearer. "Breadth of group learning" refers to the distribution of the learning within the group—how many members share understanding about the new knowledge or routine. The more group members share the learning, the greater the probability it will be retrieved in the future. Greater depth and breadth of shared understanding should be associated with stronger encoding of that learning. Accordingly, we predict the following.

Proposition 1: The depth and breadth of sharing among group members about any given knowledge, routine, or behavior improve storage and retrieval of that information.

Some current research on group learning does address the concept of sharing. Edmondson (2002) theoretically develops the specific processes by which sharing can happen—that is, by having the group take time to raise questions and to reflect on what happens as members work. She also discusses how a lack of sharing disrupts learning, as well as situations that in-

hibit sharing, such as when the group becomes too absorbed in what it is doing to take time for reflection and, thus, loses an opportunity for sharing. Laughlin and Shupe (1996) provide a good example of empirically validating whether sharing occurred when assigning groups the task of learning the correct rule for partitioning a deck of cards. By observing and modeling the processes each group used to reach a decision (e.g., voting), they could track exactly when and how individual learning became shared within the group.

For sharing to be effective, several subprocesses must take place. The group must (1) focus its attention on the information that is to be learned, (2) develop a shared understanding of the specific learning, and (3) marshal some shared understanding about using this new knowledge in the future. Most research has focused on how groups develop a shared understanding of some event (typically through discussion or observation; cf. Gruenfeld, Martorana, & Fan, 2000; Moreland & Myaskovsky, 2000). Measures of the constructs of reflexivity (Schipper, Den Hartog, & Koopman, 2007) and team learning behavior (Van der Vegt & Bunderson, 2005) are also closely related to this aspect of sharing. We know much less about the first and third subprocesses—that is, why groups focus on a particular item and how they develop a collective understanding of how and when to use the knowledge in the future. Below we show how concentrating on these neglected aspects of sharing leads to interesting research possibilities.

We start by considering how a group's focus of attention affects what it learns. Why do groups learn some things easily but repeatedly fail to learn other—often important—items? As an example, the computer emergency response teams we observed seemed quite capable of learning how to do technical processes differently but repeatedly failed to learn how to improve their internal team processes. For instance, they frequently adjusted their automated email system to send more timely alerts to the internet community but continually failed to change their staffing patterns in order to respond more effectively to serious attacks.

Biases in a group's focus of attention may provide one overlooked explanation for the observation that organizational groups often seem incapable of learning how to improve their own

functioning (Argyris, 2003; Tjosvold, Yu, & Hui, 2004). One of the contributing factors may be the phenomenon of "team halo"—the observation that both groups and individuals are more likely to attribute failures to individuals and successes to the group (Naquin & Tynan, 2003). In other words, both groups and individuals have difficulty thinking of "the group" as the source of problems. So when members of a group are faced with a problem and are inclined to think about how to change in the future, they are predisposed to think of *individual-level*—not *group-level*—actions. We observed this phenomenon in our own study of the computer emergency response teams. In group discussions regarding their performance on recent incidents, team members would be more likely to focus on individual-level changes ("Next time I will notify the Australians first") than on group-level changes ("We should hold press conferences whenever the incident spreads beyond 500 users"). This is one reason that group learning can be especially difficult.

Proposition 2: Groups are more likely to focus their attention on changes to individual-level rather than group-level routines, decreasing the probability that group learning will occur.

Even if the group effectively focuses its attention and develops a shared understanding of what it needs to learn, it still may not achieve a shared intention to behave differently in the future (the third component of sharing). Psychological safety, characterized by a willingness to confront one another and an openness to experimentation, is one predictor of group members sharing an intention to change their repertoire of behaviors. In her study of learning patterns in an office products company, Edmondson (2002) describes two teams, the Strategy team and the Radar team, both of which failed to change their repertoires because they avoided conflicts associated with committing to a specific future direction. In comparison, teams that were characterized by higher levels of psychological safety were able to commit to using their knowledge in the future. We propose that, in addition to psychological safety, there are also other variables at the group level of analysis that affect the extent to which groups will share a commitment to change their own routines. Factors such as collective efficacy and the presence of group

goals may increase the probability that groups will focus on future applications of new knowledge or routines.

Proposition 3: Groups with higher levels of collective efficacy will be more likely to share a commitment to changing their own routines.

Proposition 4: Group discussion about performance discrepancies that reflects past, present, and future scenarios increases the probability of group learning.

We recognize that the notion of how shared an understanding needs to be within a group is a complex question and has been the subject of debates in the literature on shared mental models (Cannon-Bowers, Salas, & Converse, 1993; Klimoski & Mohammed, 1994; Levine, Resnick, & Higgins, 1993). One issue is how many members have to share knowledge before group learning occurs. Shared learning in groups is not an all-or-nothing phenomenon. It is analogous to partial learning in individuals—when a person learns some, but not all, of what he or she needs to know about a particular issue. Another question is whether members may possess slightly different variations on the same knowledge yet still call it shared (in our computer emergency response teams, think of a set of FAQs versus a set of automated responses versus a set of talking points for the hotline staff). For example, in Edmondson et al. (2001), one member claimed that the group had learned to be less hierarchical because team members were referring to each other by their first names. Without asking other members if they shared this perception, however, we do not know whether they had different interpretations of what it means to address people by their first names.

Sharing is a key feature for defining group-level learning that also shapes two other important learning processes: knowledge storage and retrieval. Future research on group learning needs to be more explicit about why groups focus on certain objects in their environment but not others, and how groups develop a shared understanding about using their knowledge in the future.

Storage

Another feature of group learning is that the change in the group's repertoire needs to be stored in memory. Storage is necessary for learning to persist over time, so much so that others have defined learning as the exploitation of stored knowledge (Moorman & Miner, 1998). In our discussion of storage, we focus on how knowledge that has been learned by the group comes to be stored and retained in memory repositories or storage bins used in group-level learning.

A review of the group learning literature reveals little attention to storage processes or memory systems. The notable exception is work on transactive memory. That research focuses on who knows what within the group, how that information is acquired, and the consequences of levels of transactive memory for group functioning (Moreland et al., 1998). A growing body of work suggests that as group members gain experience with one another and gather knowledge of their fellow members' competencies, a variety of group outcomes improve (e.g., quality, satisfaction).

Unfortunately, the research literature on group-level storage is otherwise limited. Researchers primarily have focused on a single repository (human memory) and only considered the types of knowledge that are largely explicit and concrete, rather than implicit or tacit. The interaction between type of knowledge and repository remains unexplored. A variety of repositories (e.g., human memory, computer databases) and different types of knowledge (i.e., tacit or explicit) must be taken into account in order to fully understand how groups store knowledge. One unexplored area is the fit between different types of repositories and different types of knowledge and the implications of fit (or lack thereof) for group-level learning.

Under the heading of storage, there is little group research about retention—the persistence, decay, or distortion of stored knowledge, routines, or behavior over time. Because most experimental studies of group learning have taken place within a single session (e.g., Mathieu, Heffner, Goodwin, Salas, & Cannon-Bowers, 2000) or over the relatively short span of a week (e.g., Moreland et al., 1998), they provide little opportunity to test the persistence or accuracy of group knowledge over time. Since most

organizational actions and decisions rely on knowledge that has been stored for more than a few hours or a week, it seems imperative to understand the effect that time has on a group's stored knowledge, as well as the decay rates of different types of repositories.

How can we expand our theoretical understanding of group storage or memory processes? In this section we consider the common features of storage repositories that are available to groups and then the differences in these storage repositories. We also examine how the type of storage repository interacts with characteristics of the learning itself, and how different features of groups affect their storage practices. Finally, we consider how time and the group's external environment affect group storage through practice.

Because groups have access to a greater range of storage repositories than individuals, research is needed about the advantages and liabilities of each type of repository. Memories of group members constitute the most obvious group repository. This has been the focus of the research on transactive memory. In this case, creating a division of labor among the members in terms of who knows what illustrates one power of storage at the group versus individual level. Another type of repository includes formal group memory systems that emerge from groups' information technology structure. Shared databases, bulletin boards, and expert systems are examples of this type of repository (Olivera, 2000). Finally, structural storage repositories such as standard rules, procedures, and cultural artifacts can store group knowledge (Argote, 1999). Acknowledging these multiple repositories is important, because they represent different systems for storage and different functionalities for acquisition, retention, and retrieval.

Group storage repositories have a number of common features that impact group-level learning. Indexing, filtering, and maintenance functions are important components of any storage system (Olivera, 2000). Good indexing systems facilitate both where information is stored and how it is retrieved. Filtering is a process that screens out irrelevancies before information is stored. Maintaining a memory system refers to updating information, deleting obsolete data, and so forth. We expect that these features of storage systems will affect both the use and the

utility of the storage process. We propose that group storage systems with indexing, filtering, and maintenance capabilities will be used more often than systems without those features. Similarly, groups that use storage systems with strong indexing, filtering, and maintenance capabilities will ultimately exhibit higher rates of learning.

Indexing, filtering, and maintenance of stored group memories can be much more complex, and therefore potentially more interesting, than corresponding processes at the individual level of analysis. For instance, we expect that the network structure of the group will affect the indexing and updating of stored memories. We predict that members with higher centrality in the group will play a stronger role in indexing than other members of the group. A member with high centrality is more likely to know where knowledge is stored in the group and to serve as a pointer to that knowledge. This also means that the loss of group members with high centrality is likely to sever the "connective tissue" that enables many group storage systems to work. For all of these reasons we need to know much more about how groups use different storage systems.

Proposition 5: Group storage systems with strong indexing, filtering, and maintenance capabilities will be used more often than systems without those features. Groups that use storage systems with indexing, filtering, and maintenance capabilities will ultimately exhibit higher rates of learning than groups that do not.

Proposition 6: Group members with high centrality in the group will be more involved in indexing stored memories than other members of the group.

There are also important differences among the types of storage repositories available to groups. In the case of group members' memories, this repository can manage relatively complex material and handle both tacit and explicit knowledge. Of course, there are limits to the information any individual can process (Simon, 1947), but a formal division of labor can enhance the total memory capacity of the group. In comparison, formal memory systems (such as data-

bases) can store large amounts of data, if they have good indexing, filtering, and maintenance processes. But these systems are better for concrete and less complex knowledge (Goodman & Darr, 1996). In structural repositories one can embed both explicit and tacit understandings and complex ideas in rules and procedures. However, future generations will have difficulty accessing the meta-ideas that lead to these new production procedures, and may therefore find it difficult to modify the rules and procedures for new contextual situations. That is, there may be a lot of know-how (both explicit and implicit) behind the creation of a new rule or learning, but typically the production rule is stored in structural repositories on its own, without any of the knowledge that led to its creation. We propose that there is an interaction between the type of knowledge and the type of group storage repository that influences the effectiveness of storage and retrieval.

Proposition 7: Knowledge or routines that are primarily explicit can be stored in any of the three types of group storage repositories. Knowledge or routines that are more tacit can be more easily stored and retrieved in human memory systems.

In addition to having access to a range of storage repositories, groups provide unique (and largely unexplored) opportunities for practice and storage. With individual-level learning we know that elaborative processing improves memory or storage through the enhancement of retrieval cues (Anderson, 2000). "Elaboration" refers to a process by which subjects create additional ways of recalling information. For instance, nonelaborative processing at the individual level, on the one hand, might involve simply reading an assigned passage about chemistry. Elaborative processing, on the other hand, might include generating questions before reading the text and drawing a concept map of the key points.

In the context of group learning, group discussion can serve as a form of elaboration or practice. In our example of the computer emergency response teams, elaboration would include a group discussion in which members would review how best to respond to an attack and consider—as a group—the circumstances under which they might or might not use a press con-

ference to disseminate information to the outside world. The more elaborated the discussion of when to use a press conference and the greater the consensus, the more likely the group is to retrieve their shared learning regarding a press conference in the future. Through elaboration, the learning becomes stronger, and it is stored with multiple group members. Indicators of elaboration could include differentiated discussions of (1) situations in which particular learning should or should not be retrieved or (2) alternative storage and retrieval mechanisms.

Proposition 8: Elaboration by the group about when and where to use learning strengthens the memory record.

Time and a group's external environment also affect group storage—through practice (reinforcement of the learning through rehearsal). The effectiveness of storage depends on how the practice schedules are distributed over time. When practice schedules are distributed over time, it may take groups longer to acquire the knowledge or routine, but the rate of forgetting may be slower than in a more massed practice schedule (Donovan & Radosevich, 1999). Also, the better the match between the distributed practice schedule and the timing of the event evoking the learning, the better the retention of the learning. In the case of the computer emergency response teams, for instance, major incidents occurred several times a year. In this situation, if the group reviewed its routines for responding to attacks on a periodic basis, it would be more likely to successfully retrieve group learning when the next attack occurred. Matching the practice schedule to the rhythm of events in their environment improves storage for two reasons. First, group members learn that there are lags (i.e., months rather than days) in evoking these new routines; second, rehearsing or practicing periodically creates additional contextual cues that can enhance retrieval processes.

Proposition 9: The more closely a group's practice schedule matches the rhythm of events in its environment, the greater the probability of retrieving the group learning.

Retrieval

The final requirement for group learning is retrieval. Retrieval means that group members can find and access the knowledge for subsequent inspection or use. It is not unusual for members of a group to think that they have stored new learning, only to discover that the group does not access it when the next opportunity to apply the learning presents itself. We observed this multiple times with the response teams, when, for instance, members shared learning about the importance of establishing a protocol for real-time updating of all team members' technical understanding of an incident as it unfolded. Even though this learning was repeatedly shared among team members and was stored in at least one formal After Action Review document and in the memories of at least four team members, core team members failed to even mention the learning, much less enact it, during subsequent incidents. Despite the importance of retrieval for group learning and the fact that retrieval has been identified as the most critical part of the learning process at the individual level (Anderson, 2000; Loewenstein, Thompson, & Gentner, 2003), this process has been largely ignored in the literature on group learning.

Few studies of group learning even discuss the process of retrieval, and even fewer actually measure whether and how retrieval occurs. One exception is Hollingshead's (1998) study of the effects of communication during learning and recall in dyads of strangers and dyads of couples. She found that when partners could communicate during the learning task, strangers actually recalled *more* words than dating couples. She suggested that communication can impede the coordination of learning new material in groups when the members try to develop new strategies for storage and retrieval that depart from their implicit knowledge about each other's relative expertise. In another study Cohen and Bacdayan (1994) focused on how retrieval in groups can be reliable but not valid. These authors described how a feature in the task environment should cue the retrieval of particular learning, arguing that when the rules change, the same feature will cause the retrieval mechanism to misfire. Their experimental manipulation demonstrated exactly this process: they changed the rules of their task and showed that

groups still retrieved the old (and now out-moded) learning.

Unfortunately, what we know about the retrieval of group learning is not only limited by the fact that there are very few studies but also by the context of the laboratory tasks. In many laboratory studies of learning, learning is not examined over time, and the researchers control the stimuli. Consequently, we know very little about whether groups will undertake a search for stored memories over time, what their search strategies will entail, and how they will respond if their search fails. We need a framework that accounts for the complexities of group learning in an organizational setting.

For a group to effectively retrieve stored knowledge, several subprocesses must take place: (1) the group or one of its members, faced with some stimulus object, must recognize the need to access stored knowledge; (2) the group, or at least one member, must identify where the knowledge is stored; and, finally, (3) the group must actually retrieve the knowledge. Eventually, we must also consider whether the group can apply the retrieved knowledge in the new situation. Although these subprocesses also occur in individual-level learning, several unique aspects of these subprocesses are critical to understanding how learning occurs at the group level. As we will explain, there are a number of reasons why individual recall may not translate into group retrieval.

In complex organizational settings, group retrieval is difficult for a number of reasons. First, time plays an important role. The longer the time period, the greater the probability of forgetting the original learning event, the more difficult it will be to reconstruct the meaning of the earlier event for the present stimulus, and the more the contexts between the earlier and present events are likely to differ. Second, in groups, the social distance between members may also interfere with retrieval. In the case of the incident response teams we studied, it may be especially difficult for Chris to help the group retrieve what Sam has stored, because Sam's knowledge is less familiar or salient to Chris. Because Sam and Chris are not close in the social network, the group learning stored in their respective memories is less accessible to each other.

A third factor is that the group may have learned subsequent knowledge, behavior, or

routines that make it difficult to access the original learning. Groups present unique opportunities for interference (the negative relationship between learning two sets of material—when learning x interferes with the learning of y). The interesting dimensions of this problem at the group level are apparent in the case of a product development team. Effective group learning of a new quality improvement process was diminished by retrieval problems. A critical mass of group members (primarily the designers) had mental models of quality that were focused more on aesthetics than on precise measurement. Even though the knowledge of the quality improvement process was shared through training and stored in manuals on everyone's desk, the principles of the new process were not retrieved or applied. Prior learning and existing group mental models about quality interfered with the retrieval of the new learning.

Proactive interference (when previous learning interferes with the retrieval of new learning) may make group learning or adaptation to new circumstances particularly difficult. To understand why this may be so, we extend well-understood principles of group information sharing to propose new ways of thinking about group learning. Group retrieval can be thought of as a sampling problem (Stasser & Titus, 1985). To change the range of potential behavior in the group, the understandings and memories of multiple team members must be updated in the same time frame. If this is not accomplished, a group attempting to retrieve learning will be more likely to retrieve old (and possibly outmoded) learning.

Extending this reasoning to our understanding of group learning has interesting implications. It suggests that the longer groups are in existence (and the stronger their established practice effects), the lower the probability the groups will retrieve new (and updated) learning. This group tenure effect may account for differences in perceptions of group learning between laboratory studies (where new groups are formed and group learning is treated as routine; Gruenfeld et al., 2000; Moreland et al., 1998) and field studies (where groups have been in existence for some time and group learning is treated as difficult and rare; Argyris & Schön, 1995; Edmondson, 1999; Gersick & Hackman, 1990).

Proposition 10: The longer group members have worked together, the stronger the established practice effects and the lower the probability of retrieving new learning.

Social processes in groups may also interfere with effective retrieval (Finlay, Hitch, & Meudell, 2000). We know that on free recall tasks, collaborating groups retrieve fewer items than the same number of individuals working in nominal groups (Basden, Basden, Bryner, & Thomas, 1997). Although there is some preliminary evidence that cognitive social loafing does not account for collaborative retrieval problems (Weldon, Blair, & Huebsch, 2000), there may be other phenomena in groups that make retrieval in social situations more difficult. The combination of status differences and evaluation apprehension (Diehl & Stroebe, 1987) may combine to cause low-status group members to withhold knowledge or cues for retrieving collective learning. It is clear that more research is needed to understand how group members retrieve and collectively evaluate memory evidence when social issues such as status, familiarity, and group faultlines can affect the outcomes.

It is interesting to consider what features of groups or group-level learning would offset some of the difficulties of retrieval and would differentiate this literature from individual-level retrieval. First, groups represent collections of individuals and are subject to division of labor, where members divide the responsibility for storing different learning. Second and less well-recognized, group members serve not only as potential repositories for group learning but also as cues for the retrieval of particular information. We know that people serve as particularly strong cues for recall (Smith & Vela, 2001). So, when the computer emergency response teams were learning to use press conferences early in the trajectory of an incident, the team leader, Kyle, spent the most time advocating the use of press conferences originally (sharing) and was implicitly designated as the repository for this learning (storage). Traditionally, researchers would have focused on Kyle's role in personally retrieving the use of press conferences as an instance of group learning during the next incident. In our observation, however, even though Kyle was present in the early deliberations during the next incident, he did not

personally retrieve the group's agreement that they would call an early press conference in major incidents. Kyle's physical presence seemed to prime another team member (Aaron) to recall the particular group learning regarding early press conferences. This example of retrieval highlights important but neglected aspects of the group learning process, beyond the explicit division of labor for the recall of group learning. Individual team members serve two roles in group retrieval: (1) as repositories for group learning and (2) as cues to search for particular knowledge or routines. Group members have been explored as repositories of group knowledge in the transactive memory literature, but the larger literature has ignored their function as cues for retrieval.

From this expanded view of the group member as both a repository of learning and as a cue for recall, we would expect that for groups operating in complex and dynamic environments, cognitive division of labor with respect to individual members explicitly responsible for storage and recall should enhance the probability of recall. At the same time, we would expect more rapid decay of learning among group members who are not explicitly responsible for either storage or retrieval of certain information that has been learned by the group. Therefore, groups with more stable membership will have more reliable retrieval processes. In groups with explicit role assignments for storage and recall, the retrieval will be affected by the presence or absence of members, something particularly important in geographically distributed groups whose members do not necessarily "see" one another while working together. These unexplored effects of group member presence on group retrieval are highlighted below.

Proposition 11: The presence of a member with stored knowledge can cue retrieval of that knowledge without any additional priming, improving the group's chances for successful retrieval.

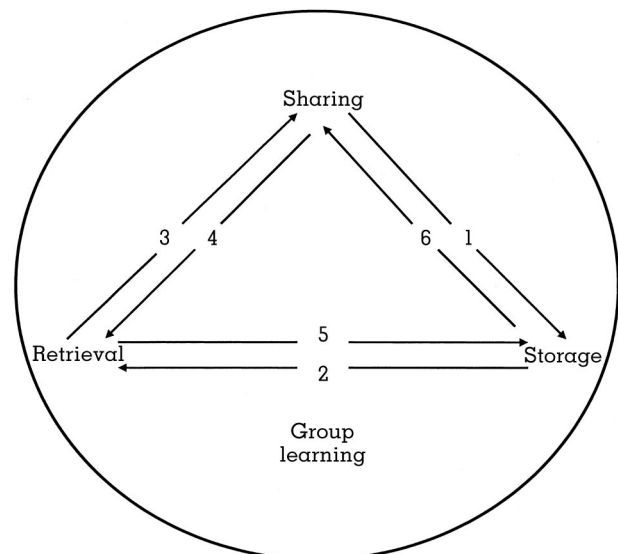
Proposition 12: The reduced "presence" of group members in geographically distributed groups will be associated with less reliable retrieval of group learning.

Overlap of Sharing, Storage, and Retrieval

Although we have dealt with sharing, storage, and retrieval separately for the sake of clarity, in practice, the three processes are intertwined. All three processes must take place for group learning to occur (represented as an equation: $GL = \text{Sharing} * \text{Storage} * \text{Retrieval}$). Without sharing there can only be individual learning in a group context. Without storage and retrieval of shared learning, the group's repertoire cannot change over time. The relationships among the three processes are illustrated in Figure 1. One advantage of our framework is that it highlights the interactions among these processes of group learning. For instance, one of the interesting possibilities we have discussed is the interaction between sharing and a group's ability to store and retrieve particular learning. One general principle about group learning is that sharing affects the robustness of the group learning through storage and retrieval. We have suggested that, in groups, increasing the breadth of sharing provides a buffer against the decay of learning because the learning is stored in multiple team members' memories (path 1 in Figure 1) and sharing creates a wider net of people able to respond to retrieval cues (path 4).

Although it may be clear that sharing affects storage (path 1) and storage affects retrieval (path 2), there are also reverse feedback loops

FIGURE 1
Interaction of Sharing, Storage, and Retrieval
in Group Learning



such that retrieval affects storage and sharing. A general principle about the relationship between retrieval and the other group learning processes is that nothing can be retrieved without at least subtly changing the content of group learning or the path to that group learning. For example, in the process of retrieving knowledge, behavior, or routines that have been learned, a group's paths to some storage mechanisms are strengthened while others decay (i.e., retrieval affects storage; path 5). As the computer emergency response team attempted to retrieve its learning about automated email updates, the group relied on Kyle to retrieve this learning, completely forgetting that the learning was stored in an After Action Review document on its intranet. The path to Kyle started out stronger for the team than its path to the intranet storage bin, and the team's continued reliance on Kyle as a retrieval mechanism further weakened its path to the intranet storage repository. This preference for retrieval paths not only affects the current instance of group learning but reinforces a general norm that the team will look to individual members to remind them of agreed-upon changes to their routines. In other words, the preference not only affects the group's current learning but also affects group norms about which storage bins it consults in the future. Thus, one of the uniquely "group" aspects of learning is that the processes of sharing, storage, and retrieval change not only the group's knowledge but the group's norms as well.

One of the other uniquely "group" aspects of learning is that the retrieval process can also serve as a sharing mechanism (path 3). In individual learning, retrieval is a largely nonverbal, cognitive process. In groups, however, retrieval often requires verbal interaction. One group member, Alex, may recognize the need to retrieve learning ("Didn't we agree on a different way to handle this?"). Another group member, Mitch, may direct the group to a particular storage bin ("Yes, I think it is in our notes from 7/18"). A third group member, Sam, may shortcut the search by saying that the group agreed to post a press release on its web site whenever an incident spreads beyond 500 end users. A fourth group member, Seth, may "second" this recall. As the group engages in sensemaking, it recreates shared understanding about the knowledge. This additional round of sharing can often strengthen the original memory record. How-

ever, it also presents an opportunity for the original learning to be distorted so that what gets retrieved (a short, one-page press release) is different from what was originally stored (a live press conference). The now-distorted retrieval can then replace or overwrite the original learning in the group's memory. So while verbal interaction upon retrieval can serve as a form of sharing, it also provides an opportunity for distortion as the group attempts to reinterpret or recontextualize the recalled learning. In this way, retrieval develops into sharing (path 3).

Group storage can also lead to sharing (path 6). Storing learning in group-accessible spaces provides the opportunity for additional sharing. In the incident response teams we studied, having clearly understood procedures was critical for effective responding during crises. During the TCP Wrappers incident, the team became concerned about what it termed *distraction protection*. Members were concerned that an intruder might launch a second attack while the group was absorbed in dealing with the first attack. The group shared learning about the importance of sending some team members home for sleep so that someone would be alert and thinking clearly. In the process of storing this learning in the group's documented crisis response procedures, the group encountered seemingly contradictory learning about the importance of increasing staffing during crises so that some team members could monitor the normal traffic of incoming emails. This prompted additional discussion about priorities and resulted in the modification of both the old and new learning.

Identifying the processes that constitute group learning not only clarifies the necessary and sufficient conditions for group learning but also focuses attention on the relationships among these processes. While it is possible to examine each of these processes independently, we think the interactions among the processes have the greatest potential for study and practice. These relationships also provide a roadmap for diagnosing group learning failures. When groups fail to retrieve knowledge, the cause is not necessarily with group retrieval mechanisms; it may be traced to problems in sharing or storage, or the interaction of the processes.

DISCUSSION AND IMPLICATIONS

The framework we have outlined in this paper is designed to change how people study group learning. Our goals were to identify (1) the features of group learning, (2) gaps in the literature, (3) new propositions for research, and (4) new ways of thinking about group learning. We have argued that change is needed because there currently is no integrated view of group-level learning. The varied definitions we highlighted earlier point to the disjointed treatment of this fundamental construct. Papers purporting to cover group learning measure everything from individual learning in the context of groups to changes in group performance (confusing other variables with group learning). We have advanced a conceptualization that reconciles different views of group learning with an emphasis on the processes that define learning at the group level. Without a shared understanding of group learning, it will be impossible to accumulate evidence in any coherent way.

To advance the study of group learning, it is especially important to be clear about the level of analysis (in both the conceptualization of group learning and its measurement). In particular, group learning should not be confused with individual learning in the context of groups. Individual-level learning in social situations is already amply covered by such concepts as social learning (Bandura, 1977) and situated learning (Lave & Wenger, 1991). Nor should learning by top management teams be confused with organizational learning. Failure to be clear about issues related to the level of analysis leads to imprecision in models, confusion in data collection and analysis, and controversy regarding conclusions (Klein et al., 1994).

We also have argued for an increased focus on the processes inherent in group learning: sharing, storage, and retrieval. There are two important reasons for this. First, when we talk about learning, our expectation is that a group confronting a new but similar situation can retrieve what it learned at an earlier time. Retrieval is intimately tied to the storage and sharing processes. We are not interested in whether a group was trained or it stored some information but, rather, whether it can enact these basic processes in an integrated way. Second, without a focus on the basic mechanisms of group learning, it can be difficult to distinguish learning

from other exogenously induced group performance changes. For instance, if a group's level of customer satisfaction improves, is this due to changes in the group's routines or changes in the group's mix of customers? Without understanding the group's processes of sharing, storage, and retrieval, it is hard to attribute the change to learning. A focus on the processes of group learning opens up new and exciting avenues for research regarding practice, decay, priming, and other processes reviewed in the propositions we have outlined.

In addition, defining the construct space brings critical gaps in the literature into focus. First, there is a gap in understanding basic concepts such as group-level storage and retrieval. For example, little work has been done on alternative storage systems and their impact on retrieval. Second, researchers have focused on a very limited set of learning outcomes. Research to date has examined the learning of fairly simple concrete knowledge, as opposed to more complex, abstract, or tacit knowledge or routines. We encourage researchers to explore the full range of group learning outcomes, including cognitive outcomes (we think differently about responding to media requests), behavioral outcomes (we act differently in situations involving media requests), and emotional outcomes (we feel differently when confronted with media requests). Third, the majority of group learning research has a positive bias. The prevailing assumption is that learning generally leads to beneficial effects. Little attention is given to how groups learn dysfunctional routines or how groups learn about incorrect relationships.

In addition to responding to gaps and testing new propositions, there are other new opportunities for research. One interesting opportunity deals with the form of groups to study. One can contrast traditional face-to-face groups, composed of members from the same organization that persist over time, with groups operating in a distributed environment, composed of members from different organizations, which by their nature have a very short life (Goodman & Wilson, 2000). Most of the larger computer emergency groups we studied functioned together for several hours or a few days, but then never met again. The question is how do the processes of sharing, storage, and retrieval occur in these very different kinds of groups? Can we talk about group-level learning for a group that lasts

for three hours, disbands, and never meets again? Thinking about groups that vary along these multiple dimensions challenges our thinking about how groups learn.

Another opportunity deals with the impact of the external environment on learning. In contrast to individual learning research conducted in the laboratory, groups in organizations typically operate in more dynamic environments; the groups' composition often is in flux; and the knowledge, routines, and behavior they must learn are more complex. For groups embedded in organizations, the environment can change—meaning that learning valid at one time can easily become counterproductive at another time. This suggests that how groups interact with their external environments may significantly affect their ability to adapt. To date, most of the research on group learning has focused on the internal workings of groups. However, we expect that a group's orientation to its external environment (e.g., Ancona & Caldwell, 1992) also will affect the quality and extent of its learning.

Our conceptualization also highlights the possibilities for *implicit* learning at the group level—that is, learning that does not depend on a group's conscious awareness of the learning (Reber, 1989). Our observations indicate that there often are instances in which groups gradually adopt new habits or procedures over time, without being consciously aware that they are changing. Just as individuals often pick up mannerisms unconsciously from their contact with others, we expect that groups will change their repertoires through observation or contact with others. However, we know very little about this form of learning in groups. Thus, key questions for future research include "Under what circumstances is implicit learning most likely to occur?" and "Are there certain group characteristics that make it more likely?"

The final important implication of our framework concerns the methods for studying group learning, which must not only account for learning over time but must also measure the construct at the appropriate level of analysis. As we have illustrated, time is a critical feature of learning. We need to move to research designs that permit different types of practice schedules. We need to examine learning over time to understand such processes as retrieval and transfer of learning. Also, we need to operationalize the concept of group learning and separate

learning from performance or other outcomes (Druskat & Kayes, 2000). In terms of data collection, understanding storage or retrieval processes will require detailed, real-time observation. For instance, identifying and analyzing implicit learning requires multiple observations of the same group over time. Because the learning is not necessarily consciously accessible, asking the group members about what they have learned will not uncover any changes. Edmondson (2002) provides a good example of the kind of observation that can identify implicit learnings. By repeatedly observing a senior team over time, she was able to identify learned patterns of behavior (e.g., using metaphors to score philosophical points) that members were not consciously aware of.

A more comprehensive approach for measuring group learning and its component processes can be illustrated using an example from the computer emergency response teams we studied. One learning repeatedly shared in the teams' After Action Reviews was the idea that they should schedule interactive press conferences, rather than simply issuing press releases, in response to major attacks on the internet. This idea was widely discussed after the Love Letters incident, and the incident response team leader was implicitly designated as the storage repository for this knowledge (although it was also stored in the After Action Review notes from both the Love Letters and the TCP Wrappers attacks). A complete measure of group learning in this situation would be spread across at least two time periods. At Time 1 it would be possible to measure sharing and preliminary storage. The *process of sharing* could be studied by observing the group's discussion of the learning at Time 1. For example, we observed the incident response team discuss interactive press conferences as group learning in April following the Love Letters virus. At that time, five group members were present; two (Seth and Alex) explicitly shared the learning during the discussion. The depth of the group's sharing about the interactive press conference learning could be coded on the following scale: 1 = discussed the learning; 3 = discussed the learning and where it should be stored; 5 = discussed the learning, where it should be stored, and the conditions under which it should be evoked in the future. The *outcome of sharing* could be measured by administering a survey

with an open-ended item (What did the group learn as a result of this incident?) at the end of time period one (April). The percentage of group members who reported the learning about interactive press conferences would be an indicator of the breadth of sharing. The Time 1 *storage* of learning could be measured on the same survey, by asking group members to indicate where each item the group learned was stored.

Measurement of longer-term storage and retrieval would require an assessment of whether the group could produce the learned response when the appropriate stimuli presented themselves (i.e., at Time 2). One practical problem with this requirement is that it is not always possible to predict when the stimuli to elicit the group learning will next occur in organizations. To deal with the problem of waiting for the next trigger event, we propose an adaptation of the scenario method. Researchers can present groups with scenarios that call for previously learned responses and observe what learning individual members retrieve, as well as the ultimate response produced by the group. So, in our example of response teams, in order to measure longer-term *storage and decay* of this learning, we would recommend waiting for the average interval between incidents (in the case of major attacks on the internet, approximately one month) and then giving each member of the group a scenario describing another major attack and asking how each thinks the group should respond to the incident. In this case, it may be that only two members of the group (Kyle and Aaron) individually recall the idea of using interactive press conferences. This Time 2 result can be compared with the Time 1 storage results to determine the decay of the learning in the group members' memories (e.g., five members stored the learning at Time 1, but it persisted in only two members' memories a month later). To measure actual group *retrieval*, we would also ask the group to agree on a collective approach to the situation in the scenario (to observe patterns in the retrieval process). By observing the group's discussion of how to respond in the situation, we would observe what storage repositories the group uses (e.g., Does anyone mention the After Action Review agreements?) and what is, or is not, ultimately retrieved. It may be that neither of the members who individually recalled the press conference learning would bring it up in the group retrieval

period. In that case, even though the learning is still stored in the memories of two group members, it is not effectively retrieved by the group as a whole. The scenario method provides the condition for identifying the persistence of the learning, as well as any anomalies in group retrieval.

The conceptualization of group learning that we have presented also has implications for practice. The fundamental processes that we have defined in this paper could be leveraged in the design and diagnosis of groups to improve their adaptability. Groups could be designed with more explicit storage repositories and procedures. Group practice schedules could be developed that more closely mirror changes in the group's environment. And assessments of group functioning could specifically include attention to how the group shares, stores, and retrieves key information. Although group learning has long been recognized as an important indicator of group functioning (Gladstein, 1984), most practical assessments of groups do not yet cover depth and breadth of sharing, the types of storage repositories groups use, or the reliability of their retrieval mechanisms (Fitz-Enz, 1997; Jones & Schilling, 2000).

Groups have been identified as the principal vehicle for learning in organizations (Edmondson, 2002; Senge, 1990). If groups are so central to the adaptive process in organizations, we need a clearer understanding of group learning. In this paper we have outlined a comprehensive conceptualization of group learning. We have specified the necessary and sufficient conditions for group learning to occur and have used the conceptualization to highlight new areas for inquiry. Our hope is to ultimately advance our ability to explain and predict group learning.

REFERENCES

- Ancona, D. G., & Caldwell, D. F. 1992. Bridging the boundary: External activity and performance in organizational teams. *Administrative Science Quarterly*, 37: 634-666.
- Anderson, J. R. 2000. *Learning & memory: An integrated approach*. New York: Wiley.
- Argote, L. 1996. Organizational learning curves: Persistence, transfer, and turnover. *International Journal of Technology Management*, 11: 759-770.
- Argote, L. 1999. *Organizational learning: Creating, retaining and transferring knowledge*. Norwell, MA: Kluwer.
- Argote, L., Beckman, S. L., & Epple, D. 1990. The persistence

- and transfer of learning in industrial settings. *Management Science*, 36: 140–155.
- Argote, L., Gruenfeld, D. H., & Naquin, C. 2001. Group learning in organizations. In M. E. Turner (Ed.), *Groups at work: Advances in theory and research*: 369–411. Mahwah, NJ: Lawrence Erlbaum Associates.
- Argote, L., & McGrath, J. E. 1993. Group processes in organizations: Continuity and change. *International Review of Industrial and Organizational Psychology*, 8: 333–389.
- Argyris, C. 2003. A life full of learning. *Organization Studies*, 24: 11–78.
- Argyris, C., & Schön, D. 1995. *Organizational learning: Theory, method and practice*. New York: Addison-Wesley.
- Bandura, A. 1977. *Social learning theory*. Englewood Cliffs, NJ: Prentice-Hall.
- Basden, B. H., Basden, D. R., Bryner, S., & Thomas, R. L. 1997. A comparison of group and individual remembering: Does collaboration disrupt retrieval strategies? *Journal of Experimental Psychology: Learning, Memory and Cognition*, 23: 1176–1189.
- Blickensdorfer, E., Cannon-Bowers, J. A., & Salas, E. 1997. *Training teams to self-correct: An empirical investigation*. Paper presented at the 12th annual meeting of the Society for Industrial and Organizational Psychology, St. Louis.
- Cannon-Bowers, J. E., Salas, E., & Converse, S. 1993. Shared mental models in expert team decision-making. In J. Castellan (Ed.), *Individual and group decision making: Current issues*: 221–247. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Carley, K. 1992. Organizational learning and personnel turnover. *Organization Science*, 3: 20–46.
- Cohen, M. D., & Bacdayan, P. 1994. Organizational routines are stored as procedural memory: Evidence from a laboratory study. *Organization Science*, 5: 554–568.
- Cook, S. D., & Yanow, D. 1993. Culture and organizational learning. *Journal of Management Inquiry*, 2: 373–390.
- Crossan, M. M., Lane, H. W., & White, R. E. 1999. An organizational learning framework: From learning to institution. *Academy of Management Review*, 24: 522–537.
- Devadas, R., & Argote, L. 1995. *Collective learning and forgetting: The effects of turnover and group structure*. Paper presented at the Midwestern Academy of Management, Chicago.
- Diehl, M., & Stroebe, W. 1987. Productivity loss in brainstorming groups: Toward the solution of a riddle. *Journal of Personality and Social Psychology*, 53: 497–509.
- Donovan, J. J., & Radosevich, D. J. 1999. A meta-analytic review of the distribution of practice effect: Now you see it, now you don't. *Journal of Applied Psychology*, 84: 795–805.
- Druskat, V. U. & Kayes, D. C. 2000. Learning versus performance in short-term project teams. *Small Group Research*, 31: 328–353.
- Edmondson, A. 1999. Psychological safety and learning behavior in work teams. *Administrative Science Quarterly*, 44: 350–383.
- Edmondson, A. 2002. The local and variegated nature of learning in organizations: A group level perspective. *Organization Science*, 13: 128–146.
- Edmondson, A. C., Bohmer, R. M., & Pisano, G. P. 2001. Disrupted routines: Team learning and new technology implementation in hospitals. *Administrative Science Quarterly*, 46: 685–716.
- Ellis, A. P., Hollenbeck, J. R., Ilgen, D. R., Porter, C. O., & West, B. J. 2003. Team learning: Collectively connecting the dots. *Journal of Applied Psychology*, 88: 821–832.
- Finlay, F., Hitch, G. J., & Meudell, P. R. 2000. Mutual inhibition in collaborative recall: Evidence for a retrieval-based account. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 26: 1556–1567.
- Fiol, C. M., & Lyles, M. A. 1985. Organizational learning. *Academy of Management Review*, 10: 803–813.
- Fitz-Enz, J. 1997. Measuring team effectiveness. *HR Focus*, 74(8): 3–7.
- Gersick, C. J. G., & Hackman, J. R. 1990. Habitual routines in task-performing groups. *Organizational Behavior and Human Decision Processes*, 47: 65–97.
- Gibson, C., & Vermeulen, F. 2003. A healthy divide: Subgroups as a stimulus for team learning behavior. *Administrative Science Quarterly*, 48: 202–224.
- Gladstein, D. 1984. Groups in context: A model of task group effectiveness. *Administrative Science Quarterly*, 29: 499–517.
- Goodman, P. S., & Darr, E. D. 1996. Computer-aided systems for organizational learning. *Trends in Organizational Behavior*, 3: 81–97.
- Goodman, P. S., & Wilson, J. M. 2000. Substitutes for socialization in exocentric teams. *Research in Groups and Teams*, 3: 53–77.
- Gruenfeld, D. H., Martorana, P. V., & Fan, E. T. 2000. What do groups learn from their worldliest members? Direct and indirect influence in dynamic teams. *Organizational Behavior and Human Decision Processes*, 82: 45–59.
- Hinsz, V. B., Tindale, R. S., & Vollrath, D. A. 1997. The emerging conceptualization of groups as information processors. *Psychological Bulletin*, 121: 43–65.
- Hollingshead, A. B. 1998. Communication, learning and retrieval in transactive memory systems. *Journal of Experimental Social Psychology*, 34: 423–442.
- Hollingshead, A. B. 2001. Cognitive interdependence and convergent expectations in transactive memory. *Journal of Personality and Social Psychology*, 81: 1–10.
- Huber, G. P. 1991. Organizational learning: The contributing processes and literatures. *Organization Science*, 2: 88–115.
- Ilgen, D. R., Hollenbeck, J. R., Johnson, M., & Jundt, D. 2005. Teams in organizations: From I-P-O models to IMOI models. *Annual Review of Psychology*, 56: 517–543.
- Ilgen, D. R., & Klein, H. J. 1989. Organizational behavior. *Annual Review of Psychology*, 40: 327–351.

- Jones, S. D., & Schilling, D. J. 2000. *Measuring team performance*. New York: Wiley.
- Klein, K. J., Dansereau, F., & Hall, R. J. 1994. Levels issues in theory development, data collection, and analysis. *Academy of Management Review*, 19: 195–229.
- Klimoski, R., & Mohammed, S. 1994. Team mental model: Construct or metaphor. *Journal of Management*, 20: 403–437.
- Lant, T. K. 1992. Aspiration level adaptation: An empirical exploration. *Management Science*, 38: 623–645.
- Laughlin, P. R., & Shupe, E. I. 1996. Intergroup collective induction. *Organizational Behavior and Human Decision Processes*, 68: 44–57.
- Lave, J., & Wenger, E. 1991. *Situated learning: Legitimate peripheral participation*. Cambridge: Cambridge University Press.
- Levine, J. M., Resnick, L. B., & Higgins, E. T. 1993. Social foundations of cognition. *Annual Review of Psychology*, 44: 585–612.
- Levitt, B., & March, J. G. 1988. Organizational learning. *Annual Review of Sociology*, 14: 319–340.
- Loewenstein, J., Thompson, L., & Gentner, D. 2003. Analogical learning in negotiation teams: Comparing cases promotes learning and transfer. *Academy of Management Learning & Education*, 2: 119–127.
- London, M., Polzer, J. T., & Omeregic, H. 2005. Interpersonal congruence, transactive memory, and feedback processes. *Human Resource Development Review*, 4(2): 114–135.
- Mathieu, J. E., Heffner, T. S., Goodwin, G. F., Salas, E., & Cannon-Bowers, J. A. 2000. The influence of team mental models on team process and performance. *Journal of Applied Psychology*, 85: 273–283.
- Mohammed, S., & Dumville, B. C. 2001. Team mental models in a team knowledge framework: Expanding theory and measurement across disciplinary boundaries. *Journal of Organizational Behavior*, 22: 89–106.
- Moorman, C., & Miner, A. S. 1998. Organizational improvisation and organizational memory. *Academy of Management Review*, 23: 698–723.
- Moreland, R., Hogg, M. A., & Hanes, S. C. 1994. Back to the future: Social psychological research on groups. *Journal of Experimental Social Psychology*, 30: 527–555.
- Moreland, R. L., Argote, L., & Krishnan, R. 1998. Training people to work in groups. In R. S. Tindale (Ed.), *Theory and research in small groups*: 37–60. New York: Plenum Press.
- Moreland, R. L., & Myaskovsky, L. 2000. Exploring the performance benefits of group training: Transactive memory or improved communication? *Organization Behavior and Human Decision Processes*, 82: 117–133.
- Morgeson, F. P., & Hofmann, D. A. 1999. The structure and function of collective constructs: Implication for multi-level research and theory development. *Academy of Management Review*, 24: 249–265.
- Naquin, C. E., & Tynan, R. O. 2003. The team halo effect: Why teams are not blamed for their failures. *Journal of Applied Psychology*, 88: 332–305.
- Olivera, F. 2000. Memory systems in organizations: An empirical investigation of mechanisms for knowledge collection, storage and access. *Journal of Management Studies*, 37: 811–832.
- Reber, A. S. 1989. Implicit learning and tacit knowledge. *Journal of Experimental Psychology: General*, 118: 219–235.
- Rosenthal, R., & Rosnow, R. L. 1991. *Essentials of behavioral research: Methods and data analysis*. New York: McGraw-Hill.
- Rousseau, D. M. 1985. Issues of level in organizational research: Multi-level and cross-level perspectives. *Research in Organizational Behavior*, 7: 1–37.
- Schippers, M. C., Den Hartog, D. N., & Koopman, P. L. 2007. Reflexivity in teams: A measure and correlates. *Applied Psychology: An International Review*, 56: 189–211.
- Senge, P. M. 1990. *The fifth discipline: The art and practice of the learning organization*. New York: Doubleday.
- Simon, H. A. 1947. *Administrative behavior*. New York: Macmillan.
- Smith, S. M., & Vela, E. 2001. Environmental context-dependent memory: A review and meta-analysis. *Psychonomic Bulletin and Review*, 8: 203–220.
- Snyder, M., & Gangestad, S. 1986. On the nature of self-monitoring: Matters of assessment, matters of validity. *Journal of Personality and Social Psychology*, 51: 125–139.
- Sole, D., & Edmondson, A. 2002. Situated knowledge and learning in dispersed teams. *British Journal of Management*, 13: 17–34.
- 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: 1467–1478.
- Sutton, R. I., & Staw, B. M. 1995. What theory is not. *Administrative Science Quarterly*, 40: 371–384.
- Tjosvold, D., Yu, Z., & Hui, C. 2004. Team learning from mistakes: The contribution of cooperative goals and problem-solving. *Journal of Management Studies*, 41: 1223–1245.
- Tolman, E. C. 1932. *Purposive behavior in animals and man*. New York: Appleton-Century-Crofts.
- Van der Vegt, G. S., & Bunderson, J. S. 2005. Learning and performance in multi-disciplinary teams: The importance of collective identification. *Academy of Management Journal*, 48: 532–547.
- Weldon, M. S., Blair, C., & Huebsch, P. D. 2000. Group remembering: Does social loafing underlie collaborative inhibition? *Journal of Experimental Psychology: Learning, Memory and Cognition*, 26: 1568–1577.
- Whetten, D. A. 1989. What constitutes a theoretical contribution? *Academy of Management Review*, 14: 490–495.
- Wittenbaum, G. M., Hubbell, A. P., & Zuckerman, C. 1999. Mutual enhancement: Toward an understanding of the collective preference for shared information. *Journal of Personality and Social Psychology*, 77: 967–978.

- Wong, S. S. 2004. Distal and local group learning: Performance tradeoffs and tensions. *Organization Science*, 15: 645–656.
- Zellmer-Bruhn, M. E. 2003. Interruptive events and team knowledge acquisition. *Management Science*, 49: 514–528.

Jeanne M. Wilson (jeanne.wilson@business.wm.edu) is an associate professor at the Mason School of Business at The College of William & Mary. She received her Ph.D. from Carnegie Mellon University. Her research focuses on new organizational forms, particularly the effects of distance and the development of trust over time in distributed groups.

Paul S. Goodman (pg14@andrew.cmu.edu) is the Richard Cyert Professor of Organizational Psychology at the Tepper School of Business, Carnegie Mellon University. His research interests are in groups, organizational errors, and learning contracts in new forms of higher education. He is also a documentary film maker about work systems.

Matthew A. Cronin (mcronin@gmu.edu) is an assistant professor of management at George Mason University's School of Management. He earned his Ph.D. at the Tepper School of Business, Carnegie Mellon University. His research focuses on the cognitive processes that promote creativity, particularly in the context of teams.