

Running Head: THEORY OF TEAM KNOWLEDGE EMERGENCE

**The Dynamics of Team Cognition:
A Process-Oriented Theory of Knowledge Emergence in Teams**

James A. Grand
University of Maryland

Michael T. Braun
University of South Florida

Goran Kuljanin
DePaul University

Steve W.J. Kozlowski
Georgia T. Chao
Michigan State University

Citation:

Grand, J.A., Braun, M.T., Kuljanin, G., Kozlowski, S.W.J., & Chao, G.T. (in press). The dynamics of team cognition: A process-oriented theory of knowledge emergence in teams. *Journal of Applied Psychology* [Monograph].

Contact Information:

James A. Grand, Department of Psychology, 4094 Campus Dr., 3147A Biology-Psychology Building, University of Maryland, College Park, MD 20742, phone: (301) 405-5935, email: grandjam@umd.edu

Michael T. Braun, Department of Psychology, 4202 East Fowler Ave, PCD 4151, University of South Florida, Tampa, FL 33620, phone: (813) 974-0354, email: mtbraun@usf.edu

Goran Kuljanin, Department of Psychology, 2219 N. Kenmore Ave., 521 Byrne Hall, DePaul University, Chicago, IL 60614, phone: (847) 207-9292, email: gkuljani@depaul.edu

Steve W. J. Kozlowski, Department of Psychology, 316 Physics Road, 309 Psychology Building, Michigan State University, East Lansing, MI 48824-1116, phone: (517) 353-8924, email: stevekoz@msu.edu

Georgia T. Chao, Department of Management, 632 Bogue Street, N435 North Business Complex, Michigan State University, East Lansing, MI 48824-1122, phone: (517) 353-5418, email: chaoq@msu.edu

Author's Note:

We gratefully acknowledge the U.S. Army Research Institute for the Behavioral and Social Sciences (ARI; W911NF-14-1-0026, S.W.J. Kozlowski and G.T. Chao, Principal Investigators; J.A. Grand, M.T. Braun, and G. Kuljanin, Co-Investigators), the National Science Foundation (NSF, 1533499, S.W.J. Kozlowski and G.T. Chao, Principal Investigators), and the Office of Naval Research (ONR) Command Decision Making (CDM) Program (N00014-09-1-0519, S.W.J. Kozlowski and G.T. Chao, Principal Investigators), for support of this research. Any opinions, findings, conclusions and recommendations expressed are those of the authors and do not necessarily reflect the views of ARI, NSF, ONR, or the CDM Program. We also thank Gilad Chen and two anonymous reviewers for their helpful comments on this paper. Lastly, we are grateful to Stanton Mak and Daniel Pickhardt for their assistance during data collection.

Abstract

Team cognition has been identified as a critical component of team performance and decision-making. However, theory and research in this domain continues to remain largely static; articulation and examination of the dynamic processes through which collectively held knowledge emerges from the individual- to the team-level is lacking. To address this gap, we advance and systematically evaluate a process-oriented theory of team knowledge emergence. First, we summarize the core concepts and dynamic mechanisms that underlie team knowledge-building and represent our theory of team knowledge emergence (Step 1). We then translate this narrative theory into a formal computational model that provides an explicit specification of how these core concepts and mechanisms interact to produce emergent team knowledge (Step 2). The computational model is next instantiated into an agent-based simulation to explore how the key generative process mechanisms described in our theory contribute to improved knowledge emergence in teams (Step 3). Results from the simulations demonstrate that agent teams generate collectively shared knowledge more effectively when members are capable of processing information more efficiently and when teams follow communication strategies that promote equal rates of information sharing across members. Lastly, we conduct an empirical experiment with real teams participating in a collective knowledge-building task to verify that promoting these processes in human teams also leads to improved team knowledge emergence (Step 4). Discussion focuses on the implications of examining team cognition processes and dynamics as well as directions for future research.

Keywords: team knowledge; team cognition; emergence; computational modeling; agent-based simulation

Since the concept of team cognition was first proposed and popularized in the organizational sciences (Cannon-Bowers, Salas, & Converse, 1990), an abundance of empirical research has accumulated examining its contribution to team effectiveness. Team cognition has been evaluated as both shared mental models (collectively shared mental representations of a team domain and task) as well as transactive memory systems (collectively distributed and organized knowledge related to team and task demands), and there is now meta-analytic evidence supporting the contribution of both conceptualizations to team performance (DeChurch & Mesmer-Magnus, 2010). There is a similarly extensive body of research demonstrating that the distribution of expertise and specialization among team members has a significant impact on the quality of knowledge-building and decision-making within teams (Mesmer-Magnus & DeChurch, 2009; Lu, Yuan, & McLeod, 2012). Despite recognition of its importance to team functioning, a coherent understanding of *how* and *why* team cognition develops remains incomplete. That is, little attention has been directed towards the dynamic processes through which knowledge is acquired, compiled, and manifests at the team level. Basic theory of how collectively held and actionable knowledge emerges from the individual to the team level is lacking.

Understanding how knowledge emerges from the individual to the team-level is critical to advancing research on team cognition. A primary objective of team cognition research is to develop insights into how the quantity and quality of knowledge held by teams as well as the capability to generate and disseminate information among team members can be improved. This goal has traditionally been pursued by examining static operationalizations of team knowledge and construct-to-construct relationships with antecedents measured at a single time-point (e.g., team ability, familiarity, task complexity) or following an intervention (e.g., cross training, crew resource management, etc., see Wildman et al., 2012). Although these efforts have expanded the nomological network of team cognition, they do not reveal the underlying *process* mechanisms responsible for team knowledge emergence. The manner by which team knowledge is created and sustained remains a “black box.”

A lack of theoretical and empirical attention towards unpacking this black box and revealing the processes responsible for team knowledge emergence is problematic for at least two reasons. First, the development and maintenance of knowledge at the team level inherently involves intra- and inter-member processes unfolding over time (Fiore, Rosen, et al., 2010; Hinsz, Tindale, & Vollrath, 1997). However, what these processes are and how they interact to facilitate team cognition are not well explicated. Efforts to enhance *team* knowledge outcomes cannot be as effectively achieved without better understanding what *individuals* do and how they work together to generate collectively held knowledge. Second, ignoring the process of team knowledge emergence limits inferences about why

teams may be more or less successful at achieving desired levels of team cognition. Focusing on construct-to-construct relationships between antecedents and team knowledge outcomes can elucidate conditions that facilitate or attenuate the development of collective knowledge (i.e., inputs → outcome). However, such research does not provide insight into *how* or *why* these conditions lead to improved knowledge outcomes. Systematically pursuing these answers necessitates directive theory on team knowledge-building *processes* that indicate how key events, actions, etc. lead to individual- and team-level knowledge outcomes (i.e., process → outcome). Unfortunately, the approaches to theory building (e.g., “box-and-arrow” models of construct-to-construct relationships) and theory testing (e.g., cross-sectional, self-report-based research) most commonly used in the organizational sciences do not permit the precision and transparency needed to specify how and why team processes shape important team outcomes (Cronin, Weingart, & Todorova, 2011; Kozlowski & Chao, 2012b; Lord, Dinh, & Hoffman, 2015). For example, a study examining the relationship between team composition and team knowledge outcomes may be able to infer that “teams with characteristic X tend to have poorer team knowledge outcomes.” However, a study that explicitly focuses on the processes of team knowledge-building would be able to infer that “teams with characteristic X have poorer knowledge outcomes *because* they have difficulty doing _____.” The latter provides a more precise account of the phenomenon as well as a better understanding of how and why team cognition is shaped—a critical and valuable contribution in this research domain (Mohammed, Tesler, & Hamilton, 2012).

To begin addressing these gaps and unpacking the black box of team cognition, we advance and test a process-oriented theory of team knowledge emergence that elaborates the core concepts and mechanisms through which team cognition develops. Developing process-oriented theory requires an explicit account of how key cognitive, affective, and/or behavioral activities interact to produce and sustain an emergent construct. To guide the development and evaluation of our theory, we advance a four-step framework distilled from recommended practices for investigating emergent phenomena (Kozlowski, Chao, Grand, Braun, & Kuljanin, 2013). Figure 1 provides a visual summary of the framework and also serves as the organizing structure for the paper. Step 1 of this protocol involves constructing a process-level account of an emergent phenomenon—that is, a narrative theory of what individuals do, think, feel, etc. that gives rise to a higher-level outcome. For our theory of team knowledge emergence, this consists of identifying a set of core constructs and process mechanisms involved in team knowledge-building and how these elements dynamically interact to produce collective knowledge over time. Step 2 entails translating the theoretical narrative into a formal computational model that specifies the coherence and logic of the core theoretical mechanisms and their interactions (Law & Kelton, 1991; Taber & Timpone, 1996; Vancouver & Weinhardt, 2012).

With respect to our theory, we translate our narrative account of knowledge emergence within teams into a formal computational model of the processes carried out by individuals to generate collective knowledge. In Step 3, the computational model is instantiated into a computer simulation to explore the theoretical space and generate propositions, insights, and/or prescriptions related to the emergent phenomenon (Harrison, Lin, Carroll, Carley, 2007). We enact this step using agent-based simulation (ABS) to advance propositions about how and why the core processes of knowledge emergence proposed in our theory influence knowledge outcomes in teams. Step 4 involves evaluating the fidelity between patterns observed in simulated and human data and exploring the utility of predictions gleaned from our simulation for influencing outcomes in real teams. We carry out this step by empirically investigating whether efforts to influence the generative process mechanisms identified in our theory and explored in the ABS improve knowledge emergence for real teams participating in a knowledge-building task. Lastly, results from the theory evaluation can be used to suggest directions for future research, model revision, and implications for practice. Our paper concludes with a discussion of the theory's potential to explain knowledge emergence in teams with distributed expertise, its utility for organizing existing work on team cognition, and its usefulness for guiding research. Overall, this cyclical four-step procedure embodies principles of scientific inquiry by both developing (Steps 1 and 2) and then systematically evaluating (Steps 3 and 4) process-oriented theory to advance research and practice.

In sum, there is significant need to unpack the "black box" of team cognition by advancing theory on the process of knowledge emergence within teams. The present research addresses this gap by (1) identifying and describing core concepts and process mechanisms of team knowledge emergence; (2) constructing a computational model to formally describe the relations among these core knowledge processes; (3) instantiating the computational model into a computer simulation to generate insights and prescriptions for enhancing team knowledge emergence; and (4) evaluating these propositions in real teams engaged in knowledge-building.

Theories of Emergence

Before elaborating our theory of team knowledge emergence, it is important to clarify the meaning of emergence and to identify key considerations for developing theories of emergent phenomena. Multilevel theory (Kozlowski & Klein, 2000) defines *emergence* as a dynamic *process* by which the interactions of elements at a lower level of analysis (e.g., learning and sharing information by individuals) give rise to constructs at higher levels. By contrast, an *emergent construct* is an *outcome* or pattern of outcomes generated and sustained by the interactions of lower-level elements over time (e.g., team knowledge). A theory of emergence therefore defines the bottom-up mechanisms and interactions that give rise to emergent constructs at higher system levels (Kozlowski et al., 2013).

Developing and testing theories of emergence requires an alternative conceptual approach to conventional theory-building paradigms employed in much of organizational research. Contemporary theory development in organizational research most often focuses on elaborating the correlational relationship among constructs (e.g., construct A positively relates to construct B; construct C mediates the relationship between constructs D and E). In contrast, the objective of a theory of emergence is to describe the nature of lower-level process mechanisms in a system/collective (e.g., if person 1 performs behavior A, then person 2 performs behavior B; if event X occurs, person 1 experiences outcome Y). Specifying a theory of emergence thus requires identifying, operationalizing, and justifying both the core concepts involved in an emergent phenomena as well as the “rules” or process mechanisms that describe what, when, and how lower-level entities think, behave, and/or react to environmental stimuli (events, other members, tasks, etc.). By meeting this demand, theories of emergence explicitly recognize the dynamic processes (e.g., feedback/feedforward mechanisms, simultaneous activation, inhibition/excitation, etc.) that generate and sustain emergent and organizationally relevant constructs. An important consequence is that theories of emergence permit multiple “pathways” through which an emergent construct may unfold. In doing so, they embrace principles of equifinality characteristic of open systems such as teams and organizations (i.e., many ways to reach the same end state, Cronin et al., 2011; Katz & Kahn, 1978; von Bertalanffy, 1950).

These considerations provide critical context for defining team knowledge emergence. Whereas *team knowledge* is an emergent construct, *team knowledge emergence* is a dynamic process. A theory of how and why knowledge emerges in teams must describe the individual-level concepts relevant to knowledge development as well as the manner by which those concepts interact to generate knowledge at the team-level. The purpose of such theory is to provide an explicit account of how key process mechanisms contribute to team cognition (process → outcome). In turn, this theoretical specification can provide a more informed conceptual framework through which interventions to enhance the level and quality of collectively held knowledge in a team can be pursued (input → process). Furthermore, precision at the process-level facilitates the ability to incorporate, revise, and extend theory on team cognition as future research continues to explore how teams generate and sustain knowledge under different conditions, constraints, and environments.

STEP 1: PROCESS-ORIENTED THEORY OF KNOWLEDGE EMERGENCE IN TEAMS

We propose that knowledge emergence in teams is primarily driven by two fundamental activities of individuals—*learning* and *sharing*. Learning is a ubiquitous feature of virtually all theories of individual and team cognition and characterizes how individuals extract information from their environment. Comparatively, sharing

reflects the foundational aspects of communication and interaction through which information is disseminated across members (Fiore, Rosen, et al., 2010; Hinsz et al., 1997; Stasser & Titus, 1985). To construct a process-level account of team knowledge emergence, we identify a core set of process mechanisms to represent learning and sharing within teams. The focus of our theory is on teams with members possessing distributed and specialized expertise who must actively acquire information from the environment and each other in order to reach a fully shared and agreed upon understanding of a problem space. Such teams are common in many organizations (e.g., top management teams, multi-disciplinary decision-making teams, interprofessional advisory committees, joint task forces, etc.). Furthermore, their capacity to efficiently and effectively develop actionable and collectively held knowledge is a crucial component of performance at multiple organizational levels (Fiore, Rosen, et al., 2010; Fiore, Smith-Jentsh, et al., 2010; Mesmer-Magnus & DeChurch, 2009).

Although it is possible to include other concepts and mechanisms in this process (e.g., motivation, affect, conflict, etc.), we adhere to tenets of complexity science and conceptual parsimony that a theory of emergence should pose the fewest and most fundamental processes capable of representing the phenomenon (Epstein, 1999; Miller & Page, 2007; Reynolds, 1987). In the sections that follow, we thus explicate the most fundamental mechanisms of learning and sharing and how they interact over time to generate emergent team-level knowledge. This specification represents our process-oriented theory of team knowledge emergence and denotes how collectively shared knowledge arises from the behavioral and cognitive activities of individual team members.

Core Concepts and Process Mechanisms Involved in Learning

A small number of fundamental mechanisms are typically implicated in the process of knowledge acquisition across a wide range of learning theories. These concepts include attending to information in the environment; evaluating and representing information; storing information for later use; and interpreting how newly acquired information fits with previously acquired information. We incorporate these core process mechanisms into our theory as *Data Selection, Encoding, Decoding, and Integration*.

Data Selection. The perception of, and attention to, environmental stimuli represents the first stage of information processing in most models of human cognition and knowledge development (e.g., Anderson et al., 2004; Love, Medin, & Gureckis, 2004; Meyer & Kieras, 1997; Newell, 1990). Research in this area describes how individuals identify, recognize, and selectively filter meaningful “data” or stimuli from the environment to subsequently be processed (Broadbent, 1958; Deutsch & Deutsch, 1963; Treisman, 1964). As Hinsz et al. (1997, p. 46) note, *data selection* mechanisms address the basic question “What information is the focus of attention for team members?”

Distilling the large literature on processes related to data selection reveals three mechanisms relevant to a theory of team knowledge emergence. First, in the absence of any external influence, individuals should be equally attentive—and therefore likely to direct learning efforts—towards any given piece of information that is accessible and relevant to their area of expertise (e.g., Deutsch & Deutsch, 1963). Second, the presence of goal hierarchies and information-seeking strategies developed as a result of expertise and experience should increase the likelihood that individuals will select information to learn that contributes to a more coherent understanding of the task environment (Duncan & Humphreys, 1989; Soto, Heinke, Humphreys, & Blanco, 2005; Wickens, 1992). That is, once a given piece of information has been learned, the presence of that information should prompt individuals to attend to other information that builds upon that existing knowledge rather than focusing on less immediately relevant information (e.g., Moores, Laiti, & Chelazzi, 2003). Finally, in groups with distributed expertise, individuals can influence the data selection processes of other team members by focusing attention towards new or important sources of information (Stasser & Titus, 1985, 1987; Dionne, Sayama, Hao, & Bush, 2010).

Encoding. Once information is brought into active awareness, individuals can attempt to internalize that data. *Encoding* encompasses how individuals transform external stimuli, events, or information into representations that are then stored in memory (Baddeley, Eysenck, & Anderson, 2009). In this sense, encoding describes how individuals translate something they perceive in their environment into something they know and remember. The encoding procedure captured in our theory of team knowledge emergence reflects the notion of “*learning-from-self*” that takes place when individuals in a team learn information available to them that requires no coordination with other team members. In groups with distributed expertise, this conceptualization mirrors members’ efforts to learn information within their unique domain of specialization (Fiore, Smith-Jentsch, et al., 2010). For example, learning-from-self is exemplified by a physician in an emergency medical team reading an EKG test to acquire information about a patient’s cardiac health, a researcher in a multidisciplinary research team evaluating findings from empirical articles in her area of specialty, or a sonar technician on a submarine assessing information about potential underwater threats using his listening equipment. In all cases, the team member encodes information that he or she had access to, unique expertise to interpret, and can learn without input from other team members.

Our theory of team knowledge emergence does not distinguish among different encoding processes for separate sources of informational stimuli (e.g., visual, acoustic, semantic, etc.; Baddeley et al., 2009; Damb et al., 1995). Rather, we instantiate a simplifying assumption that encoding happens for an entire informational stimulus (e.g., person learns all aspects of a piece of information to a certain level of completeness) rather than for specific

features of information (e.g., person encodes visual attributes of information, but not semantic, etc.). Validation evidence from previous models related to information processing that employ similarly simplified assumptions indicate that this conceptualization is reasonably accurate at capturing outcomes from empirical data (Hintzman, 1984; Dougherty, Gettys, & Ogden, 1999). Encoding in our theory of team knowledge emergence therefore embodies the notion that individuals learn an entire “unit” of information more or less completely over time.

Decoding. Individuals engage in encoding/learning-from-self by interfacing with information sources they can access directly. In team contexts with distributed expertise though, individuals also learn information from other team members that would otherwise be unavailable due to lack of accessibility or expertise. In our theory of team knowledge emergence, *decoding* characterizes “*learning-from-others*” in which members acquire and interpret information shared with them by their teammates. We posit that decoding is subject to additional demands that cause it to occur more slowly relative to encoding. This interpretation differs slightly from extant treatments of team knowledge-building and therefore represents a unique contribution of our theoretical model.

A number of conceptual and empirical sources support the decision to consider a separate and more effortful decoding mechanism. First, the relative inefficiency of decoding compared to encoding is consistent with early research in educational psychology which found that students who learned material only by reading (i.e., encoding/learning-from-self) performed better on related knowledge tests than learners who were provided the same material only through lecture (i.e., decoding/learning-from-others, Corey, 1934; Spencer, 1941; Russell, 1928). Second, the social environment in which team knowledge-building occurs influences information processing. For instance, teams may reinforce certain types of heuristic reasoning that impact how a given member perceives and interprets shared information (Hinsz et al., 1997). Similarly, members may possess different interpretations of task demands and goals (Klimoski & Mohammed, 1994; Mathieu, Heffner, Goodwin, Salas, & Cannon-Bowers, 2000; Mohammed & Dumville, 2001) that influence how, when, and why individuals internalize certain types of information. Consequently, learning-from-others is likely to involve additional processing to filter ambiguous and/or unfamiliar sources of information. Third, Fiore and colleagues’ macrocognitive framework (Fiore, Rosen, et al. 2010; Fiore, Smith-Jentsch, et al., 2010) posits that distinctive individual vs. team knowledge-building processes exist that place different demands on learning. Notably, these authors suggest that team knowledge-building processes (e.g., exchanging information, evaluating information, etc.) which operate on information external to the receiving member is subject to influences both within *and* beyond the “head” of a single member. Lastly, research on semiotic models of communication notes that information which is transmitted by others is perceived as noisier and requires additional

cognitive energies to translate and interpret (Hall, 1973). In sum, although encoding and decoding processes are functionally similar mechanisms, they are subject to different constraints and are reflected as such in our theory.

Integration. Individuals do more when learning than simply accumulate unrelated facts; they categorize and organize these concepts into coherent schemas of interrelated information (Love et al., 2004; Medin & Schaffer, 1978; Rosch, Mervis, Gray, Johnson, & Boyes-Braem, 1976). The formation of associative links among internalized information thus represents a qualitatively different and more sophisticated understanding of a task environment. In our theory of team knowledge emergence, this concept is captured through an *integration* procedure in which individuals consult their existing knowledge to identify whether any newly acquired information shares relationships or dependencies with previously internalized information.

Similar to encoding and decoding processes, we focus on the functional role played by integration in the development of team knowledge. That is, rather than model specifically *how* individuals integrate information (i.e., how new informational categories are formed, how associative/similarity judgments are reconciled, etc., Nosofsky, 1984, 1987; Nosofsky & Zaki, 1998; Love et al., 2004), we simply acknowledge *that* individuals integrate information by learning relational ties and note this process could be better or worse across individuals, situations, stimuli, etc. As exemplified by recent work examining the formation of team mental models, various characteristics and processes at the group level also have the potential to impact the integration of information by individuals. For example, Kennedy and McComb (2014) demonstrated that the type of information shared and the temporal patterning by which information is communicated influences the extent to which team members integrate their understanding of the task environment into a shared representation. Similarly, Dionne et al.'s (2010) simulations of mental model formation describe how the distribution of expertise across team members, the level of confidence members place in one another's expertise, members' sensitivity to social influence, and leadership structures influence the likelihood of mental model convergence. A critical implication of this research is that the characteristics and behaviors of individuals within a team can impact the ability for members to integrate information into a coherent structure that is then collectively shared with other members (Kozlowski et al., 2013).

Core Concepts and Process Mechanisms Involved in Sharing

As individuals in a team acquire knowledge and increase their personal understanding of a given domain, the team's aggregate knowledge also grows. However, individuals are less likely to directly benefit from increases in the team's knowledge pool unless members make efforts to share and distribute uniquely held knowledge to others within the team. A critical step in the development of team knowledge is thus the transformation of *internalized*

knowledge that has been acquired and synthesized by an individual team member into what Fiore, Rosen et al. (2010) characterize as *externalized knowledge* that is collectively held, shared, and agreed upon by multiple team members. We distill the processes involved in sharing information within a team to generate externalized team knowledge into four primary mechanisms: electing to speak and share information with others; choosing what information to share with others; communicating information to others; and confirming that shared information has been received and understood. These core concepts and mechanisms are incorporated into our theory as *Member Selection, Retrieval, Sharing, and Acknowledgment*, respectively.

Member selection. In team knowledge-building contexts without a formal hierarchy or structure for determining who has the “floor” (e.g., leaderless groups, multidisciplinary project teams, etc.), members are free to share information with others when they wish. However, research examining patterns of group discussion and problem-solving find that team members do not contribute to conversations at equal rates; rather, relatively stable patterns of engagement emerge which vary systematically across individuals (e.g., Parker, 1988; Stephan & Mishler, 1952). We incorporate this individual difference as a *member selection* mechanism that characterizes team members’ likelihood of speaking at any given point in time. In this sense, member selection does not characterize an explicit decision or democratic choice made by the team to choose who speaks; instead, it simply captures the notion that individuals exhibit differential probabilities of speaking within a team.

Communication between group members reflects the most prominent mechanism through which information is disseminated in teams. Consequently, “who” speaks and “why” they speak represents an important conceptual distinction for team knowledge emergence. To guide this decision, we note that speaking dynamics in groups are typically bounded by the nature of a team’s environment, composition, and goals. For example, committee and production teams are often comprised of members with largely similar expertise (Sundstrom, De Meuse, & Futrell, 1990); in such teams, sharing information is often used to express influence or shape preferences for particular decisions/products (Stasser, 1988; Stasser & Taylor, 1991). In contrast, individuals in action and project teams commonly possess distributed and specialized expertise and must rely on one another to relay task-critical information that no one else in the team could reliably access or interpret. Communicating information in these contexts thus directly contributes to the development of collective understanding (Dionne et al., 2010). This latter characterization is consistent with the task and team environment in many modern organizational teams in which knowledge-building is critical (e.g., Fiore, Rosen et al., 2010; Wildman et al., 2012), but it reflects an operational boundary condition for the conceptualization of member communication in our theory of team knowledge emergence.

More specifically, we focus on member selection/communication patterns in environments where individuals are aware that team members are not substitutable by other members on the team (i.e., expertise is distributed and specialized) and therefore communicate primarily to transmit information rather than exert influence.

Retrieval. Along with the encoding and storing of information, the manner by which information is accessed once it has been learned ranks among the most fundamental functions of cognition and memory (Baddeley, 2009). Various characterizations of this process exist, though common to all accounts is the assumption that information *retrieval* involves identifying a piece of information from memory and bringing it into active awareness (cf., Anderson et al., 2004; Dougherty et al., 1999; Hintzman, 1984). In our theory of team knowledge emergence, retrieval determines what piece(s) of information members select to share with other team members.

There are many potential factors that might influence whether any particular piece of learned information is retrieved by an individual to be shared (e.g., recency, saliency, validity, etc., Anderson & Schooler, 1991). However, we adopt two simplifying assumptions related to retrieval in our theory of team knowledge emergence. First, it is assumed that individuals in knowledge-building teams with distributed expertise strive to reach a collectively shared and agreed upon understanding of *all* relevant knowledge in the team task environment (McComb, 2007). The primary consequence of this intention is that members are equally likely to share *any* useful knowledge they have acquired and are therefore also equally likely to retrieve any relevant knowledge to share that is available. Second, the acquisition of integrated knowledge during learning enhances the quality of information that can be retrieved and subsequently shared. A hallmark of expertise development is the transformation of disparate and seemingly unrelated information into highly integrated mental schemata (e.g., Chase & Simon, 1973; Hunt, 1994). Of significance to our theory of team knowledge emergence, individuals who possess more integrated schemata are capable of recalling both a higher quantity as well as more coherent representations of information during retrieval (Anderson, 1993; Lipshitz, Levy, & Orchen, 2006). Taken together then, the retrieval mechanisms incorporated in our theory are consistent with team knowledge-building activities directed towards developing a comprehensive and integrative understanding of a problem domain among all team members.

Sharing. During collaborative knowledge-building, direct communication between members represents one of the few readily observable knowledge-building behaviors performed by individuals. The act of transmitting information known by a team member to one or more team members constitutes *sharing*. For purposes of characterizing what and how individuals share information in our theory, we recognize two important distinctions. First, sharing represents any overt act intended to disseminate information to others. Sharing may thus occur through

different mediums which vary along many characteristics, such as richness or temporal synchronicity (e.g., real-time/synchronous communication through face-to-face conversation versus delayed/asynchronous communication through e-mail; Daft & Lengel, 1986). Variation in the channels through which information sharing occur within a team could impact a variety of processes related to team knowledge emergence, including data selection (what gets attended to), member selection (who communicates), and decoding (how readily shared information is internalized). Precisely elaborating how differences in sharing mediums correspond to differences in these mechanisms lies beyond the scope of our theoretical specification. Nevertheless, such considerations exemplify the potential which theories of emergence hold for exploring complex and highly interdependent relations (Davis, Eisenhardt, & Bingham, 2007; Harrison et al., 2007).

Second, the quality of information which gets shared differs in teams whose members possess heterogeneous versus homogenous expertise (Hinsz et al., 1997). To the extent that team members rely on one another's unique proficiencies to provide and reliably interpret information from the task environment, communicating basic, decontextualized facts is often inefficient and unproductive (e.g., "That helicopter won't work well because it's a Bell UH-1 Iroquois"). Instead, teams with distributed expertise must often share integrated knowledge that facilitates understanding by members who do not possess the background knowledge to otherwise do so (e.g., "Based on the weather conditions, this helicopter will need to refuel twice before it can reach the target," cf., Fiore, Smith-Jentsch, et al., 2010; Hollenbeck, Ilgen, Sego, Hedlund, Major & Phillips, 1995). This distinction is incorporated in our theory of team knowledge emergence as a preference for team members to retrieve and communicate integrated knowledge rather than disparate facts during sharing attempts.

Acknowledgement. The act of sharing information within teams provides the means by which individuals can impact and contribute to one another's knowledge about a given domain. However, teams which engage only in one-way communication or do not provide feedback that information was received often fail to develop a shared understanding of the environment (Salas, Sims, & Burke, 2005). The final critical process in the emergence of team knowledge thus involves members' *acknowledgment* of receiving and internalizing information shared with them by other members. This mechanism facilitates team members' understanding of "who knows what" (e.g., transactive memory, Wegner, 1987) as well as recognition of what information has yet to be addressed (e.g., shared mental models, shared situational awareness, Klimoski & Mohammed, 1994; Salas, Prince, Baker, & Shrestha, 1995).

Previous descriptions of team information exchange have conceptualized acknowledgement processes in different ways. For example, Fiore and colleagues' characterize acknowledgement as the creation of cognitive

artifacts (e.g., reminders, notes, tables, etc.) that serve as a tangible externalization of acquired information (Fiore, Rosen, et al., 2010; Fiore, Smith-Jentsch, 2010). Alternatively, Crew Resource Management (cf., Kanki, Helmreich, & Anca, 2010) and other more general group information processing models (e.g., Hinsz et al., 1997) treat acknowledgement akin to confirmatory feedback and closed-loop communication. Our theory of team knowledge emergence integrates both these conceptualizations. More specifically, we posit that acknowledgement only occurs *after* a given piece of information has been fully decoded; in this respect, it reflects closed-loop communication that an individual has internalized information received from another team member. However, the act of acknowledgement also generates new externalized knowledge among team members who have also internalized that information. This treatment is consistent with creating a cognitive artifact that can only be produced after at least two members have exchanged information. This specification also presumes that individuals are capable of keeping track of what knowledge has been received and internalized by other members. Acknowledgement thus captures the inter-member communication exchanges between members that signal information shared by one individual has been received, comprehended, and taken into account by other members of the team.

Process-Oriented Theory of Knowledge Emergence in Teams

Having elaborated the core concepts and mechanisms of team knowledge emergence, the manner by which these processes interact dynamically to drive team knowledge-building can now be specified. Figure 2 depicts our theoretical process model of knowledge emergence in teams. The left and right panels distinguish between learning and sharing processes (respectively). The double-line horizontally bisecting the panels differentiates mechanisms and emergent outcomes at the individual (bottom of panel) and team (top of panel) levels. The solid arrows linking team knowledge concepts trace the process pathways associated primarily with learning activities, whereas the dotted arrows correspond with processes that occur during sharing. Lastly, the white box labeled *Emergent Team Knowledge* depicts how team knowledge outcomes can be operationalized as they develop over time and level.

Learning. Individuals engage in learning by first identifying and directing attention towards information that is available and relevant to their task environment (*Data Selection*). Once a piece of information has become the focus of attention, individuals can attempt to internalize it into privately held knowledge (*Encode*). As members attend to and continue encoding more information, relational associations among pieces of learned information can be formed that reflect deeper knowledge of the task environment (*Integration*). Note that throughout these activities, individuals do not benefit from the learning carried out by other team members; that is, learning occurs independently across members. However, as individuals increase their personal level of internalized knowledge, the team's pooled

knowledge as well as the manner by which that knowledge is distributed across members (i.e., proportion of information known by one or more persons) changes. Thus, learning processes in team knowledge-building contexts are conceptualized as repeated cycles of attending to, encoding, and developing an integrative understanding of knowledge that is performed in parallel by multiple team members. The result of this process is the generation of internalized knowledge held privately by members within a team.

Sharing. Sharing processes are initiated when an individual in the team elects to communicate with other members to disseminate information (*Member Selection*). In doing so, the individual selects information that has been internalized during learning (*Retrieval*). Members on teams with distributed expertise rely on the specialization of others to translate decontextualized information into more readily interpretable knowledge; consequently, members prioritize communicating integrated internalized knowledge to others on the team (*Share*). This information exchange attracts attention from receivers of the shared information (*Data Selection*) and prompts new efforts by these members to internalize the communicated knowledge (*Decode*). However, learning-from-others is a more cognitively demanding process compared to learning-from-self as a member's ability to attend to information shared by another teammate is restricted by both task specialization and expertise barriers (i.e., only certain members can reliably access or interpret information due to differences in role, training, location, background, etc.). As such, members may need to engage in multiple communications before shared information can be adequately learned and transformed into internalized knowledge by the receiving individual. Once a piece of information has been received and fully learned, the receiving team member(s) can signal that they have internalized the shared knowledge (*Acknowledge*). The explicit recognition that shared knowledge has been acquired marks the transition from privately held to publically known and actionable knowledge within a team. Sharing processes thus involve members in a team shifting dynamically between communicating, receiving/interpreting, and acknowledging information. The result is emergence of externalized knowledge held collectively across multiple team members.

Emergent team knowledge. The above characterization presents members' engagement in learning and sharing activities as a *process* of knowledge emergence in teams. However, it is also important to conceptualize team knowledge as an *emergent construct* which accumulates and can be quantified. To this end, we draw from Fiore, Rosen et al.'s (2010) macrocognitive framework and Kozlowski and Chao's (2012a) team knowledge typology to operationalize emergent team knowledge outcomes. We posit that a knowledge pool can be defined for any given task environment which consists of all relevant information within that environment. A piece of information from the knowledge pool can then be characterized by the degree to which it has been learned by one or more members on

the team (degree of overlap) and whether that information is held privately or is acknowledged collectively (internalized vs. externalized). Using this framework, the emergence of team knowledge as a construct can be quantified as the proportion of and rates of change in knowledge distribution within a team as information transitions from internalized to externalized through learning and sharing processes (Kozlowski & Chao, 2012a). An example of this conceptualization is presented in the circular diagram labeled *Knowledge Pool* at the top of Figure 2 and explained in greater detail in the accompanying note.

STEP 2: COMPUTATIONAL MODEL OF KNOWLEDGE EMERGENCE IN TEAMS

Having elaborated a narrative process theory of knowledge emergence in teams, we next translate these notional mechanisms into a computational model (Kozlowski et al., 2013). Developing theory through the use of computational modeling techniques has been described as holding great potential, yet they are rarely utilized in the organizational sciences (Ilgen & Hulin, 2000). A computational model is “a precise formulation of the processes through which the values of variables change over time based on theoretical reasoning” (Harrison et al., 2007, p. 1232). More colloquially, a computational model is a formal declaration of what, how, and when events or actions happen. The formalism of a computational model involves specifying basic algorithms (e.g., if X, then Y) and/or mathematical equations to provide a transparent account for the proposed rules characterizing a dynamic system (Law & Kelton, 1991; Miller & Page 2007). The advantage of translating a narrative theory of process into a computational model of process is the opportunity to evaluate the logic, consistency, and sufficiency of a theory’s core concepts and mechanisms (Epstein, 1999; Taber & Timpone, 1996). Computational modeling is thus particularly useful for investigating complex phenomena that involve multiple elements interacting over time to yield emergent outcomes (Kozlowski et al., 2013; Vancouver, Tamanini, & Yoder, 2010; Vancouver & Weinhardt, 2012).

Translating our theory of team knowledge emergence into a computational model entails specifying (a) a set of procedural rules dictating how, what, and when core learning and sharing process mechanisms are carried out by individuals in a team, and (b) how these processes lead to changes in internalized and externalized knowledge over time. The foundational architecture for these mechanisms is provided by the narrative process theory described in Step 1 and visualized in Figure 2. The critical development in Step 2 is provision of a complete and transparent elaboration of the proposed mechanics (e.g., How is data selected during learning? What knowledge is retrieved during sharing?) and temporal dynamics (e.g., When does internalized knowledge change? When do individuals shift between learning and sharing?) associated with the core learning and sharing mechanisms described in our theory. Table A1 in Appendix A fully summarizes the computational model used to operationalize our theory of team

knowledge emergence. This procedural algorithm provides a representation of how and when the core learning and sharing processes proposed in our theoretical framework are enacted by individuals over time. In other words, it formally describes the “process engine” that constitutes team knowledge emergence in our theory. Notably, this model affords the opportunity to systematically explore the logic of our process theory as well as examine how and where variability in its core mechanisms may influence patterns of individual and team knowledge emergence.

STEP 3: VIRTUAL EXPERIMENTATION AND KNOWLEDGE EMERGENCE IN SIMULATED TEAMS

Steps 1 and 2 focused on developing a formal, process-oriented theory of knowledge emergence in teams; in Step 3, we begin to evaluate the theory and examine its utility to generate useful insights into the dynamics of team cognition through virtual experimentation and simulation. *Virtual experimentation and simulation* involves translating a computational model into computer code that is then systematically manipulated (e.g., varying construct values, relational functions, etc.) to explore how dynamic events unfold in the theoretical space (cf., Carley, 2001; Vancouver & Weinhardt, 2012). There are many simulation approaches that could be used to instantiate computational models into computer code (see Davis et al., 2007, Harrison et al., 2007). We utilize *agent-based simulation* (ABS) to represent the processes of team knowledge emergence. ABS is useful for modeling phenomena in which multiple individuals (i.e., agents) situated in a social system influence one another through their interactions. Agents in an ABS abide by simple rules, functions, and if-then statements that specify how, what, and when to engage in particular activities. Emergent constructs (e.g., individual and team knowledge) arise from agents’ repeated engagement in these activities over time (e.g., learning and sharing information; Epstein, 1999; Macal & North, 2010; Reynolds, 1983). Notably, the process of lower-level interactions yielding higher-level outcomes over time is consistent with descriptions of emergence in multilevel theory (Kozlowski & Klein, 2000) and the nature of higher-level constructs in organizational systems (Cronin et al., 2010; Katz & Kahn, 1978; Morgeson & Hofmann, 1999).

In this step, we instantiate the underlying architecture of our process theory (Figure 2) and computational model (Table A1) into an ABS to investigate how changes in key generative processes identified in our process-oriented theory of team knowledge emergence lead to different patterns of emergent team knowledge (process → outcomes). Virtual experimentation and simulation can serve multiple purposes in theory development (see Harrison et al., 2007, for a summary of uses for simulation in theory-building). We use ABS to pursue two specific objectives: *explanation* and *prescription*. Our first objective is to evaluate the logical coherence of our process-oriented theory by assessing the degree to which changes in critical learning and sharing mechanisms generate differences in simulated emergent team knowledge outcomes (i.e., explanation; Epstein, 1999; Taber & Timpone, 1996). This is a

unique advantage of virtual experimentation as it allows us to explain whether and how the dynamic learning and sharing *processes* identified in Step 1 and formally instantiated in the computational model in Step 2 contribute to team knowledge *outcomes* (processes → outcomes). Our second objective is to use the simulated results to explore ways in which learning and sharing processes could be shaped in real teams to improve knowledge emergence. That is, we extrapolate from the simulated results to posit how knowledge outcomes in teams could be enhanced (i.e., prescription). Together, these goals permit the development of theory- and model-derived propositions about what teams and individuals could do to improve knowledge outcomes and why those activities facilitate team cognition.

To this end, we focus our virtual experiments and simulations on encoding and decoding processes as well as how member selection and information sharing are carried out by individuals. These mechanisms are central to all narrative theories of team knowledge-building (e.g., Fiore, Rosen, et al., 2010; Hinsz et al., 1997) and are common targets for intervention in research and practice (Wildman et al., 2012). As such, they represent key diagnostic indicators of the logical consistency of our computational specification (i.e., differences in encoding/decoding and communication processes should manifest as differences in emergent team knowledge) while simultaneously revealing potential chokepoints in team knowledge emergence that could be targets for future empirical studies. To facilitate comparisons between the ABS and subsequent empirical study in Step 4, we characterize the manipulations enacted in the virtual experiments as efforts to influence the *information processing skills* and *communication skills* of agents. As shown by the dark shaded boxes in Figure 2, information processing skills impact the proficiency with which agents encode and decode information (i.e., learning), while communication skills influence which and when agents convey learned information to teammates (i.e., sharing).

In addition to evaluating these core learning and sharing process mechanisms, we also examine the impact that *degree of specialization* across agents within a team exerts on team knowledge emergence. This variable captures the extent to which members possess greater versus less distributed expertise and is thus a structural factor of the team's environment and composition. Consistent with existing theory, degree of specialization is not proposed to directly influence the learning or sharing processes of individuals. Instead, it reflects an environmental condition within which individuals carry out learning and sharing that should influence the trajectory of team knowledge emergence (Lu et al., 2012; Mesmer-Magnus & DeChurch, 2009; Stasser & Titus, 1985, 1987).

METHODS

Description of Agent-Based Simulation

Tables A1 and A2 in Appendix A list the procedural rules and assumptions (respectively) that all agents followed in the ABS. In brief, the goal of agents was to fully internalize and externalize all available knowledge in the task environment. To do so, agents cycled between phases of learning and sharing activities. Each phase lasted for the same period of time, and agents carried out the sequence of core processes associated with that phase (see Figure 2 and Table A1). Similar to real knowledge-building teams, it was possible for agents to reach a point where they could no longer acquire information during learning phases (i.e., agents acquired all possible knowledge that could be internalized without input from other agents). In this situation, the team remained in a sharing phase until every agent had acquired all available information.

The passage of time in the ABS was marked by agent actions. Each agent could perform only a single action per time step, after which the simulation advanced by one iteration.¹ Agents' behaviors were organized by phases of activity that determined what actions they could perform; actions in the learning phase included either encoding or integrating knowledge, whereas sharing phase actions included either communicating or decoding shared knowledge. Agents cycled through each phase of activity together to acquire knowledge. No limit was placed on the amount of time agents had to acquire the knowledge pool, and thus every agent eventually acquired all knowledge in the task environment. This specification was purposefully intended to reflect an "ideal" world in which simulated teams performed under no time pressure, were temporally synced, and perfectly learned and shared information (e.g., agents did not forget learned information, always heard one another, etc.). If the goal of the ABS was to replicate knowledge emergence patterns from existing data, alternative parameterizations or mechanisms would likely be necessary to represent "imperfections" in agent learning and sharing behaviors. However, removing these considerations and modeling a "simpler" world maintains model parsimony and is particularly valuable when the goal of research is to evaluate the logical coherence of new theory and explore how emergent outcomes arise in response to systematic variations in key variables and boundary conditions (Harrison et al., 2007; Epstein, 1999).

Description of Virtual Experiment

The virtual experiment consisted of a 3 (information processing skill) x 3 (communication skill) x 3 (degree of specialization) fully crossed factorial design, resulting in 27 unique conditions. A total of 3000 teams were simulated with approximately 111 teams per condition; this sample size allowed for variability in knowledge emergence to arise within conditions while also achieving between-condition stability in results. Teams consisted of three simulated

agents, each of whom could communicate with all other team members. The total knowledge pool which agents attempted to acquire contained 36 pieces of information and was held constant across all simulation runs. A knowledge pool of this size permitted flexibility to vary the degree of specialization within a team while still keeping the total amount of knowledge to be learned equivalent across teams and conditions. Each learning and sharing phase lasted for 50 time steps. Thus, agents could make a maximum of 50 encoding or integration actions during a learning phase, after which point a new sharing phase would begin; similarly, each agent could make a maximum of 50 communication or decoding attempts during a sharing phase before switching to a new learning phase. Permitting 50 steps per phase ensured that even teams who were poor at learning and sharing knowledge would eventually acquire the entire knowledge pool within a reasonable number of simulated iterations. The computer code for the ABS and virtual experiment was programmed in R (R Core Team, 2016).²

Information processing skill. To operationalize individual differences in how quickly knowledge is internalized, agents were assigned an information processing skill level (high, moderate, low) that determined the rate at which they encoded and decoded information. Information processing skill was manipulated at the team-level such that all agents within a team possessed a similar skill level. An agent's encoding rate reflected the number of repetitions required to fully internalize a piece of information that could be accessed without input from any other agent (i.e., learning-from-self); alternatively, the decoding rate reflected the number of repetitions required to fully internalize a piece of information shared by another agent (i.e., learning-from-others). The encoding rate for each agent was determined by randomly sampling an integer between 3-5 repetitions for an agent in the high information processing skill condition, 6-8 for an agent in the moderate skill condition, and 9-11 for an agent in the low skill condition. Decoding rates were set proportionately to encoding rates, but were slower. The magnitude of this difference was informed by previous research reporting an average effect size of $d \sim .50$ for knowledge test performance between individuals who learned through verbal instruction (most consistent with learning-from-others/decoding) and those who learned through self-directed reading (most consistent with learning-from-self/encoding, Corey, 1934; Spencer, 1941; Russell, 1928). Decoding rates for an agent were thus increased by one repetition relative to its respective encoding rate; thus, it took an agent slightly longer to internalize a piece of information when learning it from a team member. Integration rates were held constant across ability level such that integrating data to create internalized knowledge required only one repetition for all agents.

Communication skill. Differences in how frequently agents communicated with one another were operationalized by assigning each agent a speaking probability indicating the likelihood they would elect to speak and

share information at each time step during a sharing phase. Communication skills thus influenced which member was selected to speak at each time point and therefore what information could be disseminated to team members. This variable was manipulated at the team-level to construct three experimental conditions. In the first condition, the agent with the lowest information processing skills on the team was twice as likely to speak as the remaining agents. In the second condition, all agents had the same probability of speaking. Finally, in the third condition, the agent with the highest information processing skill was twice as likely to speak as its team members.³

Degree of specialization. The final experimental manipulation concerned the degree to which information was specialized in the team's environment. Degree of specialization reflected the relative proportion of a team's knowledge pool that was "common" versus "unique" to agents. Common information represents information available and interpretable by all individuals without input from any other members; in contrast, unique information is only available and interpretable by a single team member and therefore must be shared for others to learn. Thus less unique/more common information is characteristic of lower specialization in the task environment, whereas more unique/less common information is characteristic of greater specialization. Degree of specialization was manipulated at the team-level to create three conditions that varied the proportion of unique versus common information present in the team's task environment. In the low specialization condition, 33% of the information to be learned by agents was unique while the remaining 66% was common. In the moderate specialization condition, 50% of the information was unique and 50% was common. Finally, the knowledge pool in the high specialization condition was composed of 66% unique and 33% common information.

Measures

Kozlowski and Chao's (2012a) team knowledge typology served as the foundation for measuring emergent team knowledge outcomes. The typology proposes a number of potential metrics for quantifying how internalized and externalized knowledge changes at the individual and team levels over time. Since all agents in the simulation eventually acquired all possible knowledge, these metrics were adapted to analyze the total number of actions (i.e., rates) needed by agents to achieve full knowledge acquisition. The measures described below thus provide different vantage points for conceptualizing the effectiveness and efficiency by which knowledge emerged in the team.

Individual Internalization. Individual internalization describes the proportion of the total knowledge pool that an agent has personally internalized. Given differences in their information processing and communication skills, different agents could possess different amounts of internalized knowledge at any given point within a trial. Since all agents eventually acquired 100% of the knowledge pool, individual internalization was operationalized as the number

of actions an agent needed to fully internalize the entire knowledge pool. Larger values on this measure indicated an agent took longer to acquire all the available knowledge through learning and sharing.

Team Internalization Variability. Team internalization variability reflects within-team differences in knowledge acquisition trajectories. Team internalization variability was thus operationalized as the within-team variance in the individual internalization rates of agents in each team. Larger values on this measure indicated that agents within the team tended to acquire information at different rates.

Team Knowledge Coverage. Team knowledge coverage describes the proportion of the total knowledge pool acquired by a team as a whole. Achieving total team knowledge coverage does not mean that every *individual* has internalized the entire knowledge pool; rather, it indicates that the *team as a collective* has internalized all possible knowledge (i.e., there are no empty “wedges” in the Knowledge Pool diagram shown in Figure 2). This metric was computed as the number of actions required by a team to collectively internalize the complete knowledge pool. Larger values on this measure indicated that an agent team was less effective at carrying out learning processes needed to acquire the information available to each agent.

Internalization Distribution. Internalization distribution characterizes the degree to which one or more members in a team have internalized the same pieces of information. At any given time, a single piece of information may be internalized by only one agent (*Internalized: Non-overlapping*), some but not all agents (*Internalized: Partially Overlapping*), or all agents (*Internalized: Fully Overlapping*; see Figure 2). Because all agents eventually internalized all pieces of information, this variable captured the number of actions required for a team to achieve fully overlapping internalization across all pieces of information in the knowledge pool. Larger values on this measure indicated that the agents on a team were less effective and efficient at both acquiring and distributing information to other agents.

Externalization Distribution. Similar to internalization distribution, externalization distribution describes the degree to which members in a team have externalized the same pieces of information. At any given time, a piece of information in the knowledge pool can be externalized by either some but not all agents (*Externalized: Partially Overlapping*) or all agents (*Externalized: Fully Overlapping*; see Figure 2). This measure captured the number of actions required by a team to achieve fully overlapping externalization across all pieces of information in the knowledge pool. Larger values indicated that agents were less efficient and effective at distributing and acknowledging internalized information from one another.

Analyses

Individual internalization was measured at the agent-level; thus, multilevel random coefficient modeling (MRCM; Level-1: agent, Level 2: team) was used to examine the impact of the experimental manipulations on this dependent measure. The information processing skill, communication skill, and degree of specialization conditions were modeled as dummy variables in the Level-2 equation (Appendix B provides the full MRCM specification). All other outcome variables (team internalization variability, team knowledge coverage, internalization distribution, and externalization distribution) were assessed at the team level. Univariate analysis of variance (ANOVA) tests that included information processing skill, communication skill, and degree of specialization as between-group factors were performed to evaluate the effect of each manipulation on the team-level dependent measures.⁴

RESULTS

Information Processing Skill

The MRCM analyses for individual internalization revealed that agents on teams whose members all possessed high levels of information processing skill were significantly faster at internalizing the knowledge pool relative to agents on teams composed of members with moderate ($b = 7.98, p < .001$) or low ($b = 16.38, p < .001$) skill levels (Table 1). Evaluation of the ANOVA results indicated that differences in teams' information processing skills exerted large effects on team knowledge coverage (partial $\eta^2 = .654, p < .001$), internalization distribution (partial $\eta^2 = .887, p < .001$), and externalization distribution (partial $\eta^2 = .885, p < .001$; top row of Table 2) such that teams composed of agents with better encoding/decoding rates tended to acquire the entire knowledge pool collectively as well as achieve full knowledge internalization and externalization more quickly. Additionally, Table 2 reveals a small effect of team information processing skill on team internalization variability (partial $\eta^2 = .072, p < .001$) such that teams composed of agents with higher learning rates tended to have slightly higher variability among individual internalization rates.⁵ In sum, the overall pattern of results indicated that knowledge emergence at both the individual- and team-level was enhanced when the encoding and decoding rates of team members were higher.

Communication Skills

The MRCM results for communication skills revealed that on teams where the member with the highest information processing skill spoke most frequently, agents took only slightly longer to individually internalize the full knowledge pool than agents on teams where the least skilled member spoke more often ($b = 0.48, p < 0.001$, Table 1). However, agents tended to internalize the knowledge pool much faster when all members spoke equally ($b = -8.29, p < 0.001$, Table 1). This pattern emerged because agents engaged in information sharing could not

simultaneously be engaged in learning. Since only one member could speak at a time, teams in which one individual communicated more frequently tended to impede other members from disseminating their knowledge as quickly to the team. In contrast, when information sharing was more balanced, all members were more likely to regularly participate in sharing and decoding, thus leading to quicker individual-level knowledge acquisition.

Consistent with these individual-level findings, the team-level ANOVAs indicated that teams whose agents were equally likely to engage in information sharing were also faster at fully internalizing (partial $\eta^2 = .912$, $p < .001$) and externalizing (partial $\eta^2 = .991$, $p < .001$) knowledge than teams where one agent was more likely to speak than others (Table 2, second row). Communication skills also exhibited a small impact on team internalization variability (partial $\eta^2 = .026$, $p < .001$) such that differences in knowledge acquisition rates were highest in teams whose agents had equal speaking probabilities, but did not differ when one member was more talkative than others. This pattern of results occurs because when one member tends to communicate more than others, the remaining members are more often engaged in decoding and are thus internalizing information at similar rates. On a team with only three members, this results in similar learning curves for the majority of the team. Alternatively, when members share information equally, different members learn at different times, leading to the potential for slightly higher within-team variability in internalization. Lastly, communication skills had no influence on team knowledge coverage (partial $\eta^2 = .000$, $p = .994$, Table 2). This result was unsurprising given that teams could achieve collective coverage of the entire knowledge pool entirely through individual learning activities. That is, achieving complete team knowledge coverage only requires each member to learn his/her own unique information and for one or more members to also learn the available common information—all of which can be completed without communication. Taken together, results from the communication skills manipulation indicate that team knowledge emergence was facilitated when information sharing was more balanced across members.

Degree of Specialization

The MRCM analyses and Table 1 show that the effect of team specialization on individual internalization was such that agents were less efficient at internalizing knowledge as the overall percentage of unique information increased (50% unique: $b = 3.82$, $p < 0.001$; 66% unique: $b = 16.43$, $p < 0.001$). The ANOVA results reveal a similar trend for the team-level metrics; higher proportions of unique information decreased the rate at which teams achieved collective team knowledge coverage (partial $\eta^2 = .832$, $p < .001$) as well as fully internalized (partial $\eta^2 = .914$, $p < .001$) and externalized (partial $\eta^2 = .913$, $p < .001$) knowledge (Table 2, bottom row). Higher proportions of unique information also led to slightly higher variability in internalization rates (partial $\eta^2 = .143$, $p < .001$).

Interestingly, the simulation findings suggested that the impact of team specialization on knowledge emergence may not be linear in nature. Both the MRCM and ANOVA results consistently revealed that differences in the amount of time required to internalize and externalize knowledge when the proportion of unique information increased from 50% to 66% was larger relative to when the proportion of unique information increased from 33% to 50%. In other words, as the degree of specialization across team members increased, the “bottlenecks” created by sharing processes became increasingly more difficult to overcome. The overall pattern of findings thus reflect that as teams become more distributed in their expertise, communication—and the slower, more difficult information sharing and decoding processes—play an increasingly significant role in team knowledge emergence.

DISCUSSION

The purpose of Step 3 was to (1) evaluate the logical coherence of our theory of team knowledge emergence and (2) identify prescriptions for how individual- and team-level knowledge emergence could be enhanced in real teams. To accomplish these goals, the core concepts and mechanisms of team knowledge emergence proposed in our theory were instantiated into an ABS and used to conduct a virtual experiment in which key learning and sharing processes (information processing and communication skills) of agents as well as a structural team variable (degree of specialization) were systematically manipulated. With respect to the first objective, the simulated results were consistent with conceptual and empirical work on knowledge-building in teams. Previous theory posits that individual differences in learning associated with either task role or person should be positively related to team knowledge development (e.g., Fiore, Rosen et al., 2010, Hinsz et al., 1997). This proposition was replicated in the simulations as teams whose agents possessed better information processing skills (i.e., faster encoding/decoding rates) were more effective at acquiring individual- and team-level knowledge.

The results for communication skills (i.e., who shares information and therefore what information gets shared) and degree of specialization were similarly coherent with research on “hidden profiles” in team decision-making contexts. That research tends to find that the distribution of unique information across team members as well as variability in communication rates strongly influence team cognition (e.g., Lu et al., 2012; Mohammed & Dumville, 2001). However, previous work in this area often concludes that knowledge inefficiencies in contexts with distributed expertise emerge because members *prefer* to communicate commonly held information (e.g., Stasser & Titus, 1985, 1987). Our simulations, though, provide an alternative explanation not based on member preferences or biases—a common critique of the hidden profile literature (Wittenbaum, Hollingshead, & Botero, 2004). Specifically, our results demonstrate that chokepoints in team knowledge emergence can occur when acquiring knowledge through sharing

processes is more difficult for individuals (i.e., slower decoding relative to encoding rates) and if members are not well coordinated in their communication activities (i.e., single member tends to dominate conversation). Furthermore, greater differentiation in member specialization requires individuals to rely more on decoding and communication processes to acquire knowledge, thus creating greater potential for inefficiencies during team knowledge-building. In sum, consistencies between the simulated results and those in the broader team cognition literature provide evidence that the process mechanisms specified in our theory can account for characteristic patterns of emergent knowledge outcomes in teams. The results also revealed interesting insights into the coordination demands underlying information sharing, and showed that inefficiencies in knowledge emergence can arise even when members do not possess biases towards discussing particular types of information.

With respect to the second goal of Step 3, the results of the information processing and communication skill manipulations highlight two prescriptions for improving team knowledge emergence in real teams. First, efforts to increase the speed and accuracy with which individuals internalize knowledge (i.e., encode and decode information) are likely to improve team knowledge outcomes. Improving team members' information processing skills should benefit team knowledge emergence by (a) reducing the time needed to acquire specialized knowledge during more independent learning phases of activity, thereby allowing teams to focus more on coordinating more demanding and interdependent information sharing processes; and (b) reducing the number of communication attempts needed to internalize information shared by others and thus improving the efficiency of team knowledge externalization. A second prescription involves coordinating how and when information is shared within teams. In the simulations, the best form of communication occurred when agents shared information at equivalent rates. In real teams, similar strategies that balance the regularity with which individuals engage in communication should be beneficial because they help members know when to expect incoming information, minimize process loss resulting from having to repeat missed information, and promote more rapid acquisition of specialized knowledge across all members.

These observations characterize how team knowledge outcomes could be enhanced by targeting key *processes* of team knowledge emergence. That is, the theory of team knowledge emergence summarized in Figure 2 and explored in the ABS offers a plausible account for *how* and *why* team knowledge outcomes are generated (process → outcome). The factors which potentially influence these processes (inputs → process) will vary widely and could include individual differences, training, work design, task characteristics, and group composition, among others. Consequently, the simulation results and proposed model are not intended to suggest specific methods, interventions, or antecedents that correlate with better team knowledge outcomes. Instead, the proposed theory

works to unpack the black box of team cognition by identifying core mechanisms of knowledge emergence and elaborating how these processes generate collective knowledge. Such foundational theory affords future research the opportunity to not only explore inputs that influence a team's capacity to carry out these fundamental processes, but also provide precise explanations for why such factors are effective and how they operate.

The model of team knowledge emergence summarized in Figure 2 and enacted in the ABS can be used to explore either general or context-specific prescriptive guidance for enhancing knowledge emergence in real-world teams. Given that we were not attempting to address knowledge emergence in a specific team or task setting in this initial conceptualization, we advance the following general prescriptions for enhancing knowledge emergence in real teams based on our theory and the results of the ABS:

Proposition 1: Improving the rate at which team members encode and decode information (i.e., information processing skills) facilitates team knowledge emergence by allowing individuals to more effectively internalize knowledge overall as well as devote more time to acquiring information shared by others.

Proposition 2: Information sharing strategies (i.e., communication skills) that promote more balanced and regular exchange across members facilitates team knowledge emergence by increasing the rate at which specialized knowledge is disseminated and collectively acknowledged within the team.

STEP 4: EMPIRICAL INVESTIGATION OF KNOWLEDGE EMERGENCE IN REAL TEAMS

The findings from Step 3 support the *generative sufficiency* of our theory of team knowledge emergence; that is, the core theoretical process mechanisms and interactions formalized in the computational model and ABS offer a conceptually plausible account of the macro-level occurrence (Epstein, 1999). A next important litmus test is the extent to which the patterns of emergence and predicted outcomes observed in a simulation share fidelity with the "real world" (Davis et al., 2007; Kozlowski et al., 2013). The empirical experiment pursued in Step 4 thus has two key objectives: (1) investigate the fidelity of knowledge emergence patterns observed in simulated agent teams against comparable human teams, and (2) evaluate the utility of the prescriptions derived from our theory and ABS for enhancing knowledge emergence in real teams.

Attempting to exactly replicate the design and analyses of the virtual experiments pursued in Step 3 represents only one possible strategy for accomplishing these objectives with empirical data. However, replicating an ABS with real data is frequently infeasible and impractical given the number of parameters and parameter levels that could be manipulated as well as the sample sizes that would be needed to adequately examine the full parameter space. Reproducing all possible outcomes from an ABS is fortunately not the only solution for supporting inferences

generated from a computational model. A significant advantage of developing theory through computational modeling is that simulation can be used to do the “heavy lifting” of experimentation so as to identify interesting insights into the generative mechanisms that give rise to patterns, trends, or reactions within a system. These observations can be projected forward to direct empirical investigation towards specific conditions where differences in the core process mechanisms reveal salient differences in emergent outcomes within real systems (Kozlowski et al., 2013). Support for the specification of a process-oriented theory and the validity of its prescriptions is subsequently gleaned by examining whether the emergent patterns observed in real data are similar to those observed in simulated data (Harrison et al., 2007; Wilensky & Rand, 2015).

We thus empirically investigate patterns of team knowledge emergence in human teams collaborating in a knowledge-building and decision-making task which necessitated learning and sharing processes analogous to those carried out by agents in the ABS of Step 3. To investigate the first objective of Step 4 (i.e., assessing fidelity between empirical and simulated patterns of emergent team knowledge), we examine the accumulation of individual and team knowledge as it occurs in real-time in human teams. Focusing on the development of knowledge at this level of analysis allows us to evaluate whether central assumptions of the model and the characteristic patterns of knowledge emergence observed in simulated teams (e.g., are teams slowed when transitioning from learning to sharing phases of activity?) are reflected in the actions of real teams. Producing simulated data capable of mirroring the manner by which knowledge develops and becomes distributed within real teams increases confidence in the conclusion that our theory captures concepts and process mechanisms central to team knowledge emergence.

RQ1: To what extent do human teams exhibit patterns of individual- and team-level knowledge emergence similar to those observed in simulated teams?

Accomplishing the second objective of Step 4 (i.e., assessing the validity of prescriptions for improving team knowledge emergence) necessitates systematically influencing *how* team members carry out the core processes of team knowledge-building. More specifically, examining the model-derived propositions in Step 3 involves shaping members’ learning and sharing processes and evaluating if changes to those processes contribute to differences in team knowledge outcomes. Unlike influencing the behavior of simulated agents though, it is not possible to simply “flip a switch” and immediately improve the information processing and communication skills of human teams. Morgeson and Hoffman (1999) note that such constructs reflect structured cycles of activity that unfold within and between individuals over time. Altering these processes should lead to corresponding changes in individual- and team-level outcomes, but their effects will take time to manifest as members actively reshape their interdependent

behaviors and cognitions. In other words, improving knowledge emergence in real teams requires members to “learn how to learn.” Empirically evaluating the proposed prescriptions for improving team knowledge emergence must therefore permit time for this change process to unfold and for individuals to improve their capacity to engage in effective learning and sharing processes. The theory instantiated in our computational model and the results from the ABS identify precisely *which* processes members should focus on improving as well as *how* and *why* improving those processes will contribute to team knowledge outcomes (process → outcome).⁶

Our strategy thus involves introducing interventions designed to promote better information processing and communication skills in teams engaged in a knowledge-building task.⁷ These interventions targeted the same learning (e.g., encoding/decoding) and sharing (e.g., member selection) processes as those manipulated in the ABS and were designed to make human teams more like the “best” simulated teams. More specifically, the interventions sought to improve the efficiency with which members learned-from-self and learned-from-others as well as promote more equal rates of information sharing among members. Teams performed multiple trials of the knowledge-building task to allow time for these interventions to shape their learning and sharing processes and manifest as differences in emergent team knowledge. The ability to contrast knowledge emergence in teams receiving interventions designed to improve the core process mechanisms against those not receiving similar guidance also provides further opportunity to evaluate the impact of mechanisms on team knowledge development. In sum, teams that improve their capacity to perform the core learning and sharing processes specified in our theory of team knowledge emergence should become more proficient at developing actionable and externalized team knowledge over time.

Hypothesis 1: Teams receiving an intervention to improve information processing and communication skills will be more effective at developing externalized team knowledge relative to control teams. A significant Intervention x Time interaction will be observed such that the amount of knowledge externalized will increase at a faster rate across trials for teams receiving the intervention.

METHODS

Participants

Undergraduate students ($n = 789$) recruited from psychology and management courses participated in the study for course credit. Participants were randomly assigned into three-member teams ($k = 263$ teams), and sessions were conducted with 9 to 18 participants per session. Given the virtual nature of the task environment, team members were not physically seated together nor did they know the identities of the other individuals on their team at

the start of the experiment. Thus, the teams were composed of spatially distributed members who initially had low familiarity and experience working with the other members of their team.

Experimental Task

A computer-based task simulation was developed to examine the emergence of team knowledge outcomes and capture behaviors related to learning and sharing (Kozlowski, Chao, Grand, Braun, & Kuljanin, 2016). The Crisis Relief Operation: Naval Unit Simulation (CRONUS) is a timed team learning and decision-making task in which members with different role specializations must find, learn, share, and synthesize information in order to make accurate decisions. The distribution of information within the task is similar to a hidden profile in that the total available knowledge pool is composed of both common and unique information (Stasser & Titus, 1985, 1987). Each member of the team is assigned an explicit task role (i.e., Transport, Intel, Engineer) that enables them to identify and learn certain information that no other team members can access. Individuals were thus aware they each possessed distinct expertise and sources of information and needed to work collaboratively to fully acquire the knowledge needed to make effective decisions.

The primary objective in CRONUS was for teams to select one of three possible routes shown on a map that could be used to travel between a designated starting and ending location. Making this decision effectively required teams to learn about the presence of obstacles along each route that could impede their travel. To identify and learn about obstacles, each team member had their own specialist map that could only be viewed and accessed by that member. Each member's specialist map contained information about two obstacle classes: *unique obstacles* only accessible to a single member because of their specialized task role (e.g., only the Transport member can initially learn about obstacles related to poor road conditions) and *common obstacles* accessible to all members (e.g., all members can initially learn about obstacles related to mountainous terrain). The various unique obstacles located on members' specialist maps represented the specialized expertise they brought to the team. As a result, no single individual could locate and learn all the obstacles present in a scenario without input from other members.

Once a member identified that a particular obstacle was present, they could "post" this information to their specialist map to document they had learned about the presence of an obstacle. To share information about obstacles with team members, individuals also had access to a commonly shared workspace labeled the mission map. The mission map was identical in appearance to the specialist map except that it initially contained no information about obstacles. However, all team members could post information to the mission map about obstacles they had located using their specialist map at any time. If an individual shared information by posting an obstacle to

the mission map, the other team members had a limited period of time to view that information before it disappeared and needed to be communicated again.⁸ Team members could acknowledge that they had attended to and learned about an obstacle shared on the mission map by viewing the shared post and recording that information on their own specialist map. When all team members acknowledged a particular obstacle by posting it to their respective specialist maps, the obstacle remained permanently visible on the mission map. In this sense, the specialist map provided a snapshot of how much information each member had internalized, whereas the mission map reflected how much information the team as a whole had collectively externalized.

Each scenario contained 15 total obstacles (five per route) representing the total knowledge pool to be acquired by the team. Of those 15 obstacles, three were common obstacles and the remaining 12 obstacles were unique (80% of knowledge pool was unique information, representing a high degree of specialization within the team). The total number of unique obstacles to be learned was distributed equally across members such that, in every trial, each team member was responsible for finding and learning about four unique obstacles related specifically to their task role. On any given trial an individual could thus only acquire information about seven obstacles (four unique and three common) without input from another team member. To acquire the entire knowledge pool (i.e., all 15 obstacles), members needed to share information about the unique obstacles they learned individually with others on the team. As shown in Table 3, the processes required to fully internalize and externalize knowledge in CRONUS were parallel to those specified in the ABS.

Procedure

The experimental sessions were conducted in a large computer lab and lasted 2.5 hours. Prior to arrival, all participants were asked to watch a 10-minute online instructional video introducing the objectives and task mechanics of CRONUS. Once participants arrived at the lab and teams had been assigned, all individuals engaged in a 30-minute interactive guided training session. In addition to reiterating the information presented in the instructional video, the guided training provided more detailed instructions on how to perform specific task functions (e.g., switching between maps, locating and posting obstacles, sharing obstacle information, etc.) as well as opportunities to practice these actions using CRONUS. Following training, participants were given a cover story describing their overall team mission, informed of their unique team roles, and told they would be working together with their team members in the CRONUS simulation.

Teams attempted to complete 12 trials within CRONUS. Each trial lasted 8 minutes and contained the information distribution structure described previously. The configuration and location of obstacles differed across

trials. To ensure that teams progressed in a timely fashion, a countdown timer was displayed on each member's screen indicating the amount of time remaining in the trial. After 6 minutes had elapsed within the trial, members could no longer post obstacles to the specialist or mission maps and could only enter their final decisions. At the conclusion of each trial, teams were provided with feedback on the total number of obstacles learned by the team and each member, the number of information sharing attempts made and missed, and the accuracy of individual and team decisions. After 12 trials had been completed or a 2-hour time limit had been reached, participants were provided with a debriefing form and were free to leave the session.

Experimental Manipulation

The extent to which team knowledge acquisition could be improved through targeting core learning and sharing mechanisms was examined by randomly assigning teams to either a control or experimental condition. Teams in the control condition completed the task trials with no guidance beyond the initial training. For teams in the experimental condition, a set of contextualized guidance (CG) prompts were created and embedded within CRONUS that could be triggered by specific errors or inefficiencies. The CG was delivered through a pop-up window on the screen of the individual who performed the triggering action. When the window appeared, individuals read the CG provided, closed the window, and then could correct the error and continue on with the task.

In total, ten different CG prompts were employed that provided feedback and recommendations to team members for improving their information processing and communication skills.⁹ Table 4 lists each prompt and summarizes the rationale for how these items were intended to improve the core learning and sharing processes of team knowledge emergence. Construction of the CG prompts were informed by the simulation results and propositions summarized in Step 3. They were specifically designed to (a) improve the encoding/decoding rates of individuals and (b) suggest information sharing strategies that promote frequent and equal participation by all members. To allow the experimental condition team an opportunity to gain familiarity with the basic task mechanics in CRONUS, the CG was not introduced until the start of Trial 3. The CG was removed at the beginning of Trial 10.

Measures

A log file containing every task-relevant behavior (i.e., mouse click) performed by each individual was recorded for all trials. Behaviors were automatically categorized and time-stamped, thus making it possible to track precisely when, what, and how knowledge was being acquired and distributed during each trial. The data most relevant for the present study was the posting of any correct obstacle to a participant's specialist map. This action explicitly reflects knowledge internalization in CRONUS, and could only occur if a participant correctly identified an

obstacle by learning its specific type and location from the specialist map (no communication needed from other members to learn) or the mission map (communication needed from other members to learn). Consequently, any of the operationalizations of team knowledge emergence defined in Step 3 could be computed from these data.

Analyses

Fidelity comparisons. A qualitative descriptive approach was adopted to examine RQ1 and the fidelity of team knowledge emergence patterns observed between participant data and data generated using our ABS. More specifically, representative human teams were selected from the empirical data and then compared to an agent team whose information processing skills, communication skills, and degree of specialization were parameterized under conditions similar to those in the human teams. If the general patterns of knowledge emergence observed in these data shared fidelity, then it increases confidence in the conclusion that our model and theory reflects the process mechanisms central to team knowledge-building.

One human team was randomly selected from both the control and experimental conditions during Trial 10 to serve as the representative comparison. The decision to sample data from Trial 10 was made because participants were more likely to have a better understanding of how to acquire and share knowledge through the CRONUS interface during a later trial in the experiment. As a result, the degree of random noise in team knowledge outcomes attributable to unfamiliarity operating and interacting within CRONUS should be attenuated and thus provide a better comparison between the simulated and human team knowledge outcomes. We next constructed a simulated agent team that closely paralleled the conditions under which the human teams performed.¹⁰ Similar to the participant teams' task environment in CRONUS, the comparison agent team was constructed with a high degree of specialization (75% of knowledge pool was unique information). The information processing skills of agent teams were also parameterized to reflect the within-team variability observed in the learning rates of human teams. The simulated team thus possessed one agent with high, one agent with moderate, and one agent with low information processing skills. Lastly, the communication skills of the simulated team were set such that all agents were equally to participate in information sharing (i.e., equal sharing rates). The individual internalization and internalization distribution measures previously described in Step 3 were examined for the fidelity comparisons. The individual internalization measure provides insights into the trajectories of each individual member's knowledge acquisition over time (individual-level), while the internalization distribution measure provides an indication of how knowledge transitioned from individually to collectively held over time as it is disseminated across team members (team-level).

Experimental manipulation. The impact of the CG targeting information and communication skills on team knowledge outcomes across trials for participant teams was analyzed using MRCM to account for the nested data structure (Level-1: trial; Level-2: team). The dependent variable was the total number of correctly posted obstacles which all three members held in common on the mission map at the end of each trial. This team-level metric reflected the amount of fully-overlapping externalized knowledge that had emerged within the team (cf., Figure 2). Because this measure can only increase if all team members learn and share the same pieces of information, it provides the most direct indication of how effectively teams collectively accumulated knowledge in each trial. Hypothesis 1 was assessed by regressing this dependent measure on trial number (Level-1 predictor; 0 = Trial 1) and a dummy coded condition variable (Level-2 predictor; 0 = control, 1 = experimental). Visual examination of the observed data trends also suggested the presence of non-linear changes in the dependent variable across trials; thus a quadratic trial variable was modeled at Level-1 (see Appendix B for full model specification).

RESULTS

Computer recording errors resulted in unusable data for six teams. Additionally, only 212 teams (80%) completed all 12 trials during the experimental session. The final reported analyses are thus based on teams from the control ($k = 110$) and experimental ($k = 102$) conditions who had complete data for all experimental trials.

Fidelity of Simulated and Participant Data

Figures 3 and 4 show changes in individual knowledge internalization and the distribution of internalized knowledge, respectively, for the representative simulated, experimental condition, and control condition teams.¹¹ An overall comparison of the figures reveals a number of notable similarities in patterns of team knowledge emergence between the agent and human teams. Figures 3A and 3B reveal that agents in the simulated team (gray lines) and human members in both the control and experimental conditions (black lines) experienced a relatively rapid period of initial individual internalization during which a sizeable portion of the knowledge pool was acquired. Figure 4A shows that the knowledge acquired during this initial burst was primarily non-overlapping, with a smaller proportion partially (Figure 4B) or fully overlapping (Figure 4C). These results indicate that early knowledge emergence in both the agent and human teams was largely a function of *learning* processes targeted towards acquiring unique (which emerges as non-overlapping information) and common (which emerges as partially and then fully overlapping information as multiple members learn the same thing) knowledge that is available without input from other members.

The accumulation of knowledge in agent and participant teams then slows substantially as members reach the limit for what they can learn by themselves (50% and 47% of knowledge pool for simulated agents and

participants, respectively). At this point, the more time intensive and difficult *sharing* processes began to take precedence. Figure 3 shows that the concentrated shift in emphasis towards sharing processes (occurring near the 30-45 second mark for the selected participant teams and the 120 second mark for the simulated team) led to higher variability in individual internalization rates. This pattern is attributable to two pragmatic constraints imposed on individual knowledge acquisition when teams engage in knowledge sharing: (1) when individuals “speak up” to share knowledge, they cannot also internalize new knowledge, and (2) members require more and more frequent exposures to internalize communicated knowledge. Both of these factors compound the amount of time needed to acquire knowledge during sharing, subsequently decreasing the overall internalization rate of all team members. The common internalization trajectories demonstrated by the experimental and control condition teams in Figures 3A and 3B demonstrate this universal inefficiency of sharing relative to learning processes in team knowledge emergence.

Whereas Figure 3 paints a similar pattern of individual-level knowledge internalization among the simulated, control, and experimental condition teams, changes in the distribution of internalized knowledge over time reveals two critical distinctions in team-level knowledge emergence across these conditions. First, the representative experimental condition team (solid black line) achieved total team knowledge coverage (i.e., all items in the knowledge pool known by at least one member on the team) earlier than the representative control condition team (black dashed line). This is reflected in Figure 4A as the earliest point at which the maximum proportion of non-overlapping knowledge is achieved.¹² This finding is noteworthy because it provides evidence that the CG targeting information processing skills provided to the experimental team improved the speed and accuracy with which individuals acquired unique knowledge compared to the control condition team. Consistent with Proposition 1 from Step 3, this enabled these teams to direct attention towards the more demanding sharing processes sooner—an important consideration given the time constraints faced by participant teams in the experiment.

A second noteworthy distinction concerns the manner by which privately held knowledge transitioned to collectively held team knowledge as individuals began to engage in sharing activities. Once team members began sharing information, the proportion of non-overlapping knowledge in the team decreased relatively quickly (Figure 4A). As expected, this was accompanied by increases in the amount of overlapping knowledge across members; however, whether this increase manifested as partially or fully overlapping knowledge was diagnostic of the team’s information sharing efficiency. For the simulated and experimental condition teams shown in Figure 4B, the proportion of partially overlapping team knowledge remained consistently low, rarely exceeding 20% at any point and reaching 0% by the end of the trial. In contrast, the proportion of partially overlapping knowledge held by the control

condition team varied between 30-40% for nearly two-thirds of the trial. These results indicate that a sizable portion of knowledge in the control condition team became “stuck” in the bottlenecks created by inefficient communication processes that prevented information from reaching all three members. This conclusion is also reflected in the differential rates at which fully overlapping knowledge accumulated for these teams. As Figure 4C shows, both the experimental and simulated teams were more efficient at generating collectively held knowledge compared to the control condition team. Notably, this pattern of results lends support to Proposition 2 derived from the simulated results (Step 3) as it demonstrates that teams which received guidance to engage in more regular and balanced sharing processes were more effective at generating fully shared team knowledge.

In sum, the representative simulated and participant teams demonstrated a number of key similarities in their individual- and team-level knowledge trajectories. This fidelity lends support to the conclusion that the core process mechanisms specified in our theory are capable of accounting for patterns of knowledge emergence in real teams. Differences in the knowledge outcomes observed among the two participant teams also demonstrates that the CG provided in the experimental condition was effective at stimulating improved individual and team knowledge-building. Lastly, it is important to note that the CG was explicitly designed to improve a team’s information processing and communication skills in a manner consistent with the most effective simulated teams and the propositions advanced in Step 3. That the participant team receiving the intervention demonstrated patterns of team knowledge emergence more analogous to the simulated team further supports that the process mechanisms instantiated in the ABS are critical points of leverage for enhancing knowledge emergence in real teams.

Experimental Manipulation

The qualitative assessments demonstrate that the core learning and sharing processes from our theory and ABS appear capable of sufficiently replicating patterns of knowledge emergence observed in real teams. The qualitative comparison of the representative experimental and control condition teams also provided some support for the effectiveness of the CG and the propositions generated from the simulated results for improving team knowledge outcomes. However, stronger support for the propositions derived in Step 3 would be found if teams given the CG designed to improve information processing and communication skills in *each trial* also tended to demonstrate better knowledge outcomes *across* trials (Hypothesis 1). This pattern of results would demonstrate that improving how teams carry out key learning and sharing processes (i.e., improving team knowledge emergence) leads to improved team knowledge outcomes.

The total number of obstacles fully externalized by a team at the end of each trial was analyzed. Prior to testing Hypothesis 1, a panel augmented Dicky-Fuller test was conducted (drift and lag = 1) for each team's data to evaluate the null hypothesis that the time series data resulted from a non-stationary stochastic process (Braun, Kuljanin, DeShon, 2013a; Braun, Kuljanin, DeShon, 2013b; Kuljanin, Braun, DeShon, 2011). The average test value across all teams was -1.69, exceeding the 5% critical value (-1.59) based on the number and length of the time series. Consequently, the null hypothesis was rejected and inferences from the MRCM analyses could be made without qualification.

Results from these MRCM analyses are reported in Table 5. No significant intercept differences were found across conditions ($b = .34$, 95% CI = [-.31, .99], *ns*), though significant main effects were observed for both the linear ($b = 1.10$, 95% CI = [.90, 1.29], $p < .001$) and quadratic ($b = -.07$, 95% CI = [-.08, -.05], $p < .001$) time terms. These findings indicate that the average number of obstacles fully externalized by teams in both conditions tended to be similar at Trial 1 and increased at a slowly decreasing rate on subsequent trials. However, the main effects were qualified by significant interactions between condition and both the linear ($b = -.32$, 95% CI = [-.60, -.03], $p < .05$) and quadratic ($b = .05$, 95% CI = [.02, .07], $p < .001$) time terms. Figure 5 plots the interaction effect as well as the mean number of obstacles fully externalized per trial for teams in each condition. The pattern of results reveals that the amount of fully externalized knowledge acquired by teams in the experimental condition increased at a fairly steady rate each trial with relatively little decline. In contrast, knowledge externalization for teams in the control condition increased at a similar rate early but plateaued near the mid-point of the experiment, at which point the experimental condition teams began to outpace the knowledge acquisition of control teams. Follow-up analyses revealed that significant differences in the number of obstacles learned between conditions first emerged at Trial 9 and that this difference remained significant through Trial 12 (all *t*-tests $p < .05$), thus supporting Hypothesis 1.

DISCUSSION

The results of Step 4 contribute to the evaluation of our theory of team knowledge emergence in two key ways. First, they demonstrate that patterns of knowledge emergence observed in real teams share a high degree of fidelity with the patterns of knowledge emergence produced by simulated agents enacting the core process mechanisms specified in the theory. This is significant because it provides evidence that the *processes* represented in our theory offer a plausible account for *how* emergent knowledge manifests in real teams. For example, the empirical and simulated results both exhibit the bottleneck associated with information sharing which teams must overcome for collectively held knowledge to emerge. This observation is consistent with previous conceptualizations

of team cognition in contexts with distributed expertise which posit that the dissemination of unique knowledge among members requires greater coordination on the part of the speaker and receiver to prepare for, communicate, and interpret new information (e.g., Fiore, Rosen et al., 2010). Our theory also shows that such sharing processes cannot be divorced entirely from learning. When certain roles and/or members of the team experience more difficulty (i.e., are slower at) acquiring their own information or information communicated by other members, the emergence of both individually internalized and collectively externalized knowledge will be delayed as “faster” members are limited by the internalization rates of “slower” members. Additionally, since faster members will be able to learn the maximum amount of information that can be acquired without input from others more quickly, they may be more likely to engage in knowledge sharing attempts earlier and more often. These efforts are likely to be less effective given that slower members may be unprepared to attend to new incoming information.

These findings highlight a second notable contribution from Step 4 by demonstrating the usefulness of a process-oriented theory of team knowledge emergence for explaining and advancing prescriptions to enhance team knowledge emergence. The CG provided to teams in the experimental condition was designed to facilitate effective information processing and communication skills based on the insights and propositions gleaned from the ABS in Step 3. Figures 3 and 4 provide a clear indication of how these process-focused interventions contributed to improved patterns of team knowledge emergence, while Figure 5 demonstrates that shaping the processes of emergence were effective but took time to manifest (Morgeson & Hoffman, 1999). Interestingly, Figure 5 shows a noticeable dip in the number of obstacles fully externalized by teams in the experimental condition when the CG was first introduced at Trial 3. This likely occurred as teams reacted to the CG and attempted to adjust their learning and sharing processes accordingly. In other words, teams were forced to reorient themselves and began “learning how to learn” as a coordinated and efficient unit at this time. The benefit of improving these processes grew increasingly apparent as the knowledge acquisition of experimental teams continued to improve while the control teams plateaued. Taken together, the results observed in Step 4 offer compelling support for the validity and utility of our process-oriented theory of team knowledge emergence.

GENERAL DISCUSSION

Nearly two decades ago, Hinsz et al. (1997) noted a paradigmatic shift towards the consideration of teams as key decision-making and information processing units in the workforce. That perspective has grown to emphasize the critical role of team cognition in supporting effective functioning across a variety of contexts (Mesmer-Magnus & DeChurch, 2009; Wildman et al., 2012). While the team cognition literature has witnessed significant advancements

over this period, few efforts have been made to precisely define and operationalize the fundamental process dynamics underlying the emergence of team knowledge (Mohammed et al., 2012). The lack of a process-level framework for team knowledge emergence significantly limits insights into how and why teams differ in their capacity to generate actionable and collectively held knowledge, as well as develop evidence-based recommendations for improving these outcomes (Kozlowski & Chao, 2012a; 2012b; Kozlowski, 2015).

To address this gap, we developed and evaluated a process-oriented theory of knowledge emergence in teams. This effort was explicitly organized into four steps for investigating dynamic emergent phenomena (Figure 1). In Step 1, we drew from the individual and team cognition literatures to identify foundational concepts and mechanisms to construct a process-level architecture of how, when, and why team knowledge arises from the activities and interactions of individuals over time (Figure 2). We then translated this narrative theory into a formal computational model in Step 2 to provide a precise and transparent account of how these proposed process mechanisms unfold over time to generate change in individual and team knowledge outcomes (Table A1). These two steps together encompass the development of our process theory. The remaining steps generated insights and examined predictions related to team knowledge emergence based on our theory. In Step 3, we instantiated our computational model into an ABS and systematically manipulated key behavioral and environmental factors to explore their effects on patterns of team knowledge emergence. The results of this virtual experiment provided support for the *generative sufficiency* (Epstein, 1999) of our theoretical specification to produce patterns of knowledge emergence consistent with existing research on team cognition. Furthermore, the simulated findings also suggested potential leverage points for improving knowledge emergence in human teams. Consequently, we evaluated the fidelity of the simulated results and tested these theory-derived prescriptions in Step 4 using real teams performing a knowledge-building task. Findings from the empirical study supported the generalization of our theory of team knowledge emergence to outcomes observed in human teams, and further demonstrated how targeting the learning and sharing processes of individuals contributed to improved team knowledge outcomes.

Theoretical Implications

We believe the proposed theory holds important implications for continued efforts at unpacking the “black box” of team cognition as well as more generally advancing the study of dynamic emergent phenomena in the organizational sciences. With respect to the former, the primary contribution of the process-oriented theory and computational model of team knowledge emergence is that it provides an explicit conceptualization of how, when, and why team knowledge manifests. The core concepts and processes summarized in Figure 2 offer an account of

team knowledge emergence that is both conceptually plausible and empirically defensible. We demonstrated that this process-level explanation was useful for explaining and generating prescriptions relevant to improving team knowledge outcomes. The propositions derived in Step 3 and evaluated empirically in Step 4 provide important insights into the mechanics of how and why knowledge emergence occurs in teams effectively. One particularly significant insight from this work concerns the role of information sharing in team knowledge acquisition. Research examining team communication frequently finds that failures to disseminate information to all members are a common contributor to poor team performance and decision-making. Such findings are often attributed to biases or preferences held by members for wanting to discuss commonly held information (e.g., Lu et al., 2012; Stasser & Titus, 1985). However, there is growing evidence that suggests this explanation may only hold for teams under very special circumstances (e.g., low expertise differentiation, only in judgment tasks; see Wittenbaum, Hollingshead, & Botero, 2004). In teams whose members possessed specialized skill sets and clearly definitive goals—features common in many organizational teams with distributed expertise—empirical support for such information sharing biases as explanations for inefficiencies in team cognition is lacking.

Our theory and results offer an alternative explanation for how deficiencies emerge which can influence a team's capacity to fully externalize knowledge that is not based on implicit biases or member preferences. The simulation results showed that how team members communicated during team knowledge-building activities greatly influenced the acquisition of knowledge at both the individual and team levels. Table 2 reveals that agent team's communication skills exhibited the most potent impact on the emergence of both internalized and externalized knowledge. This suggests that *which* and *how often* members share information is as—if not more—critical to the development of team knowledge emergence than *what* information gets shared. Additionally, internalizing specialized information held by other team members is time consuming and requires more effort to successfully accomplish. For example, the exemplar data from human teams shown in Figure 3 reflected that acquiring the final 40% of the knowledge pool through sharing took three to six times longer than acquiring the first 60% of the knowledge pool. Notably, this pattern also emerged within the simulated teams *even though* the difference between encoding and decoding rates was very small (one action). Together, these findings suggest that concerted efforts to elaborate when, how, and why information is exchanged between members as well as delineate factors that can improve decoding/learning-from-others are key considerations for team cognition research.

Explanation and extrapolating prescriptions from simulated results are among the most important objectives when evaluating the relevance of any process-oriented theory (cf., Davis et al., 2007; Harrison et al., 2007; Kozlowski

et al., 2013; Vancouver & Weinhardt, 2012). An additional measure of a theory's utility is its potential to organize previous research and provide guidance for future work. To this end, Figure 6 summarizes a framework for advancing research on team cognition that reinforces two important implications made salient by our theory. First, efforts to improve team cognition will benefit from examining how individual, team, and environmental factors influence and shape the core process mechanisms of team knowledge emergence (input → process). For example, meta-analytic evidence suggests that team cognition can be improved through team training (Salas et al., 2008). However, it is not clear from this body of research what processes are targeted by team training, how knowledge-building processes are changed through team training, the degree to which different interventions are effective at targeting certain process mechanisms, or how critical individual (e.g., member capabilities, personality), team (e.g., goals, leadership), and environmental (e.g., task demands, time pressures) factors interact to impact knowledge emergence. Our theory identifies eight core process mechanisms; theory and research examining the relative influence of these processes on knowledge emergence (e.g., must teams improve all facets of learning and sharing to enhance knowledge acquisition?) as well as understanding the factors that shape and influence these processes will facilitate a more complete understanding of how to improve team knowledge outcomes.

Second, our theory emphasizes the need to pursue more precise and temporally sensitive measures of team cognition that permit its generative dynamics to be examined. Figure 6 summarizes a number of indicators potentially useful for capturing individual and team knowledge emergence based on the measures employed in Steps 3 and 4. Each of these operationalizations has potential to reveal unique insights into team knowledge emergence. For example, examining the rates at which knowledge was internalized revealed stark differences in team member's effectiveness at learning-from-self/encoding and learning-from-others/decoding (e.g., Figure 3). In contrast, examining changes in the distribution of internalized knowledge across members over time (e.g., Figure 4) revealed where chokepoints in knowledge dissemination were not being effectively resolved. Concerted efforts to elaborate when, how, and why information is exchanged among team members will benefit from multiple operationalizations of collective knowledge. More broadly, team cognition research must begin to employ non-traditional methodological techniques and data sources to capture processes and outcomes indicative of team knowledge emergence beyond single-time point aggregate measures (Kozlowski & Chao, 2012b). Understanding the temporal dynamics of team cognition will require concerted efforts to explore the timing, pacing, and frequency of learning and sharing behaviors among team members. Computational modeling offers a powerful approach for simulating observations and integrating research; coupling these efforts with new empirical data sources and analytic techniques (e.g., latent

semantic analysis of transcribed data, email exchange within project teams, etc.; Cooke, Gorman, & Kiekel, 2008; Cooke, Salas, Kiekel, & Bell, 2004) represents a potent combination for unpacking the black box of team cognition. Expanding the methodologies and measurement approaches used to study team knowledge-building processes will be necessary for advancing theory, research, and practice in this domain.

Finally, in addition to implications for the study of team cognition we believe the present work also offers a template for researchers to advance theory on emergent phenomena more generally. The majority of theory in the organizational sciences lacks adequate specification of the dynamic processes inherent in human cognition, affect, and behavior (Cronin et al., 2011; Kozlowski et al., 2013; Lord et al., 2015). The field is ripe for explorations of emergence that can be supplemented with empirical research to target the bottom-up mechanisms in individual, team, and organizational functioning. Advancing process-oriented theories supplements relational, construct-to-construct covariance approaches to theory-building with alternative approaches that target the interactive elements of systems. The four-step procedure presented in Figure 1 and used to structure the present work offers a replicable and accessible meta-approach for integrating narrative theory, computational modeling, simulation, and traditional experimentation to investigate any emergent phenomena in the organizational sciences.

Practical Implications

As with the development of any new theory, advancing strong recommendations for practice from the present work should be tempered pending further study and validation. Nevertheless, we note two promising avenues for enhancing knowledge acquisition within distributed expertise teams. First, the various operationalizations of team knowledge emergence described in Step 3 as well as their use for monitoring patterns of knowledge emergence of real teams in Step 4 demonstrate the potential diagnostic utility of measuring knowledge-building at the process- rather than outcome-level (Kozlowski & Chao, 2012a). For example, an examination of only the individual knowledge internalization trajectories observed in the exemplar experimental and control teams (Figures 3A and 3B, respectively) did not reveal many noticeable differences in knowledge emergence between these groups. However, consideration of the internalization distribution metrics (Figure 4) revealed a number of key differences in how collectively held knowledge emerged within these teams. Integrating a task paradigm similar to CRONUS with a similar suite of metrics as a training and/or assessment tool could thus prove useful for pinpointing deficiencies in a team's learning and sharing processes. These could then be targeted for improvement through instruction or with the development of adaptive feedback tools (cf., Kozlowski, Toney, Mullins, Weissbein, Brown, & Bell, 2001).

Second, results from both the simulated and empirical studies indicate that in teams with a commonly shared goal and whose members are specialized, every effort should be made to promote knowledge sharing by all individuals as early and often as possible. One method for facilitating this goal is by ensuring that members have the proper training and level of expertise to efficiently and effectively complete learning processes (i.e., data selection, encoding, integration). In a team knowledge-building context with even a moderate degree of specialization, a single member who does not keep pace with the team can slow the dissemination and accumulation of collectively held knowledge. A second method involves establishing procedural norms, protocols, and/or reinforcement strategies that encourage teams to regularly share learned knowledge. Depending on the nature of task demands and the pace at which new information appears and dissipates in the environment, the frequency with which teams may need to share information could differ greatly (e.g., emergency medical teams may require many frequent communication exchanges, whereas multidisciplinary project teams may require far fewer exchanges). As such, appropriate windows in which disseminating knowledge is most beneficial should be established and adhered to.

Limitations and Future Directions

The development of our process-oriented theory of team knowledge emergence was guided by a foundational principle of complexity science—efforts to describe emergent dynamics should identify a core set of process mechanisms capable of representing the targeted phenomenon (Epstein, 1999; Miller & Page, 2007; Reynolds, 1987). This tenet places parsimony as the most critical goal for a theory of emergence. The theoretical model summarized in Figure 2 and investigated in Steps 3 and 4 almost certainly falls short in capturing all the possible influences related to team knowledge emergence. However, the goal was to elaborate a small number of *core* concepts and processes that constitute team knowledge emergence, and then evaluate their sufficiency to replicate patterns of knowledge development in teams with distributed expertise. Consequently, we view our theory of team knowledge emergence as “lean” and a point of departure for future research to continue unpacking the dynamics of team cognition rather than a definitive declaration of knowledge production in teams. One of the major advantages of specifying a theory of emergence into a computational model is that it necessitates transparency among one’s theoretical assumptions. With respect to the present research, this means *all* the model specifications summarized in Tables A1 and A2 in Appendix A characterize limitations and boundary conditions of our theory—and therefore opportunities for extending its precision and generalizability.

We identify two directions that are particularly fruitful for future research. First, our model and simulations required team members to be temporally synched in their learning and sharing activities; that is, every agent

transitioned into and out of learning and sharing phases together. However, the behavioral patterns of teams observed in Step 4 suggest this sequence does not always reflect the pacing of activity in real teams. It was not uncommon to observe instances in which one team member was engaged in learning processes (e.g., searching the specialist map, posting to the specialist map, etc.) while another performed sharing process (e.g., posting to the mission map, monitoring the mission map for new posts). Despite these differences, the qualitative comparisons in Figures 3 and 4 demonstrate that the trajectories of team knowledge emergence observed in simulated and human teams were quite similar. This suggests that temporal synchronicity among learning and sharing activities may be only one path towards achieving knowledge externalization and that integrating other mechanisms (e.g., adaptability, monitoring, etc.) could represent a useful theoretical extension.

An additional extension of the theory would be to consider more flexible end states with respect to team knowledge emergence. Hinsz et al. (1997) posit that groups possess processing objectives that define the goals, purpose, and frame of reference for how teams engage in cognitive tasks. The processing objective of teams in the ABS and CRONUS was that every member sought to internalize and externalize all available information in the environment. Although this objective is a feature of many distributed expertise teams engaged in knowledge-building (Fiore et al., 2010), there are contexts in which full knowledge externalization among members may be unnecessary or unproductive. For example, some types of teams exhibit better performance when they possess strong transactive memory systems (i.e., low overlap in internalized knowledge between members coupled with shared awareness of who knows what) as opposed to commonly shared team knowledge structures (i.e., high overlap in internalized and externalized knowledge; DeChurch and Mesmer-Magnus, 2010). Compilational team knowledge may therefore represent a more effective processing objective in certain contexts. Such extensions could be incorporated into our theory by specifying alternative process mechanisms (e.g., filtering data, knowing who to communicate with, etc.) that facilitate the development of both distributed transactive memory systems and shared mental models.

Conclusion

The study of team cognition—and many other similarly emergent phenomena in the organizational sciences—have frequently relied on post hoc examinations of cross-sectional, static data gathered after emergence has already happened (Cronin et al., 2011; Kozlowski et al., 2013; Lord et al., 2015). Unfortunately, this has limited the ability to draw inferences about the dynamic processes that unfold within and between individuals to generate and sustain team knowledge. The present research offers one of the first efforts in the organizational sciences to pair a theoretical model with empirical data to capture a commonly cited but rarely observed tenet of multilevel theory:

changes in team-level outcomes result from changes in individual-level processes. One of the most useful and promising directions for theories of emergence is the opportunity to explore *how* and *why* organizationally relevant outcomes emerge rather than focus only on differences in *what* has already emerged. Such insights enable both researchers and practitioners to achieve a clearer understanding of why particular outcomes tend to occur and where interventions could be directed to facilitate or correct them. We hope that the conceptual and methodological approach exemplified in the present paper can serve as a model for advancing the study of both team cognition and dynamic team phenomena in general.

References

- Anderson, J. A., (1993). Problem-solving and learning. *American Psychologist*, 48, 35-44.
- Anderson, J.R., Bothell, D., Byrne, M.D., Douglass, S., Lebiere, C., & Qin, Y. (2004). An integrated theory of the mind. *Psychological Review*, 111, 1036-1060.
- Anderson, J.R., & Schooler, L.J. (1991). Reflections of the environment in memory. *Psychological Science*, 2, 396-408.
- Baddeley, A., Eysenck, M.W., & Anderson, M.C. (2009). *Memory*. London: Psychology Press.
- Braun, M.T., Kuljanin, G., & DeShon, R.P. (2013a). Spurious prediction of growth in organizational research. *Organizational Research Methods*, 16, 302-330.
- Braun, M.T., Kuljanin, G., & DeShon, R.P. (2013b). Spurious relationships in growth curve modeling: The effects of stochastic trends on regression-based models. In J.M. Cortina & R.S. Landis (Eds.), *Modern research methods for the study of behavior in organizations* (pp. 161-198). New York: Routledge.
- Broadbent, D.E. (1958). *Perception and communication*. London: Pergamon.
- Cannon-Bowers, J.A., Salas, E., & Converse, S.A. (1990). Cognitive psychology and team training: Shared mental models in complex systems. *Human Factors Bulletin*, 33, 1-4.
- Carley, K.M. (2001). Computational approaches to sociological theorizing. In J. Turner (Ed.), *Handbook of sociological theory* (pp. 69-84). New York: Kluwer Academic/Plenum.
- Chase, W.G. & Simon, H.A. (1973). Perception in chess. *Cognitive Psychology*, 4, 55-81.
- Cooke, N.J., Gorman, J.C., & Kiekel, P.A. (2008). Communication as team-level cognitive processing. In M.P. Letsky, N.W. Warner, S.M. Fiore, & C.A.P. Smith, *Macrocognition in teams: Theories and methodologies*. (pp. 51-64). Hants, UK: Ashgate Publishing Ltd.
- Cooke, N.J., Salas, E., Kiekel, P.A., & Bell, B. (2004). Advances in measuring team cognition. In E. Salas & S.M. Fiore (Eds.), *Team cognition: Understanding the factors that drive process and performance* (pp. 83-106). Washington, DC: American Psychological Association
- Corey, S. M. (1934). Learning from lectures vs. learning from readings. *Journal of Educational Psychology*, 25(6), 459-470.
- Cronin, M.A., Weingart, L.R., & Todorova, G. (2011). Dynamics in groups: Are we there yet? *The Academy of Management Annals*, 5, 571-612.
- Daft, R.L., & Lengel, R.H. (1986). Organizational information requirements, media richness, and structural design. *Management Science*, 32, 554-571.
- DeChurch, L.A., & Mesmer-Magnus, J.R. (2010). The cognitive underpinnings of effective teamwork: A meta-analysis. *Journal of Applied Psychology*, 95, 32-53.
- Davis, J.P., Eisenhardt, K.M., & Bingham, C.B. (2007). Developing theory through simulation methods. *Academy of Management Review*, 32, 480-499.
- Deutsch, J.A., & Deutsch, D. (1963). Attention: Some theoretical considerations. *Psychological Review*, 70, 80-90
- Dionne, S.D., Sayama, H., Hao, C., & Bush, B.J. (2010). The role of leadership in shared mental model convergence and team performance improvement: An agent-based computational model. *The Leadership Quarterly*, 21, 1035-1049.

- Duncan, J., & Humphreys, G.W. (1989). Visual search and stimulus similarity. *Psychological Review*, 96, 433–458.
- Epstein, J. M. (1999). Agent-based computational models and generative social science. *Complexity*, 4, 41-60.
- Fiore, S.M., Rosen, M.A., Smith-Jentsch, K.A., Salas, E., Letsky, M., & Warner, N. (2010). Toward an understanding of macrocognition in teams: Predicting processes in complex collaborative contexts. *Human Factors*, 52, 203-224.
- Fiore, S.M., Smith-Jentsch, K.A., Salas, E., Warner, N., & Letsky, M. (2010). Towards an understanding of macrocognition in teams: Developing and defining complex collaborative processes and products. *Theoretical Issues in Ergonomics Science*, 11, 250-271.
- Hall, S. (1973). *Encoding and decoding in the television discourse* (Vol. 7). Centre for Cultural Studies, University of Birmingham.
- Harrison, J. R., Lin, Z., Carroll, G. R., & Carley, K. M. (2007). Simulation modeling in organizational and management research. *Academy of Management Review*, 32, 1229-1245.
- Hinsz, V.B., Tindale, R.S., & Vollrath, D.A. (1997). The emerging conceptualization of groups as information processors. *Psychological Bulletin*, 121, 43-64.
- Hintzman, D.L. (1984). MINERVA2: A simulation model of human memory. *Behavior Research Methods, Instruments, and Computers*, 16, 96-101.
- Hollenbeck, J.R., Ilgen, D.R., Segoe, D.J., Hedlund, J., Major, D.A., & Phillips, J. (1995). Multilevel theory of team decision-making: Decision performance in teams incorporating distributed expertise. *Journal of Applied Psychology*, 80, 292-316.
- Hunt, E. (1994). Problem-solving. In R. J. Sternberg (Ed.), *Thinking and problem-solving*. (pp. 215-232). San Diego: Academic Press.
- Ilgen, D.R., Hollenbeck, J.R., Johnson, M., & Jundt, D. (2005). Teams in organizations: From input-process-output models to IMO models. *Annual Review of Psychology*, 56, 517-543.
- Ilgen, D.R., & Hulin, C.L. (2000). *Computational modeling of behavioral processes in organizations: The third scientific discipline*. Washington, DC: American Psychological Association.
- Kanki, B., Helmreich, R., & Anca, J. (2010). *Crew resource management* (2nd ed.). San Diego, CA: Academic Press.
- Katz, D., & Kahn, R.L. (1978). *The social psychology of organizations*. New York, NY: Wiley
- Kennedy, D.M., & McComb, S.A. (2014). When teams shift among processes: Insights from simulation and optimization. *Journal of Applied Psychology*, 99, 784-815.
- Klimoski, R., & Mohammed, S. (1994). Team mental model: Construct or metaphor? *Journal of management*, 20, 403-437.
- Kozlowski, S.W.J. (2015). Advancing research on team process dynamics: Theoretical, methodological, and measurement considerations. *Organizational Psychology Review*, 5, 270-299.
- Kozlowski, S.W.J. & Chao, G.T. (2012a). Macrocognition, team learning, and team knowledge: Origins, emergence, and measurement. In E. Salas, S. Fiore, & M. Letsky (Eds.), *Theories of team cognition: Cross-disciplinary perspectives* (pp. 19-48). New York: Taylor & Francis.
- Kozlowski, S.W.J., & Chao, G. T. (2012b). The dynamics of emergence: Cognition and cohesion in work teams. *Managerial and Decision Economics*, 33, 335-354.

- Kozlowski, S.W.J., Chao, G.T., Grand, J.A., Braun, M.T., & Kuljanin, G. (2016). Capturing the multilevel dynamics of emergence: Computational modeling, simulation, and virtual experimentation. *Organizational Psychology Review, 6*, 3-33.
- Kozlowski, S.W.J., Chao, G.T., Grand, J.A., Braun, M.T., & Kuljanin, G. (2013). Advancing multilevel research design: Capturing the dynamics of emergence. *Organizational Research Methods, 16*, 581-615.
- Kozlowski, S.W.J., & Ilgen, D.R. (2006). Enhancing the effectiveness of work groups and teams. *Psychological Science in the Public Interest, 7*, 77-124.
- Kozlowski, S.W.J., & Klein, K.J. (2000). A multilevel approach to theory and research in organizations: Contextual, temporal, and emergent processes. In K.J. Klein & S.W.J. Kozlowski (Eds.), *Multilevel theory, research, and methods in organizations: Foundations, extensions, and new directions* (pp. 3-90). San Francisco, CA: Jossey-Bass.
- Kozlowski, S.W.J., Toney, R.J., Mullins, M.E., Weissbein, D.A., Brown, K.G., & Bell, B.S. (2001). Developing adaptability: A theory for the design of integrated-embedded training systems. In E. Salas (Ed.), *Advances in human performance and cognitive engineering research* (Vol. 1, pp. 59-123). Amsterdam: JAI/ Elsevier Science.
- Kuljanin, G., Braun, M.T., & DeShon, R.P. (2011). A cautionary note on modeling growth trends in longitudinal data. *Psychological Methods, 16*, 249-264.
- Law, A.M., & Kelton, D.W. (1991). *Simulation modeling and analysis* (2nd ed.). New York: McGraw-Hill.
- Lipshitz, R., Levy, D.L., & Orchen, K. (2006). Is this problem to be solved? A cognitive schema of effective problem-solving. *Thinking and Reasoning, 12*, 413-430.
- Lord, R.G., Dinh, J.E., & Hoffman, E.L. (2015). A quantum approach to time and organizational change. *Academy of Management Review, 40*, 263-290.
- Love, B.C., Medin, D.L., & Gureckis, T.M. (2004). SUSTAIN: A network model of category learning. *Psychological Review, 111*, 309-332.
- Lu, L., Yuan, Y. C., & McLeod, P. L. (2012). Twenty-five years of hidden profiles in group decision making: A meta-analysis. *Personality and Social Psychology Review, 16*, 54-75.
- Macal, C.M., & North, M.J. (2010). Tutorial on agent-based modeling and simulation. *Journal of Simulation, 4*, 151-162.
- Mathieu, J.E., Heffner, T.S., Goodwin, G.F., Salas, E., & Cannon-Bowers, J.A. (2000). The influence of shared mental models on team process and performance. *Journal of applied psychology, 85*, 273-283.
- McComb, S.A. (2007). Mental model convergence: The shift from being an individual to being a team member. In F. Dansereau & F.J. Yammarino (Eds.), *Research in multilevel issues* (Vol. 6, pp. 95-145). Amsterdam: Elsevier.
- Medin, D.L., & Schaffer, M.M. (1978). Context theory of classification learning. *Psychological Review, 85*, 207-238.
- Mesmer-Magnus, J. R., & DeChurch, L. A. (2009). Information sharing and team performance: A meta-analysis. *Journal of Applied Psychology, 94*, 535-546.
- Meyer, D.E., & Kieras, D.E. (1997). A computational theory of executive cognitive processes and multiple-task performance: Part 1. Basic mechanisms. *Psychological Review, 104*, 2-65.
- Miller, J. H., & Page, S. E. (2007). *Complex adaptive systems*. Princeton, NJ: Princeton University Press.

- 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.
- Mohammed, S., Tesler, R., & Hamilton, K. (2012). Time and shared cognition: Towards greater integration of temporal dynamics. In E. Salas, S. Fiore, & M. Letsky (Eds.), *Theories of team cognition: Cross-disciplinary perspectives* (pp. 87-116). New York: Taylor and Francis Group, LLC.
- Moore, E., Laiti, L., & Chelazzi, L. (2003). Associative knowledge controls deployment of visual selective attention. *Nature neuroscience, 6*, 182-189.
- Morgeson, F.P., & Hoffman, D.A. (1999). The structure and function of collective constructs: Implications for multilevel research and theory development. *Academy of Management Review, 24*, 249-265.
- Newell, A. (1990). *Unified theories of cognition*. Cambridge, MA: Harvard University Press.
- Nosofsky, R.M. (1984). Choice, similarity, and the context theory of classification. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 10*, 104-114.
- Nosofsky, R.M. (1987). Attention and learning processes in the identification and categorization of integral stimuli. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 13*, 87-108.
- Nosofsky, R.M., & Zaki, S.R. (1998). Dissociations between categorization and recognition in amnesiac and normal individuals: An exemplar-based interpretation. *Psychological Science, 9*, 247-255.
- Parker, K.C. (1988). Speaking turns in small group interaction: A context-sensitive event sequence model. *Journal of Personality and Social Psychology, 54*, 965-971.
- Patrashkova, R., & McComb, S.A. (2004). Exploring why more communication is not better: Insights from a computational model of cross-functional teams. *Journal of Engineering and Technology Management, 21*, 83-114.
- R Core Team (2016). *R: A language and environment for statistical computing*. R Foundation for Statistical Computing. Vienna, Austria.
- Reynolds, C. (1987). Flocks, herds and schools: A distributed behavioral model. *SIGGRAPH '87: Proceedings of the 14th annual conference on Computer graphics and interactive techniques* (Association for Computing Machinery): 25-34.
- Rosch, E., Mervis, C.B., Gray, W.D., Johnson, D.M., & Boyes-Braem, P. (1976). Basic objects in natural categories. *Cognitive Psychology, 8*, 382-439.
- Russell, R.D. (1928). A comparison of two methods of learning. *The Journal of Educational Research, 18*(3), 235-238.
- Salas, E., DiazGranados, D., Klein, C., Burke, C.S., Stagl, K.C., Goodwin, G.F., & Halpin, S.M. (2008). Does team training improve team performance? A meta-analysis. *Human Factors, 50*, 903-933.
- Salas, E., Prince, C., Baker, D.P., & Shrestha, L. (1995). Situation awareness in team performance: Implications for measurement and training. *Human Factors, 37*, 123-136.
- Salas, E., Sims, D.E., & Burke, C.S. (2005). Is there a "Big Five" in teamwork? *Small Group Research, 36*, 555-599.
- Spencer, E.M. (1941). The retention of orally presented materials. *Journal of Educational Psychology, 32*(9), 641-655.
- Soto, D., Heinke, D., Humphreys, G.W., & Blanco, M.J. (2005). Early, involuntary top-down guidance of attention from working memory. *Journal of Experimental Psychology: Human Perception and Performance, 31*, 248-261.

- Stasser, G. (1988). Computer simulation as a research tool: The DISCUSS model of group decision making. *Journal of Experimental Social Psychology, 24*, 393-422.
- 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.
- Stasser, G., & Titus, W. (1987). Effects of information load and percentage of shared information on the dissemination of unshared information during group discussion. *Journal of Personality and Social Psychology, 53*, 81-93.
- Steiner, I. D. (1972). *Group process and productivity*. New York: Academic Press.
- Stephan, F.F., & Mischler, E.G. (1952). The distribution of participation in small groups: An exponential approximation. *American Sociological Review, 17*, 598-608.
- Sundstrom, E., De Meuse, K.P., Futrell, D. (1990). Work teams: Applications and effectiveness. *American Psychologist, 45*, 120-133.
- Taber, C.S., & Timpone, R.J. (1996). *Computational modeling*. Thousand Oaks, CA: Sage.
- Treisman, A. (1964). Selective attention in man. *British Medical Bulletin, 20*, 12-16.
- Wegner, D.M. (1987). Transactive memory: A contemporary analysis of the group mind. In B. Mullen & G.R. Goethals, *Theories of group behavior* (pp. 185-208). New York, NY: Springer-Verlag.
- Wickens, C.D. (1992). *Engineering psychology and human performance* (2nd ed.). New York: Harper Collins.
- Wildman, J., Thayer, A.L., Pavals, D., Salas, E., Stewart, J. E., & Howse, W. R. (2012). Team knowledge research: Emerging trends and critical needs. *Human Factors, 54*, 84-111.
- Wilensky, U., & Rand, W. (2015). *An introduction to agent-based modeling: Modeling natural, social, and engineered complex systems with NetLogo*. Cambridge, MA: MIT Press.
- Wittenbaum, G.M., Hollingshead, A.B., & Botero, I.C. (2004). From cooperative to motivated information sharing in groups: Moving beyond the hidden profile paradigm. *Communication Monographs, 71*, 286-310.
- Vancouver, J. B., Tamanini, K. B., & Yoder, R. J. (2010). Using dynamic computational models to reconnect theory and research: Socialization by the proactive newcomer as example. *Journal of Management, 36*, 764-793.
- Vancouver, J.B., & Weinhardt, J.M. (2012). Modeling the mind and the milieu: Computational modeling for micro-level organizational researchers. *Organizational Research Methods, 15*, 602-623.
- von Bertalanffy, L. (1950). The theory of open systems in physics and biology. *Science, 111*, 23-29.

Footnotes

¹ This specification is equivalent to assuming that all actions performed by agents take the same amount of time as measured in seconds, minutes, etc.

² The full code for running the agent-based simulation is included in the supplementary material for this article.

³ Note that the information processing skill of each agent in each team was determined by randomly sampling an integer from an interval consistent with the information processing skill conditions. All agents on a given team thus tended to have *similar* but not identical encoding/decoding rates, and thus it was possible to manipulate the relative speaking rates of agents within each team in accordance with this manipulation. For example, a team in the low information processing skill condition might be generated such that agent one's encoding rate = 9, agent two's encoding rate = 10, and agent three's encoding rate = 10. If this team were in the communication skill condition in which the lowest skilled agent spoke most frequently, agent one would be given the highest probability of speaking (50%) followed by agents two and three (both 25%). If this team were in the communication skill condition in which the highest skilled agent spoke most frequently, either agent two or agent three would be selected to have the highest probability of speaking.

⁴ A primary analytic goal of this experiment was to evaluate the internal coherence and logical consistency of the ABS instantiating our theory of team knowledge emergence. Doing so involves assessing the degree to which manipulating theoretically meaningful parameters produces changes in model outcomes that are interpretable and consistent with conceptual rationale (Taber & Timpone, 1996; Epstein, 1999). As such, although the statistical models included all possible two- and three-way interaction effects, only the main effects for the experimental manipulations are evaluated and reported. These results provide the most straightforward interpretation of how changes in individual-level learning and sharing behaviors contribute to team knowledge emergence.

⁵ This result was likely an artifact of the internalization process and stopping protocol used in the ABS. Since all teams were comprised of members with fairly homogenous encoding rates (e.g., agents on high ability teams required 3, 4, or 5 actions to encode data), there was relatively little variability among knowledge acquisition trajectories within teams. Nevertheless, because the simulations ran until all agents internalized the entire knowledge pool and higher ability agents could acquire more knowledge during each phase, idiosyncratic patterns of communication during sharing phases meant that high ability agents reached maximum internalization much more quickly than their teammates in some cases. Because an agent's knowledge acquisition trajectory stopped growing

once it had learned all possible knowledge, this had the potential to result in slightly larger variability in internalization rates in higher compared to lower ability teams.

⁶ It should be noted that modeling *changes* in the information processing and communication skills of agents over time was not examined in the virtual experiments nor is this mechanism incorporated into our theory of team knowledge emergence. If such expertise development were of conceptual interest, the computational model summarized in Table A1 could be adapted to include feedback mechanisms that change specific parameters or mechanisms (e.g., encoding, decoding, integration, speaking, etc.) in response to exogenous or endogenous factors (e.g., experience, training, accumulated task knowledge, etc.). However, identifying the factors that influence learning and sharing processes or how those process mechanisms change over time is beyond the scope of our theory and does not directly contribute to advancing a general process-oriented account of team knowledge emergence.

⁷ Degree of specialization within the team's task environment was not manipulated in the experiment with human teams. As described previously, this structural characteristic was not expected to directly impact *how* members carried out the core learning and sharing process mechanisms of team knowledge emergence. Although the simulation results demonstrate that this factor has an influence on team knowledge outcomes, the effect occurs because it requires team members to rely more heavily on one another for information and therefore places greater emphasis on the relatively more demanding knowledge sharing processes. The actual learning and sharing processes enacted by team members did not differ under differing degrees of specialization and thus was held constant in the empirical experiments.

⁸ This mechanism is intended to emulate the process of information exchange that occurs in a face-to-face context. To learn information communicated that is communicated by another team member, the individual receiving information must direct their attention to the message that is being conveyed to them. Failure to do so results in the information sharing attempt going unnoticed, and it must be communicated again to be acknowledged.

⁹ In addition to the CG for learning and sharing, CG describing how to make accurate route and asset decisions were also administered to some teams in both the control and experimental conditions. These prompts only provided guidance on how to combine acquired information to make task decisions. Subsequent analyses (available from the first author upon request) show that the decision-focused CG had no significant influence on learning/sharing behaviors or outcomes (i.e., amount of information internalized, externalized, etc.). Consequently, we do not discuss the role or design of the embedded decision CG further.

¹⁰ The purpose of the fidelity comparison was to examine whether the *patterns* of knowledge emergence produced by an agent team following the process rules in our simulation were consistent with the *patterns* of knowledge emergence observed in human teams. It was only necessary to simulate a single agent team to accomplish this goal; because all agents enacted the same learning and sharing processes (i.e., Table A1 in Appendix A), the same global trends with respect to how knowledge accumulates over time will always be the same regardless of the model parameters selected. Although there will be idiosyncrasies in the specific trajectories of knowledge acquisition for any one simulated agent/team under a given parameterization, such differences are not substantively meaningful. Consequently, it was unnecessary to simulate and then average across many teams under the same parameterization for purposes of the fidelity comparisons.

¹¹ Because the units of time differed between the simulated (number of learning and sharing phases) and participant (elapsed seconds) teams, the time scale for the simulated data was transformed to facilitate comparisons with the observed data. This conversion was performed by dividing the total amount of time participant teams were given to engage in knowledge-building during a trial (360 seconds) by the total number of phases observed for the selected simulated team to fully externalize the entire knowledge pool (40 phases). This resulted in equating all learning and sharing phases for the simulated team as lasting approximately 9 real-time seconds in Figures 3 and 4. This conversion is not intended to imply a “true” duration for the learning and sharing phases represented in our ABS. The time scale transformation simply provides a convenient method for comparing the overall patterns of team knowledge emergence implied by our theory and reflected in the simulated and participant team data.

¹² Although the simulated team appears to take much longer to achieve complete team knowledge coverage compared to the participant teams, this delay is an artifact of the time scale transformation described in Footnote 10. Just like the real human teams, simulated teams were initially acquiring knowledge primarily during learning phases of activity. However, the ABS forced agents to perform at least one sharing phase after completing a learning phase, even if there was no new knowledge to share (see Table A1, Appendix A). Because *any* and *all* simulated phases—even a sharing phase with just one iteration and no knowledge exchange—were treated as 9 seconds in length in Figures 3 and 4, this had the effect of artificially inflating the apparent number of seconds before simulated agents reached complete team knowledge coverage. Although this does not affect interpretations of the overall patterns/trajectories of team knowledge emergence in the simulated data, it cautions against drawing conclusions about the implied time scale of knowledge emergence expressed in the ABS.

Table 1
Coefficient estimates from MRCM analyses for individual internalization rates by agents (Step 3)

Variables	<i>b</i>	Std. Error	95% CI
<i>DV: Number of time steps required by individual agents to internalize all knowledge</i>			
Intercept	32.47**	.11	[32.25, 32.69]
Low information processing skill	7.98**	.15	[7.69, 8.27]
Moderate information processing skill	16.38**	.15	[16.09, 16.67]
Equal speaking rate	-8.29**	.18	[-8.64, -7.94]
Most skilled agent speaks most	.48**	.05	[.38, .57]
50% unique information	3.83**	.15	[3.54, 4.12]
66% unique information	16.42**	.15	[16.13, 16.71]

** $p < .001$

Note. Parameter estimates are from the full two-level MRCM model including all two- and three-way interactions (interaction effects not reported, see Appendix B). All predictor variables were dummy coded Level-2 categorical variables:

Low information processing skill (0 = high skilled team, 1 = low skilled team)

Moderate information processing skill (0 = high skilled team, 1 = moderate skilled team)

Equal speaking rate (0 = lowest skilled agent speaks most, 1 = all agents speak equally)

Most skilled agent speaks most (0 = lowest skilled agent speaks most, 1 = most skilled agent speaks most)

50% unique information (0 = 33% unique information, 1 = 50% unique information)

66% unique information (0 = 33% unique information, 1 = 66% unique information)

Table 2.

Summary of estimated marginal means, ANOVA tests, and effect sizes for information processing skills, communication skills, and degree of specialization on simulated team knowledge outcomes (Step 3)

Independent Variable	Team Internalization Variability	Team Knowledge Coverage	Internalization Distribution	Externalization Distribution
Information processing skills				
Low skill	0.000158 _a	15.843 _a	47.141 _a	47.525 _a
Moderate skill	0.000189 _b	13.254 _b	41.625 _b	41.748 _b
High skill	0.000255 _c	11.036 _c	36.539 _c	36.670 _c
Sig. level for ANOVA	$p < .001$	$p < .001$	$p < .001$	$p < .001$
Effect size (partial η^2)	.072	.654	.887	.885
Communication skills				
Least skilled speaks most	0.000188 _d	13.374 _d	45.180 _d	45.320 _d
All speak equally	0.000234 _f	13.378 _d	34.711 _e	34.824 _e
Most skilled speaks most	0.000179 _d	13.381 _d	45.414 _f	45.525 _f
Sig. level for ANOVA	$p < .001$	$p = .994$	$p < .001$	$p < .001$
Effect size (partial η^2)	.026	.000	.912	.991
Degree of Specialization				
33% unique/67% common	0.000121 _g	9.699 _g	37.481 _g	37.615 _g
50% unique/50% common	0.000213 _h	12.900 _h	38.926 _h	39.051 _h
66% unique/33% common	0.000267 _i	17.533 _i	48.898 _i	49.004 _i
Sig. level for ANOVA	$p < .001$	$p < .001$	$p < .001$	$p < .001$
Effect size (partial η^2)	.143	.832	.914	.913

Note. Within a column, means with identical subscripts do not significantly differ (Tukey HSD test $p > 0.05$).

Table 3
Operationalization of core concepts & mechanisms from theory of team knowledge emergence in virtual (Step 3) and empirical (Step 4) experiments

Core Concepts & Mechanisms	Description	Representation in ABS (Step 3)	Representation in CRONUS (Step 4)
Data Selection	Identifying data to be learned from the task environment	Agents select an accessible piece of data to learn from knowledge pool or attend to data shared by another agent	Members search Specialist Map or Mission Map for obstacles to learn
Encoding	Transforming data observed from the environment into internalized data ("learning-from-self")	Agents spend some number of actions to internalize a selected piece of data	Members post an obstacle identified from Specialist Map to the Specialist Map
Decoding	Transforming knowledge received from other team members into internalized knowledge ("learning-from-others")	Agents spend some number of actions to internalize a piece of communicated knowledge	Members post an obstacle identified from Mission Map to Specialist Map
Integration	Transforming internalized data into relationally organized internalized knowledge	Agents spend some number of actions to learn an association between two pieces of internalized data	(Not examined in current study)
Member Selection	Choosing to speak to other members in the team	Agents probabilistically selected to speak	Members visit Mission Map to communicate information
Retrieval	Identifying internalized knowledge from memory to be shared	Agents choose one piece of non-externalized knowledge from knowledge they have internalized	Members choose an obstacle posted on Specialist Map to share with team members
Sharing	Communicating of internalized knowledge to other team members	Agents communicate a piece of knowledge to all other agents on the team	Members post an obstacle to Mission Map
Acknowledgement	Generating externalized knowledge by confirming knowledge shared by another team member is internalized	Agents repeat a piece of communicated knowledge once it has been internalized	Members post a shared obstacle to Specialist Map to make obstacle permanently visible on Mission Map

Note. ABS = agent-based simulation; CRONUS = Crisis Relief Operation: Naval Unit Simulation.

Table 4.

List, description, and rationale for the embedded contextualized guidance (CG) prompts provided to experimental condition teams in Step 4

Focal Target	CG Name	Description	Rationale
Information Processing Skills	Avoid Distractors	Triggered when a member posted non-mission relevant obstacle to the SM. Advised to delete the post and not post this obstacle type in the future.	Improve data selection and encoding by focusing members only on task-relevant information
	Repeat Common Obstacle	Triggered when a member posted a common obstacle to the MM. Advised to delete the post and not post this obstacle type to the MM in the future.	Improve data selection and decoding by reducing ambiguity over presence of information
	Incomplete Specialist Map	Triggered when a member posted an obstacle to the MM before learning all available obstacles on the SM. Advised to search SM further before posting additional obstacles to MM.	Improve encoding and decoding by focusing members on the quicker learning-from-self process first and ensuring no unique information is forgotten
	Incorrect Post	Triggered when a member posted an obstacle to the SM that did not exist. Advised to delete the post and confirm identity of the obstacle being posted.	Improve data selection and encoding by removing incorrect information
Communication Skills	Share Too Quickly	Triggered when a member posted an obstacle to the MM before all other members had learned every obstacle on their SMs. Advised to use the in-game chat box to notify teammates when ready to share and receive information.	Improve member selection and sharing by reducing number of communication attempts missed by teammates
	Accidental Mission Map Post	Triggered when a member posted a unique obstacle not within his/her domain of expertise to the MM that was likely meant to be posted to the SM instead. Advised to undo the post from the Mission Map and post to the SM.	Improve decoding and acknowledgement by reducing number of incorrect communication attempts and improving learning-from-others
	Missed Post: Receiver	Triggered when a member missed an obstacle post shared by a team member on the MM. Advised to monitor communication notification that indicate obstacle information has been shared by a team member.	Improve data selection and sharing by reducing number of communication attempts missed by teammates
	Missed Post: Sender	Triggered when a post shared by the team member on the MM was missed by another team member. Advised to re-post the missed obstacle on MM.	Improve sharing by reducing number of communication attempts missed by teammates
	Sharing Strategy	Triggered when a member posted an obstacle to the MM out of order. Advised all members to follow a prescribed strategy for posting obstacles to the MM.	Improve member selection and sharing by developing a common strategy for communicating information
	Learning Strategy	Triggered when a member posted new obstacles to the MM before posting obstacles shared with him/her on the MM to the SM. Advised all members to follow a prescribed strategy for transferring obstacles from the MM to the SM.	Improve decoding and sharing by developing a common strategy for acquiring shared information

Note. SM = Specialist Map; MM = Mission Map. All CG prompts were delivered through a pop-up window that appeared on screen of the member(s) who triggered the CG. Members were required to click an "OK" button to remove the window and continue with the task.


Table 5
Coefficient estimates from MRCM analyses for amount of fully externalized knowledge acquired by teams across trials (Step 4)

Variables	<i>b</i>	Std. Error	95% CI
<i>DV: Number of fully externalized obstacles at end of trial</i>			
Intercept	1.93**	.23	[1.48, 2.38]
Trial	1.10**	.10	[.90, 1.29]
Trial ²	-.07**	.01	[-.08, -.05]
Condition	.34	.33	[-.31, .99]
Trial x Condition	-.32*	.14	[-.60, -.03]
Trial ² x Condition	.05**	.01	[.02, .07]

* $p < .05$, ** $p < .001$


Note. Condition is a dummy coded variable (0 = control; experimental = 1). Trial variable was coded such that 0 = first trial. All coefficients are reported in original (unstandardized) units.

Figure 1. General framework for developing and evaluating theories of emergence.




Note. CM = computational model; ABS = agent-based simulation. Descriptions in the light gray boxes summarize the tasks carried out for each step in the present research.

Figure 2. Process-oriented theory of knowledge emergence in teams.




Note. Solid arrows signify pathways engaged during learning processes, while dotted arrows signify pathways engaged during sharing processes. The boxes labeled *Information Processing* and *Communication* highlight the core process mechanisms targeted by the manipulations in the virtual (Step 3) and empirical (Step 4) experiments. The diagram labeled *Emergent Team Knowledge* exemplifies how team knowledge outcomes were operationalized in the simulated and empirical data as the distribution of internalized and externalized knowledge across members in a team. Members in this diagram are represented by the letters A, B, C. Each “wedge” of the knowledge pool represents a distinct piece of information that could be learned. Letters within a wedge indicate members who have internalized a piece of information. Letters separated by an arc within each wedge indicate members who do not yet share externalized knowledge about that piece of information; letters not separated by an arc within each wedge indicate members who share externalized knowledge about a piece of information. For example, the wedge labeled *Internalized: Full Overlap* shows that all three members of the team have internalized this same piece of information, but no one has acknowledged this information with other team members (i.e., no externalization). In contrast, the wedge labeled *Externalized: Partial Overlap* shows that all three members have internalized this information, but members B and C have also acknowledged and externalized this information.

Figure 3. Comparison of individual knowledge internalization for a single trial between (A) a simulated team and a representative experimental condition team and (B) a representative control condition team (Step 4).




Note. The same simulated team is shown in both panels. Agents in the simulated team had heterogeneous information processing skills, equal speaking rates, and operated in an environment with high specialization (75% unique information). Data from the control and experimental teams were from Trial 10.

Figure 4. Comparison of changes in internalization distribution for (A) non-overlapping, (B) partially overlapping, and (C) fully overlapping knowledge for simulated, experimental, and control condition teams for a single trial (Step 4).




Note. Agents in the simulated team had heterogeneous information processing skills, equal speaking rates, and operated in an environment with high specialization (75% unique information). Data from the control and experimental teams were from Trial 10.

Figure 5. Observed and regression estimated number of fully externalized obstacles by teams each trial in the control and experimental conditions (Step 4).



Note. * $p < .05$ for difference in observed means between experimental and control condition. CG = contextualized guidance. Error bars show 95% confidence interval of observed means. The CG was present only in the experimental condition and was introduced at start of Trial 3 and removed at start of Trial 10.

Figure 6. Integrative framework for advancing research on team cognition and knowledge emergence.



APPENDIX A

Table A1 outlines the procedural algorithm of the computational model used to formally specify our process-oriented theory of knowledge emergence in teams. These rules were also enacted by agents in the ABSs used during the virtual experiment conducted in Step 3 of the paper.

Table A2 summarizes the list of assumptions and specifications instantiated in the ABSs reported in Step 3 of the paper. When translating a theory into a computational representation and operationalizing it into a dynamic simulation, researchers must make a number of assumptions about the nature, sequence, and boundary conditions implied by the proposed theoretical processes (cf., Epstein, 1999). Some assumptions are more operational in nature and are dictated by the demands of programming a dynamic process into computer language (i.e., How/when does the amount of knowledge an agent possesses increment? How is “memory” quantitatively represented?). However, many assumptions are conceptual in nature and may carry implications for the nature of the theoretical process at hand (i.e., Do agents act simultaneously or sequentially? Do agent parameters change or remain constant throughout a simulation?). Regardless of their nature, making such assumptions transparent is critical to understanding and interpreting a computational model and accompanying simulations.

Table A1.

Computational model of team knowledge emergence

Learning Process
<ol style="list-style-type: none"> 1. Set timer for Learning phase to 0 2. Select piece of data to learn from knowledge pool that is accessible <ol style="list-style-type: none"> a. If no accessible data partially or completely encoded, randomly select new data to encode, go to Step 3 b. If any accessible data partially encoded, select that data, go to Step 3 c. If any accessible data completely encoded, go to Step 4 3. Encode selected data at encoding rate <ol style="list-style-type: none"> a. If data encoding not completed, go to Step 7 b. If data encoding completed, go to Step 5 4. Identify whether any internalized data can be integrated <ol style="list-style-type: none"> a. If no integration can occur, go to Step 2a or 2b b. If integration can occur, encode link between data pieces, go to Step 6 5. Increment amount of internalized data by 1, go to Step 7 6. Increment amount of internalized knowledge by 1, go to Step 7 7. Increment timer for learning phase by 1 <ol style="list-style-type: none"> a. If learning phase timer < length of learning phase, return to Step 2 b. If learning phase timer = length of learning phase, go to Step 8
Sharing Process
<ol style="list-style-type: none"> 8. Set timer for Sharing phase to 0 9. Select one agent to speak according to speaking rate 10. Determine whether selected agent has any internalized knowledge to share that is not yet fully externalized <ol style="list-style-type: none"> a. If no new knowledge to share, go to Step 15 b. If new knowledge to share, go to Step 11 11. Speaking agent retrieves random piece of internalized knowledge and shares with other agents 12. Receiving agents decode shared knowledge at decoding rate <ol style="list-style-type: none"> a. If shared knowledge not fully decoded by receiving agent, agent goes to Step 15 b. If shared knowledge fully decoded by receiving agent, agent increments internalized knowledge by 1, goes to Step 13 13. Receiving agent(s) acknowledge newly internalized knowledge by repeating it to other agents <ol style="list-style-type: none"> a. If any other agents have not yet internalized this knowledge, those agents repeat Step 12 b. If knowledge is acknowledged by all agents, go to Step 14 14. Increment amount of externalized knowledge by 1 15. Increment timer for sharing phase by 1 <ol style="list-style-type: none"> a. If sharing phase timer < length of sharing phase, return to Step 9 b. If sharing phase timer = length of sharing phase or no new information to share, go to Step 16
Team Knowledge-building Cycle
<ol style="list-style-type: none"> 16. Determine whether all agents have internalized entire knowledge pool <ol style="list-style-type: none"> a. If entire knowledge pool not internalized, go to Step 1 b. If entire knowledge pool internalized, go to Step 17 17. Determine whether all agents have externalized entire knowledge pool <ol style="list-style-type: none"> a. If entire knowledge pool not externalized, go to Step 8 b. If entire knowledge pool externalized, end simulation

Note. The length of both the learning and sharing phase timers was set to 50 time steps. All steps of the pseudo-code were evaluated individually by each agent, but all agents proceeded through each phase synchronously (i.e., the next time increment within a phase did not begin until all agents reached the final step of the phase).

Table A2.

Specifications and assumptions for agent-based simulation of team knowledge emergence

-
1. Time is equated with agent actions such that each agent only performs one action per time step. Consequently, all possible actions an agent can take during a single time step (e.g., encoding attempt, decoding attempt, sharing attempt) take the same amount of time.
 2. All teams initially begin in a learning phase in which agents are assumed to be complete novices who know none of the information in the knowledge pool.
 3. Agents fully encode a single piece of data during learning before selecting another piece of data to learn. Thus, agents “remember” what they have started encoding and always opt to finish learning partially encoded data.
 4. Encoding and decoding is operationalized as the number of time steps needed for an agent to fully internalize a piece of data.
 5. Decoding rates are slower than encoding rates.
 6. Creating internalized knowledge requires agents to integrate two pieces of internalized data that share an a priori association. Once agents internalize both pieces of data, they integrate the data into internalized knowledge.
 7. Integration rates are the same for all agents. All agents require only one time step to completely integrate two pieces of internalized data into a piece of internalized knowledge.
 8. Member selection (i.e., speaking rate) is operationalized as the probability a given member will be selected to speak in the team. Member selection at each time step is an independent random sampling process and is not influenced by any previous events (e.g., speaking in a previous turn, learning new information, etc.).
 9. Each agent’s speaking, encoding, decoding, and integration rates remain constant throughout an entire simulation. Agents did not improve (or become worse) at performing certain actions over time.
 10. When an agent shares a piece of knowledge, it is simultaneously communicated to and received by all other agents.
 11. Agents know which pieces of knowledge have been externalized/acknowledged by other agents, but they do not know which pieces of knowledge have been internalized by other agents.
 12. Agents only share internalized, integrated knowledge that has not yet been fully externalized by all other agents.
 13. Agents always acknowledge a piece of internalized knowledge after it is fully decoded.
 14. When a piece of knowledge is repeated during acknowledgement, it is considered a normal sharing attempt by any agents that have not yet internalized that knowledge and prompts an additional decoding attempt for those agents.
 15. If all agents have internalized but not yet externalized the entire knowledge pool, agents no longer enter learning phases and will continually cycle through sharing phases until all knowledge is fully externalized by all agents.
 16. Agents act with perfect memory and without bias; that is, agents never forget something they have internalized.
-

APPENDIX B

MRCM Analyses for Step 3

The purpose of the MRCM analyses in Step 3 was to evaluate the overall main effects of agent ability level, agent speaking rate, and the distribution of information in the task environment on knowledge internalization rates observed at the individual-/agent-level. A 2-level (Level-1 = individual, Level 2 = team) MRCM was thus used to account for the fact that individual agents were nested within teams:

$$\text{Level 1: } \text{Ind.Int}_{ij} = \beta_{0j} + r_{ij} \quad (1)$$

$$\begin{aligned} \text{Level 2: } \beta_{0j} = & \gamma_{00} + \gamma_{01}(\text{Low.Ability}_j) + \gamma_{02}(\text{Mod.Ability}_j) + \gamma_{03}(\text{Equal.Spkrate}_j) + \gamma_{04}(\text{High.Spkrate}_j) + \\ & \gamma_{05}(\text{Mod.InfoDist}_j) + \gamma_{06}(\text{High.InfoDist}_j) + \gamma_{07}(\text{Low.Ability}_j * \text{Equal.Spkrate}_j) + \\ & \gamma_{08}(\text{Mod.Ability}_j * \text{Equal.Spkrate}_j) + \gamma_{09}(\text{Low.Ability}_j * \text{High.Spkrate}_j) + \gamma_{10}(\text{Mod.Ability}_j * \text{High.Spkrate}_j) \\ & + \\ & \gamma_{11}(\text{Low.Ability}_j * \text{Mod.InfoDist}_j) + \gamma_{12}(\text{Mod.Ability}_j * \text{Mod.InfoDist}_j) + \\ & \gamma_{13}(\text{Low.Ability}_j * \text{High.InfoDist}_j) + \gamma_{14}(\text{Mod.Ability}_j * \text{High.InfoDist}_j) + \\ & \gamma_{15}(\text{Equal.Spkrate}_j * \text{Mod.InfoDist}_j) + \gamma_{16}(\text{Equal.Spkrate}_j * \text{High.InfoDist}_j) + \\ & \gamma_{17}(\text{High.Spkrate}_j * \text{Mod.InfoDist}_j) + \gamma_{18}(\text{High.Spkrate}_j * \text{High.InfoDist}_j) + \\ & \gamma_{19}(\text{Low.Ability}_j * \text{Equal.Spkrate}_j * \text{Mod.InfoDist}_j) + \\ & \gamma_{20}(\text{Mod.Ability}_j * \text{Equal.Spkrate}_j * \text{Mod.InfoDist}_j) + \gamma_{21}(\text{Low.Ability}_j * \text{High.Spkrate}_j * \text{Mod.InfoDist}_j) + \\ & \gamma_{22}(\text{Mod.Ability}_j * \text{High.Spkrate}_j * \text{Mod.InfoDist}_j) + \gamma_{23}(\text{Low.Ability}_j * \text{Equal.Spkrate}_j * \text{High.InfoDist}_j) + \\ & \gamma_{24}(\text{Mod.Ability}_j * \text{Equal.Spkrate}_j * \text{High.InfoDist}_j) + \gamma_{25}(\text{Low.Ability}_j * \text{High.Spkrate}_j * \text{High.InfoDist}_j) + \\ & \gamma_{26}(\text{Mod.Ability}_j * \text{High.Spkrate}_j * \text{High.InfoDist}_j) + u_{0j}, \end{aligned}$$

where Ind.Int_{ij} indicates the individual internalization rate for agent i on team j , Low.Ability_j and Mod.Ability_j are dummy-coded condition variables indicating teams of all low-ability agents or moderate-ability agents (respectively), Equal.Spkrate_j and High.Spkrate_j are dummy-coded condition variables indicating teams in which the speaking rates for agents were equal or biased towards the highest ability agent (respectively), and Mod.InfoDist_j and High.InfoDist_j are dummy-coded condition variables indicating teams in which 50% or 66% of the knowledge pool was composed of unique information (respectively).

It should be noted that although the fully-crossed factorial design employed in the Step 3 simulations permitted analysis of all two- and three-way interactions among the ability level, speaking rate, and information distribution manipulations from our results, only the main effect findings for these variables were relevant to the conceptual argument advanced in our paper. Accordingly, the MRCM (Table 1) and ANOVA (Table 2) results discussed in the Results section of the paper do not present or describe any interaction results from these analyses. The rationale behind this decision was twofold. First, the primary purpose of Step 3 was to evaluate the generative sufficiency of our computational model; that is, the goal of virtual experimentation was to establish that changes in key exogenous variables in our model (i.e., encoding/decoding rates, speaking rates, information distribution) were capable of producing coherent and tractable patterns of change in the key outcome variables of interest. Establishing the validity of these core process mechanisms is a necessary precondition to exploring more nuanced and subtle patterns of interactions among exogenous variables (cf., Taber & Timpone, 1996). Thus, a thorough and systematic review of the main effect findings was of greatest significance.

Second, a related goal of Step 3 was to provide insight into points of leverage that could be manipulated in real teams to enhance team knowledge emergence. In this respect, interpretation of the interaction effects was less useful for identifying interventions that could be empirically tested in human teams. The design of our contextualized guidance interventions in Step 4 targeted behavioral analogues of the processes specified in our computational theory and carried out by agents in our ABS (see Table 4 in paper). Consequently, our interest in Step 4 was to implement interventions based on the main effect findings from Step 3 to evaluate whether teams that improved the core mechanisms specified in our theory resulted in better team knowledge outcomes. Notably, the objective was not to evaluate how specific processes/interventions affect knowledge emergence in conjunction with or in comparison to one another. The goal was to assess whether those interventions were effective at making human teams more like the “best” agent teams observed in the ABSs.

MRCM Analyses for Step 4

The purpose of the MRCM analyses in Step 4 was to evaluate changes in knowledge externalization across trials for human participant teams. The rationale for these analyses is that if teams in the experimental condition improved their learning and sharing processes as a result of the CG intervention, this should manifest as

improvements in the number of obstacles fully externalized by teams each trial. A 2-level (Level-1 = time, Level 2 = team) MRCM was thus used to account for the fact that trial was nested within teams:

$$\text{Level 1: Ext.Know}_{ij} = \beta_{0j} + \beta_{1j}(\text{Trial}_{ij}) + \beta_{2j}(\text{Trial}_{ij}^2) + r_{ij} \quad (2)$$

$$\text{Level 2: } \beta_{0j} = \gamma_{00} + \gamma_{01}(\text{Condition}_j) + u_{0j}$$

$$\beta_{1j} = \gamma_{10} + \gamma_{11}(\text{Condition}_j) + u_{1j}$$

$$\beta_{2j} = \gamma_{20} + \gamma_{21}(\text{Condition}_j) + u_{2j} ,$$

where Ext.Know_{ij} indicates the number of obstacles fully externalized on trial t by team j , Trial_{ij} reflects the trial number ($\text{Trial}_1 = 0$), Trial_{ij}^2 reflects the quadratic trial number, and Condition_j is a dummy-coded condition variable indicating whether a team was in the control ($\text{Condition} = 0$) or experimental ($\text{Condition} = 1$) condition.