

ABSTRACT

Title of Thesis: SIFTING THROUGH DECISIONS: A
COMPUTATIONAL MODEL OF
INFORMATION SHARING

Alisha C. Tronetti, Master of Science, 2025

Thesis Directed By: Associate Professor James Grand
Department of Psychology

This thesis explores a novel theory of how individuals decide to share information with their team, based on the characteristics of information (certainty, relevance), the team environment (shared knowledge, voice) and individual attributes (competence and warmth). Drawing from hidden profile research, the present work advances our understanding of team dynamics through an agent-based computational model. The SIFT model explores the emergence of knowledge within teams over time, which observes how teams cultivate, grow, and share information necessary to make optimal team decisions toward a task. Findings revealed that information distribution was the strongest predictor of unique information sharing: teams with shared access communicated more fully, while asymmetrical structures limited what was surfaced. Voice dynamics also played a central role in participation. Agents who spoke early, especially in low-voice environments—experienced positive feedback that increased their voice and led to disproportionate speaking opportunities over time. Conversely, others received fewer chances to contribute. While both certainty and voice influenced deliberation time, each had a

distinct and sustained effect: low-certainty and low-voice conditions imposed greater cognitive demand throughout. Together, these findings suggest that knowledge emergence is shaped less by individual dispositions than by structural and social feedback loops. The model offers both theoretical and practical insights into when and why team members share, highlighting the importance of early voice development and information symmetry in promoting effective team decisions.

**SIFTing Through Decisions:
A Computational Model of Information Sharing**

by

Alisha Carol Tronetti

Thesis submitted to the Faculty of the Graduate School of the
University of Maryland, College Park, in partial fulfillment
of the requirements for the degree of
Master of Science
2025

Advisory Committee:
Associate Professor James Grand, Chair
Professor Paul Hanges
Professor Rellie Derfler-Rozin

© Copyright by
Alisha C. Tronetti
2025

Table of Contents

Table of Contents.....	ii
Introduction and Overview	1
Chapter 1: Team Knowledge Emergence and Information Sharing.....	5
Dynamic Theory of Knowledge Emergence (Grand et al., 2016)	8
DISCUSS Model of Group Information Sharing (Stasser, 1988).....	12
Chapter 2: Advancing Theories of Team Knowledge Emergence	17
Chapter 3: Utility of Information.....	20
Task Relevance	20
Certainty.....	21
Expected Utility of Information.....	23
Chapter 4: Utility of Sharing Information	27
Concordance of Information	27
Voice.....	29
Self-Image.....	33
Chapter 5: Sharing Information For Teams (SIFT) Model.....	37
Model Initialization.....	38
Model Process/Algorithm	47
Chapter 6: Proposed Questions	68
Chapter 7: Methods.....	71
Simulation Parameters	71
Simulation Design.....	76
Simulation Measures.....	77
Chapter 9: Results.....	79
Surfacing Information: What Teams Choose to Share	79
Shaping the Conversation: Who Takes the Floor	87
Chapter 10: Discussion	101
Conclusion	108
References.....	110

Introduction and Overview

Teams are integral aspects of modern organizations, allowing for specialized attention on diverse tasks and efficient flow of operations (Mathieu et al., 2008, 2017, 2019). At minimum, teams are composed of two or more individuals united by shared goals and formed to complete organizationally relevant tasks (Kozlowski & Ilgen, 2006). A hallmark of team contexts is interdependence among member actions, outputs, and resources that require individuals to coordinate their knowledge, skills, and effort to complete team-relevant tasks (Courtright et al., 2015; Steiner, 1972). Furthermore, effective team functioning typically necessitates that members gather and disseminate information to facilitate task work and develop common mental representations of goals, members, and resources. This phenomenon broadly describes *team cognition*, or the “knowledge-building processes and/or the emergent mental representations characterizing the degree of team-relevant knowledge content and structure” (Mohammed et al., 2021, p. 458).

Although a large body of literature exists demonstrating the correlations and consequences of team cognition (DeChurch & Mesmer-Magnus, 2010; Mesmer-Magnus & DeChurch, 2009), comparatively less attention has been directed towards understanding how team cognition emerges, develops, and is maintained. For collective knowledge structures to emerge, a variety of specific processes must unfold within and between individuals. Broadly speaking, the core processes involved in the emergence of team cognition concern individuals acquiring (i.e., learning) and distributing (i.e., sharing) information in team contexts (Grand et al., 2016). Through

learning and sharing information alongside and with their teammates, members develop an understanding of their team, tasks, and resources relevant to shared goal accomplishment (Argote et al., 2001). As these processes play out over time and across members, they can give rise to different knowledge structures (i.e., mental representations reflecting associations among information relevant to a team and its tasks; Johnson-Laird, 2004; Chi et al., 1988) and knowledge distributions within a team (Mohammed et al., 2000). For example, some learning and sharing processes can lead to the emergence of knowledge structures that are highly variable across members and/or reflective of differences in team members' expertise and roles (i.e., transactive memory systems; Wegner, 1985) whereas others may lead to more convergent and commonly held representations (i.e., shared mental models; Mohammed et al., 2001).

Although both learning and sharing information are critical to the emergence of team cognition, the demand for interpersonal interaction and exchange typically required of teams to effectively coordinate places significant importance on information sharing processes (Grand et al., 2016; Janis, 1973; Stasser, 1999). Efficient team knowledge structures arise through continuous revision and convergence of individual knowledge structures as team members engage in collective sensemaking and understanding (Dionne et al., 2010). The extent to which team members adequately share and incorporate task-relevant information into their knowledge structures is thus critical to both the development of team cognition and overall team effectiveness. However, there is ample research evidence demonstrating that individuals working in teams are often ineffective at sharing and distributing

information across members. For example, research examining group decision-making contexts in which information is differentially distributed across individuals has revealed that members often fail to share uniquely held task-relevant information with others during team discussions (e.g., the hidden profile paradigm; Stasser & Titus, 1985, 2003). Work in this area has also revealed that the failure to share and incorporate task-relevant information held by members into a collectively shared understanding can lead to suboptimal decision and performance outcomes (e.g., groupthink; Janis, 1973).

From the perspective of team cognition, information sharing is a purposeful act carried out by an individual to disseminate knowledge pertinent to their team's performance/task environment in the service of generating, maintaining, and updating collective understanding (Fiore et al., 2010). Despite its recognition as a ubiquitous and central component of team cognition (Cooke et al., 2013; Mohammed et al., 2021), comparatively little attention has been directed to unpacking the process by which individuals make decisions about whether, when, and which information to share with their team members and the implications of these choice mechanisms for team knowledge emergence (Wang & Noe, 2010).

To this end, the purpose of this proposed work is to generate and advance formal theory on the process of information sharing in team knowledge development. Building upon existing theories of shared knowledge, this work seeks to specifically explore the processes through which task-relevant information is shared in teams. The present work proposes to develop a novel process-oriented account of how, when, and why team members share task-relevant information with team members. Furthermore,

the underlying theoretical assumptions and propositions of this theory will be instantiated in a computational model that will be used to explore its logical coherence, generate empirically testable predictions, and probe implications for influencing the emergence of shared team knowledge through a series of simulations and virtual experiments.

Chapter 1: Team Knowledge Emergence and Information Sharing

In their review of the research literature, Mohammed et al. (2021) conclude that research on team cognition is diffuse, poorly integrated, and (ironically) has struggled to reach shared consensus on how team cognition is defined. Across these varied examinations, empirical and theoretical treatments of team cognition have tended to focus on conceptualizing and operationalizing the structure/representation of team knowledge. From this perspective, much of the research on team cognition has concerned efforts to describe, measure, and quantify “what” teams know, the ways in which such information is organized and distributed within and across team members, and the implications that different distributions/patterns of cognitive representations across members hold for team functioning (Mohammed et al., 2000). Numerous different structures/representations have been proposed to reflect team cognition (e.g., team mental models, Cannon-Bowers et al., 1993; team situation awareness, Endsley, 1995; transactive memory systems, Wegner, 1987; representational gaps, Cronin & Weingart, 2007). However, the characteristic distinctions between these conceptualizations can generally be described by differences in (a) the extent to which the mental representations of team-/task-relevant information are similar (i.e., convergent) versus dissimilar (i.e., divergent) across members and (b) the content/type of information stored in individuals’ mental representations (e.g., facts, goals, procedures, assumptions, preferences).

Conceptualizations and efforts to empirically document the structure/representation of team cognition tend to approach the study of team knowledge as an *emergent construct*—a descriptive state or property of a team that is

brought about by the actions of constituent members (Marks et al., 2001). Implicit in such considerations is that team cognition is an outcome of a team knowledge *emergence* process—the series of actions, evaluations, and interactions carried out within and between individuals that give rise to the states, patterns, and distributions observable at higher levels (Grand et al., 2016). In contrast to the study of team cognition as an emergent construct and its structural representations, considerably less attention has been directed towards examining team knowledge emergence.

The limited work that has been directed towards examining team knowledge emergence has generally highlighted the significance of unpacking how team members interact, coordinate, and communicate in ways that influence their individual cognitive structures and knowledge representations (Cooke et al., 2000; Cooke et al., 2004; Kozlowski & Klein, 2000; Kozlowski et al., 2013). For example, Cooke and colleagues (2004) propose that team cognition is synonymous with team interaction, going so far as to suggest that collective (aggregated) knowledge is best understood by examining what and how members communicate and coordinate their actions across individuals to efficiently accomplish their task rather than its structural representation. Dionne and colleagues (2010) similarly expand on the position that collective knowledge is a function of communication in teams and demonstrate how information exchange networks among members (i.e., density and configuration of relationships between individuals) can affect the emergence of shared knowledge. Finally, Grand and colleagues (2016) proposed and empirically examined a formal theory of the team knowledge emergence process detailing how processes related to

information acquisition, dissemination, and acknowledgement can lead to unique patterns of collective knowledge in teams.

Understanding how team knowledge emerges requires examining the mechanisms through which individuals acquire, share, and integrate information within a team. In this vein, previous work has sought to develop formal computational models to propose and more precisely elaborate how these processes might unfold under different circumstances and their implications for individual and team functioning. For example, Stasser's (1988) model (Distribution of Information and Structuring of Shared Understanding in Social Systems, or DISCUSS) demonstrates how the distribution of information and decision preferences across team members interact with several cognitive mechanisms used by individuals during collective decision-making tasks are capable of producing suboptimal team outcomes. Similarly, Grand et al.'s (2016) formal theory and computational model of team knowledge emergence demonstrates how the unique behaviors and choices carried out by individuals in teams with distributed expertise/access to information can lead to unique patterns of collectively held knowledge. Though there are important differences across these perspectives, they exhibit several foundational similarities and highlight two core processes as central to team knowledge emergence—learning and sharing information. Furthermore, although both information acquisition and exchange processes are recognized as important to the emergence of team knowledge, the emphasis on team communication and coordination as particularly critical to the emergence of collectively shared knowledge suggests the need to further explore how information sharing processes unfold within and between team members. To this end,

the present work seeks to expand upon the foundations and assumptions of Grand et al.'s (2016) theory of team knowledge emergence and Stasser's (1988) DISCUSS model of group decision-making to propose a novel formal model that elaborates how individuals decide to share information in team contexts and to explore the implications of this theory for team functioning. The core processes and mechanisms of the theories proposed by Grand et al. (2016) and Stasser (1988) are briefly summarized below to provide context for development of the present work.

Dynamic Theory of Knowledge Emergence (Grand et al., 2016)

Grand and colleagues (2016) posit that knowledge emerges within teams through a dynamic, bottom-up process which includes learning and sharing of information. Learning information concerns how individuals select information within their environment, encode information into memory, and integrate learned information with their prior knowledge/information. Sharing information concerns how individuals decide when/whether to share information with others, choosing who to share information with, retrieving information from memory, communicating information to others, and finally acknowledging information shared by others.

Learning Behaviors

Unlike prior models of team knowledge-building (e.g., Dionne et al., 2010; Stasser, 1988), Grand et al.'s (2016) work focused on teams operating in contexts in which members do not initially hold any specific goal-/task-relevant information. Consequently, Grand et al. (2016) emphasize that how team members attend to and learn information from their performance environment and team members is an important component of collective knowledge development. To this end, Grand et al.

(2016) posit that individuals attend to information from their environment consistent with three assumptions. If operating in isolation of external influences (e.g., goals, organizational policies, managerial priorities), individuals will equally attend to any information relevant and accessible to their expertise. Second, individuals prefer to seek out information which contributes to a more comprehensive understanding of their task environment, such that information within their environment is prioritized when it builds upon prior knowledge. Finally, individuals can influence others to attend to information by directing attention toward and communicating new/important sources of information.

Depending on whether task-relevant information is found by searching the performance environment versus communicated by a team member, individuals are presumed to interpret that information through either encoding or decoding (Grand et al., 2016). Encoding information entails when an individual transforms external information which they have directly accessed from their performance environment into something they know or can remember. In this situation, the individual is described as “learning-from-self” to build knowledge because they are interpreting information in the absence of other team members. However, information is also acquired in team contexts that is shared by another member. Grand et al. (2016) differentiate this action/event as “learning-from-others” and posit that it tends to place additional demands on interpreting information (e.g., listening, filtering, translating, etc.) which can make learning slower.

A final action/event in the learning process described by Grand et al. (2016) involves an individual integrating information into categories of related concepts. In

this step, newly acquired information may become associated with prior knowledge, allowing for coherent schemas and better understanding of the task environment. These associations are crucial in forming cognitive structures that facilitate interpreting and sharing information with other team members.

Sharing Behaviors

As individuals build their own knowledge structures and personal understandings, the team's collective (aggregate) knowledge grows (Cooke et al., 2004; Grand et al., 2016). However, teams can only benefit from a comprehensive knowledge pool if individuals are able to access the unique knowledge held within their team (Mohammed et al., 2021). Thus, sharing information becomes an important step in converting *internalized knowledge* of the individual into holistically held, shared, and agreed upon *externalized knowledge*. Grand and colleagues (2016) propose sharing processes occur through individuals deciding to share information with others, choosing which information to share, communicating information to others, and confirming the information has been acknowledged and understood.

When sharing information, individuals elect to communicate with their team members. Grand and colleagues (2016) noted individuals tend to form patterns of engagement over time across team members, denoting a rate of engagement or participation (Stephan & Mischler, 1952). However, "who" and "why" members share is often bounded by the nature of an individual's information, task environment, team environment, and composition. For example, sharing information when team members have similar expertise can be used as a form of influence; by comparison, when team members have diverse knowledge areas, individuals must rely on each

other to share information that is task-critical which would otherwise be inaccessible or difficult to interpret by other members of the team. For purposes of their theory, Grand and colleagues (2016) do not elaborate how the decision to share information is made by an individual but note that it is likely a function of individual differences (e.g., extraversion), the team environment (e.g., norms/procedures), and other temporal states (e.g., available knowledge of individual, individual's current understanding of topic).

In addition to deciding whether to share information in their team, individuals must also decide which information to retrieve from memory and communicate. The theory of team knowledge emergence proposed by Grand and colleagues (2016) posited two assumptions related to how individuals make this decision: individuals prefer to share integrated and contextualized information over disjointed and/or decontextualized information and all members strive to have a complete understanding of all available information. Together, these assumptions engender an important consequence and boundary condition for the scope of Grand et al.'s (2016) theory in that they imply the likelihood that any known piece of (relevant) information being retrieved and shared by an individual should be equal because all information contributes towards the same single goal held by all individuals (i.e., learn all possible information available to the team).

After deciding to share and retrieve information, an individual can then transmit the information to their team members. This may occur through different mediums, such as face-to-face or email, which can affect when and/or how information is disseminated (e.g., real-time or delayed). Additionally, Grand et al.

(2016) propose that shared information must be acknowledged by other members of the team before it can be considered team knowledge. Acknowledgement is proposed to occur only if one or more of the receiving team members learn/internalize the shared information and results in the creation of “externalized” knowledge reflecting that the shared information is collectively held and agreed upon by members (Fiore et al., 2010).

DISCUSS Model of Group Information Sharing (Stasser, 1988)

Stasser’s (1988) DISCUSS model also provides a formal, process level framework for understanding how information becomes collectively represented in group-decision making contexts. In contrast to Grand et al.’s (2016) theory, which emphasizes knowledge emergence through bottom-up learning and sharing behaviors in unstructured environments, the DISCUSS model simulates how differences in initial information distribution and individual decision preferences interact to shape team discussions and final choices. The model assumes that individuals begin with access to different subsets of information and that each piece of information carries a valence indicating how strongly it supports a given decision alternative. Through a series of probabilistic rules, the model proposes how information is recalled, selected, shared, and integrated across group members over time.

Information Distribution and Selection

The DISCUSS model begins a pre-discussion phase, during which each individual attempts to recall pieces of information from their memory. The likelihood of recalling a particular item decreases as the total number of items known increases, reflecting a probabilistic memory sampling process. Once information is recalled,

individuals evaluate the relative support each item provides for competing decision options. Individuals then sum or average the valences of recalled items and form an initial decision preference.

A critical element modeled in DISCUSS concerns how group decision-making is affected by the initial distribution of information across members in the environment. Most notably, DISCUSS was the first model to attempt to formally represent the biased information sampling effect documented by Stasser and Titus (1985). Despite its name, the biased information sampling effect describes that even if individuals do not hold any inherent bias in what information they recall or choose to communicate, information initially known by multiple members (i.e., shared/common information) tends to be communicated more frequently than information initially known by only one member (i.e., unique information). Stasser and Titus (1985) note one explanation for the apparent preference for teams to communicate common rather than unique information could be a statistical artifact of the initial distribution of information within teams. Specifically, common pieces of information are more likely to be communicated at the group level simply by virtue of being known by multiple individuals. In other words, if each member has a pool of knowledge and sample randomly across all pools, items known by many individuals will appear more frequently in the discussion due to their higher base rate in the collective memory system. By formally modeling the dynamics of the information sampling bias effect, several boundary conditions and additional psychological accounts for the effect have also been advanced (e.g., Wittenbaum et al., 1999; Wittenbaum et al., 2004).

Information Sharing and Learning

Once individuals have formed an initial decision preference based on their own information, the DISCUSS model transitions to representing the dynamics of group discussion, where agents speak and share information with one another. At each time step, one individual is probabilistically selected to speak based on their participation rate, a parameter that reflects how frequently an individual tends to contribute to discussion. The selected speaker then chooses one of their previously recalled items to share. Stasser (1988) incorporates an advocacy mechanism during this stage which allows one to change the probability agents share items when they are consistent with their current decision preference. Thus, rather than randomly selecting from all recalled items, agents are more likely to advocate for their position by sharing information which tends to support their currently held position.

Once shared, a piece of information becomes available to all other group members. If the recipient already knew the information, they retrieve it from memory; if not, they learn it as new. The DISCUSS model also incorporates a norm parameter during the learning phase that governs whether information sharing serves to resolve differences in members' understanding of a piece of information (i.e., recipients update their evaluation of the item's valence based on the speaker's expressed view) or simply serves to communicate information but does not resolve any potential discrepancies.

Information Integration and Preference Updating

After receiving new information, DISCUSS proposes that group members recompute their decision preferences based on their current knowledge and decision

rules. This cyclical process of sharing, receiving, updating, and recalculating preferences continues across multiple model iterations (i.e., conversation turns) until the team either reaches consensus over their preferred decision option or enters a stalemate in which consensus does not exist but no new information has been introduced over several periods.

Stasser (1988) demonstrates that the DISCUSS model was able to reproduce a number of empirical findings in the group decision-making literature. Most notably, the mechanisms of DISCUSS replicated the difficulty with uncovering “hidden profiles” during team decision-making and the failure of groups to identify the optimal decision even when all necessary information is initially available somewhere within the team (Stasser & Titus, 1985). Stasser’s (1988) simulations further observed that DISCUSS produced more accurate results when it assumes members rely on information sharing to improve/resolve their understanding of information (i.e., information sharing is normative and contributes to the existence of shared understanding). Lastly, Stasser (1988) notes the simulated teams in DISCUSS communicate more common relative to unique information and tend to make suboptimal decisions irrespective of whether individuals advocate for/demonstrate bias towards choosing to share information based on their personal decision preferences. Taken together, the DISCUSS model highlights the path-dependent nature of team decision-making: early preferences, participation dynamics, and the structure of the discussion all interact to influence which knowledge becomes collectively held.

Implications for Team Knowledge Emergence

The DISCUSS model demonstrates how information distribution, participation rates, and preference-driven sharing shape collective knowledge construction. In contrast to Grand et al. (2016), who emphasize the role of individual learning, encoding, and interpersonal communication in shaping emergent knowledge, Stasser (1988) emphasizes how the structural constraints of group decision-making can shape the nature of knowledge emergence in teams. In many respects, these models reflect important differences in the assumed context and “processing objective” of the teams which they model and the influence these assumptions can hold for how theory on team knowledge-building is developed (Hinsz et al., 1999). Grand et al.’s (2016) work focused on knowledge-building in teams with distributed expertise and responsibilities, in which members with unique skillsets and access to resources are tasked with actively searching their environment for task-critical information, communicating that information to their team, and assuring all members were knowledgeable and agreed upon the available information. By comparison, Stasser’s (1988) work focused on individuals of roughly equal expertise but with unequal access to information whose primary goal was to reach consensus over a single course of action. By integrating insights from both models, researchers may gain a more complete picture of team knowledge emergence, capturing both the dynamic processes of team knowledge construction (Grand et al., 2016) and dynamics of group decision-making (Stasser, 1988).

Chapter 2: Advancing Theories of Team Knowledge Emergence

Grand et al. (2016) and Stasser (1988) both offer useful process-oriented accounts of how team cognition and knowledge can emerge within teams. Nevertheless, there are several assumptions and boundary conditions of both models and similar theories of team knowledge-building that provide opportunities for further development. For example, Grand et al.'s (2016) account is primarily limited to circumstances in which all members on a team are motivated, capable, and/or required to attain fully shared externalized knowledge such that all available information should be acquired and acknowledged by the entire team. Although this premise applies to many team contexts, it also serves as a basis for several important assumptions in Grand et al.'s (2016) theory (e.g., all pieces of information are equally likely to be shared, all information should eventually be shared given sufficient time) and places less emphasis on other mechanisms/considerations often highlighted as important in group contexts involving information sharing (e.g., impression management, team climate/culture; Morrison, 2023). Revisiting these assumptions and integrating new/alternative mechanisms into existing models of team knowledge emergence thus opens the potential for improving the explanatory quality of theories of team cognition as well as expanding the range of phenomena to which they may apply (e.g., why teams often do not share all information with each other; Janis, 1973; Stasser & Titus, 1985, 2003).

The present proposal seeks to advance theories of team knowledge emergence by directing attention to developing and specifying the mechanisms underlying two actions/events critical to the development of collective knowledge: whether and/or

when individuals choose to share information (i.e., who shares information) and how individuals choose which information to share with their teammates (i.e., what information is shared). These processes are fundamental to the emergence of knowledge in teams, as they are the primary causal determinants of what and how information is transmitted between individuals in a team, shape the completeness and coherence of both individual and shared knowledge, and can directly impact individual and team effectiveness (Mesmer-Magnus & DeChurch, 2009; Mesmer-Magnus et al., 2011). Furthermore, better specifying the mechanisms underlying who shares information and which information is shared can inform potential interventions and guidance for improving team cognition.

To this end, the present work proposes two additional considerations for characterizing how individuals determine whether and which information is shared in the context of team knowledge-building. The first posits that information sharing is influenced by an individual's evaluation of the *utility of information* they possess and describes how perceptions of the relevance and certainty of information that an individual has available to share can affect which information they may elect to share. The second posits that information sharing is also influenced by an individual's evaluation of the *utility of sharing information* and describes how an individual's experiences and perceptions of their team may affect whether they choose to share information with their teammates. Although these mechanisms have received significant research attention in the broader judgment/decision-making and social cognition literatures, they have largely not been considered in the work on team cognition or group information sharing. The following sections review seminal

research and advance several core concepts relevant to integrating these considerations into the formal theory of team knowledge emergence advanced in the present work.

Chapter 3: Utility of Information

When considering any given piece of information, individuals are proposed to hold perceptions of the usefulness of that information to the team's task. In this context, the utility of information can be defined with respect to its perceived consequence for both team goals and individual goals. When deciding which information to share then, people are proposed to formulate and behave in accordance with their perceptions of the utility of information they hold. In the present theory, evaluations of the utility of information is defined as a function of its perceived relevance and certainty to the task.

Task Relevance

The task relevance of information has been defined and operationalized in various ways across different areas of research; however, existing conceptualizations can broadly be categorized into those that define relevance subjectively versus objectively (Xu & Chen, 2006). Subjective definitions of task relevance tend to focus on differences in how individuals interpret the consequences of information attributable to, for example, differences in exposure to information, trust, or perceptions of source reliability (Xu & Chen, 2006; Schamber et al., 1990). In contrast, objective definitions of information relevance tend to focus on properties of the information and its relation to the task environment, including the “number of different tasks that make use of the information, the frequency of occurrence of those tasks, and the criticality of those tasks” (Durso et al., 2008, p. 755).

A common theme across discussions of information relevance though is the assumption that evaluations of relevance are based on the extent to which the

information contributes to and/or influences the accomplishment of specific goals and tasks (Espinosa & Clark, 2014; Huang & Cummings, 2011). For purposes of the present theory, the task relevance of information is defined in relation to the accomplishment of goals/tasks shared by team members and central to the team's purpose (Kozlowski & Ilgen, 2006). In this respect, the exchange of information in the present work maintains consistency with prior team cognition work which holds that information sharing establishes priorities and criteria for guiding collective decisions/behavior, stimulates individual and shared sense-making, and is used to inform decisions and coordinate behaviors in ways that holds consequences for team performance (Mohammed et al., 2021). Although the point of comparison for defining information relevance is fixed to shared goals/tasks in the current work, this does not mean all members necessarily hold the same interpretation of the perceived relevance of any given piece of information. Instead, members may hold different perceptions of the relevance of information to a team's task (e.g., Huang & Cummings, 2011). Differences in the perceived task relevance of information have the potential to contribute to differences in how individuals evaluate the utility of information and therefore their likelihood of sharing certain information.


Certainty

In research on judgment and decision-making, (un)certainty has been broadly characterized as an individual's beliefs about the propensity for an outcome/event to occur (Hastie, 2001). This definition is extrapolated to the present context to describe the propensity with which a statement, conclusion, and/or consequence associated with a piece of information is perceived to be true. Perceptions of certainty have long

been shown to influence how people integrate, react to, and use information in decision-making contexts (e.g., Kahneman & Tversky, 1984; Tversky & Kahneman, 1991). When considering how this may affect information sharing activities, the present theory posits individuals consider the certainty of information available to them when evaluating its utility and deciding which information to communicate to their teammates. The impact of information certainty has received little empirical or conceptual attention in prior work on team cognition or team knowledge emergence. However, the presence of and manner by which individuals resolve uncertainty has frequently been cited as an important contributor and challenge to group and organizational decision-making (e.g., Lipshitz & Strauss, 1997).

Beyond the incorporation of information certainty into theories of team knowledge emergence, the present work also seeks to account for the ways in which certainty is commonly expressed during information exchange. Specifically, although empirical research in decision-making has primarily operationalized and presented uncertainty statements as numerical probabilities (e.g., 90% chance that observation A will occur), individuals often summarize and communicate perceptions of uncertainty to others using probabilistic language (i.e., “likely”, “unlikely”; Kent, 1993; Wintle et al., 2019). Past work attempting to equate numeric probability values with probabilistic language has observed that people associate a wide range of potential numeric probability values with various natural language expressions of certainty (Figure 1). For example, Wintle and colleagues (2019) reported that the numeric probability values participants associated with the term “likely” ranged from 55% to 80%. The imprecision associated with interpreting and conveying the

Figure 1.
Natural Language Expressions of Certainty (Wintle et al., 2019)



Note. Figure reproduced from Wintle, B.C., Fraser, H., Wills, B.C., Nicholson, A.E., & Fidler, F. (2019). Verbal probabilities: Very likely to be somewhat more confusing than numbers. *PLoS One*, 14(4), e0213522. Copyright 2019 by Wintle et al. and distributed under the terms of the Creative Commons Attribution License.

certainty of information through natural language as so often occurs in team contexts thus reflects an additional important consideration for understanding how information sharing behaviors could affect the emergence of team knowledge.

Expected Utility of Information

The provided definitions for the relevance and certainty of information and their role in shaping which information individuals choose to share with their team members closely parallels similar concepts in economic/normative models of decision-making based on *expected utility* (EU). EU models of decision-making describe individual choice behavior as one in which individuals seek to maximize the EU of outcomes (Schoemaker, 1982). In its canonical form, these decision models define the EU of a potential choice option as the sum of the utility associated with a potential outcome state weighted by the probability of that state occurring for all possible outcome states associated with an option (e.g., $EU_{choice} = \sum p_{state} * u(outcome_{state})$; von Neumann & Morgenstern, 1947).

Many variants of the classic EU model have been proposed and examined. However, more contemporary models influenced by psychological research suggest that human decision-making is better predicted by models of *subjective expected utility* (SEU; e.g., Edwards, 1954). In contrast to EU models, SEU models hold that human decision-makers do not make decisions consistent with the explicitly stated value/utility and probability of outcome states, but rather the individual's subjective beliefs about those components. The subjective interpretations of the utility and probability of outcome states for a choice option represent the unique effects of a decision-maker's experiences, biases, and other cognitive processes that arise when they assess and select among potential courses of action. At their core, though, the formulation of SEU models is similar to EU models in that they still define the SEU of a choice as the sum of the (perceived) utility of outcomes associated with a given option weighted by their (perceived) probability of occurring.

Prospect theory (Kahneman & Tversky, 1979) and its revised variant cumulative prospect theory (CPT, Tversky & Kahneman, 1992) is one of the most widely cited and studied formal decision-making models of SEU. In CPT, individuals are proposed to interpret the value of potential outcomes relative to a reference point that is typically defined with respect to an individual's current standing on the value in question (e.g., money, satisfaction, etc.). An important implication of this assumption is that the utility of the outcomes associated with a choice option are interpreted in terms of gains and losses rather than final value of the outcome itself. Empirical research examining how individual decision-makers make choices amongst options concerning varying levels of gains versus losses finds that individuals tend to

be more sensitive to losses relative to equivalent gains (Kahneman & Tversky, 1979). The formalization of CPT thus posits that perceptions of the value of potential outcome states follows an asymmetric function that is shallower and concave for gains but steeper and convex for losses.

Research on decision-making under uncertainty also reveals that individuals do not appear to evaluate choice options in accordance with the stated probability of outcomes states. More specifically, individuals tend to make choices in ways that systematically overestimate the likelihood of lower probability events occurring while systematically underestimating the likelihood of higher probability events (Kahneman & Tversky, 1974; Tversky & Kahneman, 1992). To account for this pattern of observations, CPT proposes that decision-makers consider their “reactions/attitudes” towards the level of uncertainty associated with a potential outcome state rather than its stated likelihood of occurrence when evaluating potential courses of action. These reactions/attitudes are termed decision weights in CPT. The formalization of CPT thus transforms the likelihood of possible outcome states into decision weights through a sigmoidal function that biases lower probability events upwards and higher probability events downwards from their stated values.

Although CPT (and SEU models more generally) has been criticized as lacking a psychological account for its formulations, it has proven to be successful at accounting for empirical observations of decision-making under uncertainty (Schoemaker, 1982). The present work therefore uses the formalization of CPT to represent how individuals consider the perceived value (i.e., relevance) and likelihood (i.e., certainty) of information they have available to share with others to formulate

perceptions of its SEU. The proposed model further assumes that individuals generally prefer to maximize the usefulness of information they share with their team members. In this respect, an individual's perceptions of the SEU of information within the pool of potential information available for them to share is proposed to be a core mechanism for determining which pieces of information are likely to be shared with others in team knowledge-building contexts.

Chapter 4: Utility of Sharing Information

The prior discussion posits that generating and evaluating perceptions of the utility of information affects *which* pieces of information an individual may consider sharing with their teammates. An additional consideration, however, concerns the utility of engaging in information sharing as a behavior and thus *whether* an individual is likely to exchange information with their teammates. In contrast to evaluating the utility of information, evaluating the utility of sharing information is proposed to be influenced by an individual's beliefs about the consequences of their information sharing actions in the team/social context (e.g., "sharing information to assist my team requires effort/energy, so I should only be speaking if I believe it will be received/used by my team"; "my information sharing behavior is reflective of my disposition/image, so I should behave in ways that are consistent with my self-image"). In the present theory, evaluations of the utility of sharing information are a function of an individual's perceptions of the concordance of information, voice, and self-image goals. Critically, the present theory acknowledges and directly incorporates how these perceptions and their impact on the likelihood of sharing information unfold dynamically during team knowledge-building and shape how individual and collective knowledge emerges (Grand et al., 2016; Kozlowski et al., 2013).

Concordance of Information

Past research suggests that individuals consider how the information they could contribute to their team pertains to and may be incorporated into the group decision-making process (Xiao et al., 2016). One of the most commonly reported

evaluations in decision-makers' choices of whether to share a given piece of information is the extent to which that information is believed to be consistent with and/or likely to contribute to team members converging on a mutual understanding of task-relevant actions, decisions, or conclusions (e.g., Larson et al., 1996; Postmes et al., 2001; Stasser & Titus, 1985; 2003; Stasser et al., 1989; Wittenbaum et al., 2004). A robust body of empirical research reports that information which promotes concordance among the beliefs and conclusions of members is more likely to be communicated by individuals working in team decision tasks than discordant information that challenges or presents alternatives to existing knowledge. In contexts where information is distributed and/or restricted to certain members, one resultant pattern of this mechanism is that "unique" information (i.e., information known/understood by few members) tends to be less often shared than "common" information (i.e., information known/understood by many members) during team knowledge-building and discussion. Although the preference for expressing concordant information can lead to more widely shared and consensual team knowledge (Dionne et al., 2010), it can also result in poorer team decision-making and outcomes when unique/discordant information is relevant for successfully accomplishing team goals (Janis, 1973; Lu et al., 2011; Mesmer-Magnus & DeChurch, 2009).

Researchers have examined multiple explanations for why and when individuals are more likely to express concordant information. A common theme underlying accounts of the preference to share concordant information concerns individuals' motivations to regulate how they are perceived by others in the group or

perceive themselves (Stasser & Titus, 2003). For example, some accounts emphasize social validation and the potential “cost” that expressing unique/discordant information can hold for an individual’s status or inclusion in the group (e.g., Gigone & Hastie, 1993; Larson et al., 1996; Stewart & Stasser, 1995). Other accounts posit that exchanging concordant information engenders a positive feedback loop between establishing their social-image and enhancing their self-image (e.g., Kameda et al., 1997; Wittenbaum & Park, 2001; Wittenbaum et al., 1999). Specifically, sharing concordant information that is more readily understood/accepted by others tends to lead to being perceived as more credible and competent by one’s teammates; at the same time, receiving concordant information from others that is consistent with one’s existing knowledge contributes to more positive self-evaluations of credibility and competency. In either case, existing research highlights that individuals in team knowledge-building contexts appear to consider the concordance of the information they could potentially share when determining the utility of sharing that information. Importantly though, perceptions of concordance and the extent to which information is perceived as consistent with existing conclusions/beliefs in the team can only emerge as information exchange unfolds within the team over time and individuals infer what others in their team do and do not know. Consequently, the proposed theory and model seeks to incorporate the dynamics of this time varying perception into conceptualizations of team knowledge emergence.

Voice

The research and rationale underlying the effects of information concordance primarily imply motives related to social validation, comparison, and belonging as the

core drivers of whether individuals engage in information sharing during team knowledge-building. However, there are numerous motives and/or considerations that may influence individuals' perceptions of the utility of engaging in information sharing (e.g., Wang & Noe, 2010; Wittenbaum et al., 2004). For example, individuals in many organizational decision-making contexts are motivated to ensure that they and their teams attain the best possible outcomes. Communicating *any* information believed useful to that goal holds utility irrespective of its alignment with the current knowledge/beliefs of others in the teams. Indeed, such conclusions are consistent with research which finds that members encouraged to adopt norms of critical/counterfactual thinking (e.g., Kalinsky & Gray, 2002; Postmes et al., 2001) as well as those with more experience/expertise in the decision context (e.g., Stasser & Stewart, 1995; Stasser et al., 1995) are less likely to disproportionately favor sharing concordant over discordant information.

To this end, voice is posited as an important consideration in how individuals formulate perceptions of the utility of sharing information. Voice reflects the intent to communicate and bring about improvement or change, usually resulting from an individual's belief that something productive and/or desired (e.g., goal accomplishment, action, change) will come from their sharing information (voice instrumentality) and perceptions of psychological safety (Morrison, 2023). In organizational contexts, perceptions of voice are closely related to perceptions of autonomy, safety, and the extent to which employees feel they can exert control over their work behaviors and outcomes (Tangirala & Ramanujam, 2008). Importantly, research related to voice introduces several unique and competing

predictions/findings compared to those observed in prior research examining information sharing in group decision-making and knowledge-building.

For example, Hussain et al. (2019) found that employees were *less* likely to share information already known by others (i.e., common information) in teams. These authors suggest that, in organizational contexts, information redundancy tends to reduce perceptions of voice and therefore the perceived utility of communicating something already known. Relatedly, Tangirala and Ramanujam (2008) report evidence of a nonlinear/U-shaped relationship between supervisors' perceptions that an employee is willing to voice task-relevant but discordant information and an employee's self-reported perceptions of personal control (i.e., voice instrumentality). The authors interpret this pattern of results as stemming from two separate motivational mechanisms believed to operate at differing levels of personal control. When perceptions of personal control are low, dissatisfaction with the current state of events is proposed to motivate individuals to "assert" their autonomy and thus increase the likelihood of sharing discordant information; in contrast, when perceptions of personal control are high, individuals associate greater expectancies with their actions and are therefore motivated to exert their potential for change and share discordant information. At "moderate" levels of personal control, however, neither motivation is proposed to dominate and thus individuals are expected to be the least likely to express discordant information.

Just as an individual's understanding of whether a piece of information is concordant with existing knowledge can only develop over time as team members exchange information, an individual's perceptions of their voice instrumentality is

also posited to dynamically change and emerge as teams communicate. In the present theory, an individual's perception of their voice instrumentality is directly correlated with the extent to which the information they share is learned by others and subsequently incorporated into the team's actionable knowledge structure (Dionne et al., 2010; Fiore et al., 2010; Grand et al., 2016; McComb, 2007). Specifically, how other members acknowledge and react to shared information is proposed to be a signal/cue to the speaking team member about their voice within the team.

These signal cues can also affect the perceptions of safety within the team, another important factor related to someone speaking up. In addition to whether their ideas are useful, the belief that an individual will be safe, or protected, in their job, position, or relationships is imperative when an individual decides to speak.

Perceptions of safety can form through individual and group/relational factors. Individual characteristics which may affect safety are high core self-evaluations, relating to greater sense of confidence and approach motivation (Aryee et al., 2017). Safety also relates to an individual's personal initiative and felt responsibility. People will speak up when they feel more confident, have a higher approach motivation, and more responsibility (Ng & Feldman, 2015).

Relationships at work also affect psychological safety. Employees who believe they are respected by their coworkers are more motivated to speak up (Ng et al., 2021). Supportive colleagues also increase the sense of obligation for constructive change, increasing promotive voice, particularly around safety issues (Xie et al., 2015; Curcuruto et al., 2020). When witnessing other employees safely share information with leaders, an individual is more likely to share their own information

(Ng et al., 2021). This is particularly important for women, who tend to share less information unless they observe other women speaking up (Yan et al., 2022).

Individuals' centralization within their team network can also increase safety and encourage individuals to share ideas and thoughts (Venkataramani et al., 2016).

In sum, the present theory integrates how the dynamics of interpersonal perceptions affect the information exchange process in team knowledge emergence contexts. Specifically, when the information shared by an individual is positively received/evaluated by their team members, the perceived voice of the individual is proposed to increase (Liu et al., 2015). Teams may also be relatively neutral in their acknowledgement of the information, where they learn and utilize the information in discussion without any overt valence in reactions. In this instance, an individual may not have enough information to update voice perceptions and there is no change in voice. Finally, an individual's information sharing behavior may also be met with negative reactions from their members wherein, although members of their team may still learn/acknowledge the shared information, the sharing attempt received backlash, push back, or generated discord among other team members. This instance may leave the individual questioning their felt safety or usefulness within the team and thereby decrease their perceptions of voice within the team (Edmonson & Bransby, 2023; Hood et al., 2016).

Self-Image

The prior sections propose that dynamic perceptions related to information concordance, voice instrumentality, and psychological safety directly influence an individual's evaluation of the utility of sharing information. In addition to these

emergent sources of influence, the proposed theory also acknowledges that stable individual differences related to managing one's self-image can also affect a team member's information sharing behaviors (Bolino et al., 2008). Impression management theory describes the "process by which individuals attempt to control the impressions others form of them" (Leary & Kowalski, 1990, p. 34) and closely overlaps with research conducted on self-presentation. Such frameworks often assume that people hold relatively stable ideas (or self-concepts) of who they are and are motivated to produce an impression to others that is comparable to this image (Schlenker, 1975).

Fiske et al. (2002) propose that competence and warmth make up two universal dimensions upon which people evaluate themselves and others in ways that shape self-presentation and social interaction. Within the impression management context, these dimensions represent distinct motives: the desire to be perceived as competent (i.e., intelligent, effective, agentic) and the desire to be perceived as warm (i.e., trustworthy open, prosocial). These impression motives influence how individuals present themselves in social interactions, including decisions about what information to share with others. Evaluations of competence are linked to the motivation to be seen as capable of achieving performance-oriented goals, while evaluations of warmth reflect a concern with being seen as aligned with relational or communal goals (Cuddy et al., 2011; Fiske et al., 2007).

As previously noted, the preference to share concordant information during team knowledge-building tasks is shaped in part by concerns over how such exchanges affect others' perceptions (Kameda et al., 1997; Stasser & Titus, 2003;

Wittenbaum & Park, 2001; Wittenbaum et al., 1999). Impression management concerns related to warmth capture an individual's desire to avoid breaking norms and team preferences and in pursuit of social acceptance, they must be convinced that what is being shared is valuable (CITATION). Individuals motivated to appear warm may only choose to share information that strongly reinforces this image—favoring contributions that align with perceived team norms around openness, transparency, and regard for others. Since deviating from these norms may signal a lack of warmth, such individuals become more selective, preferring to share information that signals solidarity and interpersonal sensitivity. Prior work suggests that warmth impressions are especially fragile: while competence impressions tend to allow occasional lapses, warmth is more easily damaged and less readily restored (Skowronski & Carlston, 1987; Tausch et al., 2007). Accordingly, individuals with a strong desire to maintain a warm image may scrutinize the relational impact of the information they consider sharing (Goffman, 1959; Leary & Kowalski, 1990).

In contrast, the desire to be perceived as competent influences the breadth and risk profile of information agents consider for sharing. This behavior reflects a willingness to surface ideas that signal intellectual agility, problem solving capacity, or deeper thought. Drawing on decision-making research (Payne et al., 1993), such individuals may engage in broader or more integrative exploration of content to enhance their perceived competence. In contrast, those less motivated to appear competent may restrict their contributions to preference-consistent information, minimizing the risk of appearing ineffective or misinformed. This pattern is consistent with heuristic-based selection strategies (Gigerenzer & Gaissmaier, 2011).

As a result, those prioritizing competence motives may be more likely to introduce novel or high-value content in service of maintaining or enhancing a competent self-image.

Chapter 5: Sharing Information For Teams (SIFT) Model

Building upon the prior rationale and existing theories of team cognition (e.g., Dionne et al., 2010; Grand et al., 2016; Stasser, 1988), this theory seeks to integrate how individuals in contexts where team knowledge-building is critical to task performance formulate perceptions of the utility of information and the utility of sharing information and explore the implications of these mechanisms for understanding knowledge emergence, information exchange, and decision-making in teams. Specifically, a computational model was developed to represent the dynamic processes and mechanisms involved in information sharing, learning, and knowledge-building in team contexts. 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). Computational modeling provides several advantages for advancing theory and understanding dynamic and emergent phenomena, including more precise specification of the logic/rationale of a proposed explanatory account and the potential to evaluate their consequences critically examined through computer simulation methods (Davis et al., 2007; Epstein, 1999; Grand et al., in press; Harrison et al., 2007; Kuljanin et al., 2024; Rand & Rust, 2011; Vancouver & Weinhardt, 2012).

The proposed model is represented using agent-based modeling (ABM), which translates the hypothesized processes into computer code and allows for manipulation of key parameters of interest. ABM is particularly well-suited for representing and exploring systems in which macro-level outcomes emerge from the micro-level behaviors and interactions of individual entities over time (Rand & Rust,

2011). Unlike systems dynamics or equation-based modeling approaches—which typically focus on aggregated flows and feedback loops, ABMs allow researchers to explicitly model individual agents, each with their own attributes, decision rules, and capacity for adaptation. This granularity is especially important given that the current theory is centrally concerned with how individual-level perceptions, motivations, and impression management strategies shape patterns of information sharing in teams. Because these behaviors are dynamic, heterogeneous, and context-sensitive, an agent-based approach provides the necessary flexibility to simulate how such micro-level processes give rise to collective knowledge-building phenomena. The strength of ABM lies in its grounding in behavioral theory: when well-specified and empirically informed, ABMs can serve as powerful tools for both explaining and predicting emergent outcomes in real-world social systems.

The subsequent sections present the formal specification of the computational model, entitled SIFT (Sharing Information for Teams). Table 1 summarizes the *pseudocode* (i.e., written description of the model steps/computations) for SIFT that highlights how, when, and where the core concepts, actions/events, and mechanisms of the proposed theory are represented.

Model Initialization

The task environment modeled in the simulation closely mirrors the structure used in the DISCUSS model (Stasser & Titus, 1985), in which a team is presented with multiple discrete options, only one of which is objectively correct. While the original DISCUSS framework involves team members evaluating candidates and making a final collective decision, the current simulation does not include a formal

selection phase. Instead, it focuses on the unfolding process of information sharing over time, modeling how agents distribute and respond to shared knowledge within the team. Each agent holds a subset of the total information relevant to evaluating the options, and the simulation tracks how this information is shared or withheld. By adapting the DISCUSS task structure in this way, the model preserves the core epistemic challenge—accurately identifying the correct option based on distributed knowledge—while enabling a deeper exploration of how self-presentation concerns (e.g., the desire to appear competent or warm) shape communication behavior in teams. This process-oriented focus allows for investigation into the emergent properties of group knowledge development without requiring agents to converge on a final decision.

The initialization procedure for SIFT begins with construction of the information environment (Steps 1-2, Table 1). An information environment $\{J, I\}$ is comprised of j_{total} alternatives representing a finite set of decision options ($j \in J$) that a team is tasked with choosing between and i_{total} information pieces reflecting a finite set of information ($i \in I$) available to inform the team’s decision. Additionally, a mapping between information and decision options is initialized, defining how many (k_j) and which pieces of information are directly relevant to evaluating the utility of a given decision option (i.e., $i_{total} = \sum_{j=1}^{j_{total}} k_j$). Given these parameter settings for the information environment, an *objective knowledge pool* is initialized that defines (a) whether each information piece is positive, negative, or neutral towards each of the potential decision options and (b) the associated attributes of each information piece.

In SIFT, each information piece within a set of k_j information pieces is explicitly defined as being either supportive (i.e., positive), unsupportive (negative), or uninformative (neutral) towards option j , and neutral towards all other options $j \neq j$ (e.g., a piece of information defined as positive towards option A in the objective knowledge pool is considered neutral towards options B and C). Given this definition, information that is explicitly defined as neutral in a set of k_j information pieces is considered neutral towards all options (e.g., a piece of information explicitly defined as being neutral towards option A would also be considered neutral towards options B and C). Because this definition effectively renders neutral information uninformative with respect to any comparisons between decision options, the option associated with the set of information that contains the greatest net difference between the number of positive and negative pieces of information typically defines the “optimal” decision alternative in the information environment (i.e., choice with the strongest support given all available information).

After initializing the mapping between information pieces and decision options, the attributes for all information pieces in the objective knowledge pool are assigned. Each information piece is first given a “true” numeric probability level (p_i) reflecting its overall strength of certainty in the information environment:

$$p_i \sim Unif(p_{i_{lower}}, p_{i_{upper}}), \quad 0 < p_{i_{lower}} \leq p_{i_{upper}} \leq 1 \quad (1)$$

The lower and upper limits of this sampling distribution are manipulable parameters in SIFT and thus allow for constructing information environments in which the mean and/or range of informational uncertainty can vary. A corresponding verbal probability category is also assigned to each piece of information in the objective

Table 1
SIFT Model Pseudocode

Initialization

1. Create information environment with j_{total} decision options, i_{total} information pieces
2. Create objective knowledge pool that maps information to decision options and assigns information attributes (relevance, certainty)
3. Create a team with n members
4. Assign initial voice, competence, and warmth values to each member
5. Create initial knowledge pools for each agent based on objective knowledge pool and distribution of information in the environment
6. For each agent, assign numeric probability to a verbal probability category
7. For each agent, determine individual decision preference

Sharing Phase

8. Set time $t = 1$
9. For each agent, compute the probability of sampling each known information piece based on its expected utility and the agent's competence
10. Initialize decision iteration counter $d = 1$
11. Each agent n samples a piece of known information based on sampling probability
12. Compute perceived concordance based on agent's perception of the decision option preferred by the team
13. Compute utility of sharing the sampled information based on its perceived concordance and the agent's voice
14. Compute cumulative utility of sharing each piece of known information
15. Check whether any agent's cumulative utility of sharing for any information exceeds its speaking threshold based on agent's warmth:
 - a. If no agent exceeds threshold, increment decision iteration $d = d + 1$ and return to Step 11
 - b. If one or more agents exceed threshold, go to Step 16
16. Check whether information that has exceeded sharing threshold is eligible to share based on whether it has been shared recently:
 - a. If no agent has eligible information to share, increment decision iteration $d = d + 1$ and return to Step 11
 - b. If one or more agents has eligible information to share, go to Step 17
17. Select speaking agent:
 - a. If only one agent has identified an eligible piece of information to share, select that agent as speaker
 - b. If multiple agents have identified an eligible piece of information to share, select the agent with the highest voice as speaker
18. Speaking agent shares the eligible information piece with the highest utility to share
19. Speaking agent communicates the verbal probability and relevance of selected information with all other agents

Learning Phase

20. All agents update their knowledge pools with the cumulative number of times the spoken piece of information has been shared
21. Receiving agents learn and/or revise their interpretation of the verbal probability, numeric probability, relevance, and expected utility of the shared information

Acknowledgment and Reaction Process

22. Receiving agents react to shared information based on their competence, warmth, and perceived concordance of shared information

23. Speaking agent updates voice based on aggregate team reaction
 24. Each agent updates which decision option they individually prefer and which option they believe the team prefers based on their current knowledge pool
 25. If $t < t_{max}$, increment time $t = t + 1$ and return to Step 9; otherwise, end model
-

knowledge pool based on its sampled numerical certainty using the US Intelligence Community Directive ICD 203 guidelines (Wintle et al., 2019; Figure 1):

$$verbal_{p_i} = \begin{cases} \textit{Very likely}, & \textit{if } p_i > .80 \\ \textit{Likely}, & \textit{if } .55 < p_i \leq .80 \\ \textit{Roughly even}, & \textit{if } .45 < p_i \leq .55 \\ \textit{Unlikely}, & \textit{if } .20 < p_i \leq .45 \\ \textit{Very unlikely}, & \textit{if } p_i \leq .20 \end{cases} \quad (2)$$

Information pieces in the objective knowledge pool are next assigned a vector of relevance values (x_{ij}) that reflect the extent to which information piece i is supportive of decision option j :

$$x_{ij} \sim Unif(x_{ij_{lower}}, x_{ij_{upper}}), \quad -1 \leq x_{ij_{lower}} \leq x_{ij_{upper}} \leq 1 \quad (3)$$

The lower and upper limits of the distributions used to sample the relevance values for an information piece are manipulable parameters in SIFT. Furthermore, different lower and upper limits are used to determine each x_{ij} value depending on the mapping between information piece i and decision option j such that $x_{ij} > 0$ for positive information, $x_{ij} < 0$ for negative information, and $x_{ij} \approx 0$ for neutral information.

Following construction of the information environment, the team environment is initialized (Steps 3-7). A team environment is comprised of n agents representing the members of a team, each of which is uniquely defined by a set of attributes, knowledge, and decision preferences. To begin, each agent a is randomly assigned a numeric value representing its initial level of voice and fixed warmth and competence self-image perceptions:

$$\begin{aligned} voice_{at} &\sim Unif(voice_{lower}, voice_{upper}), \\ -1 &\leq voice_{lower} \leq voice_{upper} \leq 1 \end{aligned} \tag{4}$$

$$\begin{aligned} warmth_a &\sim Unif(warmth_{lower}, warmth_{upper}), \\ -1 &\leq warmth_{lower} \leq warmth_{upper} \leq 1 \end{aligned} \tag{5}$$

$$\begin{aligned} competence_a &\sim Unif(competence_{lower}, competence_{upper}), \\ -1 &\leq competence_{lower} \leq competence_{upper} \leq 1 \end{aligned} \tag{6}$$

The lower and upper limits of the sampling distributions used to assign these attributes to agents are manipulable parameters in SIFT and thus allow for creating different team compositions¹.

The next step in initializing the team environment in SIFT involves distributing the i_{total} information pieces from the information environment across members of the team to generate each agent's initial knowledge pool. The meaning and implications of how information is initially distributed within teams has been a central focus in the research on information exchange and decision-making in groups (e.g., Grand et al., 2016; Lu et al., 2012; Stasser & Titus, 2003), and there are myriad ways in which information could be distributed across members of a team. Prior research suggests that two considerations are especially significant with respect to their effects on team knowledge development and decision outcomes: (a) the ratio of common (information initially known/held by all members) to unique (information initially known/held by only one member or a subset of members) information reflected in members' initial knowledge pools (i.e., distribution of expertise, Grand et

¹ The present operationalization of SIFT treats agents' standing on initial voice, warmth, and competence as independent and uncorrelated, though this assumption can be relaxed in future versions of the model.

al., 2016); and (b) the extent to which the distribution of information across members is likely to result in individuals preferring decision alternatives that differ from one another and/or from the overall optimal decision alternative in the information environment based on the entire pool of available information (i.e., existence of a “hidden profile,” Stasser & Titus, 1985).

Given these considerations, SIFT currently permits a limited number of possible information distributions to be represented and used for constructing agent’s initial individual knowledge pools. Further details on the specific information distribution patterns implemented and the characteristics of the initial knowledge pools that can be created for agents in SIFT are provided in Chapter 7. Irrespective of the specific information distribution pattern initialized, the initial knowledge pools created for agents in SIFT are always generated such that each piece of information is initially known by at least one agent to ensure that the entire information environment is represented in the team’s collective/pooled knowledge. Additionally, the total number of information pieces initially known is constrained to be the same across all agents in a team.

The final steps of initializing the team environment involve specifying agents’ perceptions of the certainty, relevance, and utility of known information and formulating their initial decision preferences. With respect to the informational certainty of known information, the verbal probability category assigned to each piece of known information is reproduced exactly from the objective knowledge pool. However, the numerical certainty level an agent assigns to an information piece is

determined by randomly sampling from the numeric probability ranges associated with its verbal probability (cf., Equation 2):

$$numeric_{p_{ia}} = \begin{cases} \sim Unif[0.8, 1), & \text{if } verbal_{p_i} = \text{Very likely} \\ \sim Unif[.55, 0.8), & \text{if } verbal_{p_i} = \text{Likely} \\ \sim Unif[.45, 0.55), & \text{if } verbal_{p_i} = \text{Roughly even} \\ \sim Unif[.2, 0.45) & \text{if } verbal_{p_i} = \text{Unlikely} \\ \sim Unif(0, .2), & \text{if } verbal_{p_i} = \text{Very unlikely} \end{cases} \quad (7)$$

This operationalization means that all agents in SIFT are assumed to hold accurate knowledge of the “true” verbal probability associated with each piece of information. However, agents may differ from both the objective knowledge pool and other members with respect to their interpretations of the numeric probability they associate with a given piece of information.

The initial relevance values assigned to the known information in an agent’s knowledge pool are similarly derived from the corresponding relevance values in the objective knowledge pool plus a small degree of noise controlled by a constant model parameter (*encoding_x*):

$$x_{ijat} = x_{ij} + \sim N(0, encoding_x) \quad (8)$$

This operationalization again reflects that the relevance values individuals associate between a piece of information and the available decision alternatives may differ both from its “true” relevance as well as the relevance interpreted by other members.

SIFT proposes that individuals integrate their perceptions of informational certainty and relevance to formulate beliefs about the expected utility that each piece of known information holds for evaluating and choosing between potential decision alternatives at a given time point (expected utility of information; *EUI_{ijat}*). Perceptions

of EUI_{ijat} are represented in SIFT using the formalization of subjective expected utility proposed in cumulative prospect theory (Tversky & Kahneman, 1992):

$$EUI_{ijat} = w(\text{numeric}_{pia}) * v(x_{ijat}), \text{ where:} \quad (9)$$

$$w(\text{numeric}_{pia}) = \begin{cases} \text{numeric}_{pia} / (\text{numeric}_{pia}^\gamma + (1 - \text{numeric}_{pia})^\gamma)^{1/\gamma} & , \text{if } x_{ijat} \geq 0 \\ \text{numeric}_{pia} / (\text{numeric}_{pia}^\delta + (1 - \text{numeric}_{pia})^\delta)^{1/\delta} & , \text{if } x_{ijat} < 0 \end{cases}$$

$$v(x_{ijat}) = \begin{cases} x_{ijat}^\alpha & , \text{if } x_{ijat} \geq 0 \\ -\lambda(-x_{ijat})^\beta & , \text{if } x_{ijat} < 0 \end{cases}$$

For purposes of the present work, the empirical estimates reported in Tversky and Kahneman (1992) were used to define the free parameters of the decision weight (γ , δ) and value (α , β , λ) function curves.

The final element involved in initializing the team environment entails determining each agent's decision preference based on its initial knowledge pool. The procedure implemented in SIFT for computing an agent's initial decision preference replicates the procedure employed in Stasser's (1988) DISCUSS model. Notably, Stasser (1988) posits that individuals do not formulate their decision preferences using all of their known information, but rather a subset of this information that they are able to recall from memory. A stochastic sampling procedure defined by a recall probability parameter (r_{at}) is therefore used to determine which initially known information pieces are recalled by an agent a when forming its individual decision preference at time t . Based on research examining the effects of information load/interference, Stasser (1988) posits that the likelihood of recalling any given piece of information decreases as the amount of information a person knows/could recall

increases. The final formulation of recall probability implemented in Stasser's (1988) DISCUSS model and reproduced in SIFT is given as:

$$r_{at} = \frac{\mu}{known_{at}} + \tau, \quad 0 \leq r_{at} \leq 1 \quad (10)$$

where $known_{at}$ is the number of information pieces known by agent a at time t . The free parameters in Equation 10 (μ, τ) were defined using empirically derived estimates reported in Stasser (1988).

The formation of agent's initial decision preferences in SIFT thus proceeds by first computing r_{at} for each agent and then using this value to sample which pieces of known information each agent recalls from its initial knowledge pool. Using this set of recalled items ($recalled_{at} \subset \{known_{at}\}$), agents then independently identify which decision alternative they prefer according to the following formulation:

$$j_{\epsilon_{at}} = \arg \max_j (EU_{j_{at}}), \quad \text{where:} \quad (11)$$

$$EU_{j_{at}} = \sum_{i \in recalled_{at}} EU_{i_{j_{at}}}$$

This operationalization reflects that members are proposed to formulate their initial individual decision preference ($j_{\epsilon_{at}}$) by choosing the option for which the set of known information they can recall is associated with the highest aggregate expected utility.

Model Process/Algorithm

Consistent with prior theories and computational models (cf., Dionne et al., 2010; Grand et al., 2016; McComb, 2007; Stasser, 1988), the overall process of team information exchange and knowledge development modeled in SIFT is organized into a recurring sequence of actions/events organized into three phases: *sharing, learning,*

and *acknowledgement*. The sharing phase in SIFT proposes how team members decide whether to share information with their team, including which information members consider communicating and/or do communicate during decision-making deliberations. The learning phase in SIFT describes how team members incorporate and integrate information shared during team discussion into their individual knowledge pools. Lastly, the acknowledgement phase in SIFT represents how team members react to information shared during team discussion that subsequently informs interpretations about themselves, their team, and their decision preferences. Assumptions of the model are summarized in Table 2.

Sharing. The sharing phase in SIFT encompasses two of the most consequential elements of knowledge-building in teams as well as the areas where the conceptualization of SIFT expands most significantly upon prior work: determining (1) which agent chooses to engage in information sharing and (2) what information is shared at a given point in time (Step 9-19). For example, the models proposed by Grand et al. (2016) and Stasser (1988) assign agents time-invariant probabilities of speaking meant to reflect stable individual differences in a person's propensity to speak in a group. Once an agent elects to speak, Grand et al. (2016) and Stasser (1988) also posit that individuals randomly choose among known pieces of information when choosing what to share with their teammates². The model of team mental model convergence proposed by Dionne et al. (2010) posits a slightly more

² Grand et al. (2016) technically model the information selection process as a conditional probability in which an agent may only choose to share known information that their team has not yet fully learned and acknowledged. In contrast to Stasser (1988), this effectively eliminates the potential for agents/teams to select the same piece of information to communicate ad infinitum.

sophisticated procedure that posits an individual's likelihood of speaking can vary over time based on their overall "confidence" in the knowledge they possess and that individuals are also more likely to choose to share pieces of information in which they have the most confidence. However, changes in the propensity to speak in Dionne et al.'s (2010) model are posited to occur very slowly and do not consider how other individual- or team-level factors may affect a person's likelihood of participating in information sharing. As such, prior models have largely tended to conceptualize the speaking and information selection activities which unfold during team knowledge-building in terms of a relatively stable probabilistic sampling procedure consistent with an individual differences perspective.

In contrast, a core proposition of SIFT is that who speaks and what information they choose to share entails a more dynamic decision-making process. The foundational mechanics of the proposed information sharing process in SIFT draws from preference accumulation models developed in the judgment and decision sciences in which decision-makers are modeled as "internally deliberating" over potential courses of action by evaluating information for and against available choices until a single satisficing option is identified (e.g., Busemeyer & Townsend, 1993; Oppenheimer & Kelso, 2015; Weber & Johnson, 2008). In the context of team knowledge-building, SIFT posits that team members similarly evaluate and accumulate preferences towards sharing specific pieces of information at a given time point and that this procedure terminates once any single agent decides upon and communicates a piece of information they are satisfied with sharing.

Consistent with prior preference accumulation models, the information sharing procedure in SIFT begins by computing a dynamic sampling probability for each piece of information an agent knows at the current time point (s_{iat}). Conceptually, this probability can be interpreted as the likelihood that a team member considers sharing a given piece of known information with their teammates relative to the other pieces of information they have available to potentially share. The sampling probability for a piece of information in SIFT is computed as follows:

$$s_{iat} = \frac{e^{(1/e_{competence_a}) * |EUI_{ij\epsilon at}|}}{\sum_{i=1}^{i=known_{at}} e^{(1/e_{competence_a}) * |EUI_{ij\epsilon at}|}} \quad (12)$$

Equation 12 is a standard softmax function commonly used in reinforcement learning and similar machine learning paradigms to generate a probability distribution over a set of possible states (Sutton & Barto, 1998). In SIFT, this function is used to transform the perceived $EUI_{ij\epsilon at}$ an agent holds for a piece of information into a value that can be interpreted as a probability (i.e., $0 \leq s_{iat} \leq 1$; $\sum s_{iat} = 1$) such that information perceived as more useful generally has a higher probability of being selected/considered.

However, the formalization of the information sampling probability in SIFT also embeds some additional assumptions intended to acknowledge certain considerations of collective knowledge-building in organizations. First, SIFT assumes that individuals working in organizational contexts are motivated to share information they perceive as more consequential to the team's task (e.g., Espinosa & Clark, 2014; Huang & Cummings, 2011; Xu & Chen, 2006; Schamber et al., 1990). The implied team processing objective (Hinsz et al., 1999) in SIFT is for agents to make a well-


informed collective decision. The prior definition of *EUI* as the utility that a given piece of information is perceived to hold for evaluating and choosing between potential decision alternatives thus suggests that the *magnitude* of this value reflects the degree to which a piece of information is perceived as consequential, irrespective of whether that piece of information is positive, negative, or neutral towards a potential option. Equation 12 therefore relies on the absolute value of *EUI* to reflect that agents will generally be more likely to consider sharing a piece of information that has an expected utility further from zero.

A second key assumption in the information sharing process represented in SIFT is that the extent to which more consequential information is more likely to be considered is proposed to depend on the agent's perceptions of its competence. Specifically, SIFT proposes that individuals with lower competence are more likely to prioritize sharing information they perceive as consequential, whereas individuals with higher competence are less sensitive to differences in the perceived importance of information when considering which information to share. Figure 2 provides an illustration of this functional relationship. The top panel of the figure plots hypothetical *EUI* values for information pieces known by an agent, whereas the bottom panel plots the corresponding information sampling probabilities that an agent with a high competence perceptions versus an agent with low competence perceptions in SIFT would assign to those pieces of information based on Equation 12. As Figure 2 demonstrates, agents with lower competence perceptions assign sampling probabilities to information pieces in direct proportion to the magnitude of their *EUI* values (i.e., information with higher $|EUI|$ values have much larger sampling

probabilities than information with lower $|EUI|$ values). In contrast, the sampling probabilities assigned by agents with higher competence perceptions are more

Figure 2.

Example of Information Sampling Probabilities in SIFT Based on Agent Competence Perceptions and Expected Utility of Information



Note. Panel A depicts the absolute expected utility (EU) values for a hypothetical set of information. Panel B demonstrates how the EU values from Panel A would be transformed into information sampling probabilities for an agent with a very high competence self-image (competence_a = 1; green bar) and an agent with a very low competence self-image (competence_a = -1; orange bar) in SIFT based on Equation 12.

uniform across the set of known information (i.e., information with higher $|EUI|$ values have similar sampling probabilities to information with lower $|EUI|$ values). A key implication of this operationalization is that agents with higher competence perceptions in SIFT will generally tend to consider sharing information with more varied EUI during team knowledge-building than agents with lower competence perceptions.

A final important assumption represented in Equation 12 is that an agent's interpretations of which piece of information they might consider sharing is proposed to be conditioned on the decision option it believes is best at the current time point ($j_{\epsilon_{at}}$). This assumption reflects that an individual's perceptions of what it means for a piece of information to be "useful"/consequential may change over time as their understanding of and preferences for different decision alternatives evolves (Payne et al., 1993). Note that conditioning the likelihood of sharing a piece of information on an agent's current decision preference could be interpreted as modeling a form of "advocacy" in that it means agents in SIFT will tend to demonstrate a preference towards sharing information that is most consequential to their preferred decision alternative (cf., Stasser, 1988). However, by incorporating the absolute (as opposed to the signed) value for the expected utility of information in Equation 12, agents in SIFT are not more predisposed to consider sharing information *supportive* of their preferred decision alternative (i.e., information in which $EUI > 0$ towards the agent's current preference) than information that is *not supportive* of their preferred decision alternative (i.e., information in which $EUI < 0$ towards the agent's current preference). In this sense, SIFT does not conceptualize an agent's decisions about which information to share as a means of influencing or attempting to convince others to endorse a particular decision. The information sharing process in SIFT instead treats an individual's belief about their preferred decision alternative as emergent and responsive to the information exchanged within the team. This dynamic belief updating subsequently affects the set of information an individual perceives as more versus less useful to share in the current moment (cf., Thomas et al., 2008).

Once the dynamic information sampling probabilities are determined, agents in SIFT begin the process of internally evaluating pieces of information to potentially share at the current time point (Steps 10-16; Table 1). As a general overview, SIFT models each agent as computing momentary dynamic perceptions about the utility of sharing various pieces of known information (expected utility of sharing, EUS_{ijatd} ; cf., Busemeyer & Townsend, 1993). These momentary perceptions are proposed to accumulate as an agent internally deliberates over d evaluation periods and are used to determine whether an agent chooses to speak at a given time point and, if so, which piece of information it elects to share. The first step in the computation of EUS_{ijatd} during the internal deliberation process involves each agent selecting a single piece of known information from its knowledge pool according to its sampling probability (S_{iat}). The agent then interprets the concordance of the sampled information according to the following formulation:

$$concordance_{iatd} = \begin{cases} 1 & , \text{if } t = 1 \\ EUI_{ij_{\omega}at} & , \text{if } t > 1 \end{cases} \quad (13)$$

Note that the operationalization of information concordance proposes that individuals consider the decision option they believe their *team* currently prefers when interpreting the utility of potentially communicating a given piece of information to others (i.e., $j_{\omega at}$; see Equation 21). In this respect, Equation 13 represents one pathway through which the contextual influence of the team environment dynamically shapes the cognitive and behavioral activities of an individual in SIFT (e.g., Kozlowski & Klein, 2000). Because interpretations of concordance require an agent to first formulate a belief regarding the team's preferred decision option, the

information concordance value for any information considered during the first simulated period (i.e., prior to any team communication) is set to one.

The momentary expected utility of sharing the selected information during the current deliberation period is then computed by weighting its perceived concordance with the agent's voice at the current time point. Consistent with empirical research on voice (Hussain et al., 2019; Morrison, 2023; Tangirala & Ramanujam, 2008), SIFT models that as an agent's perceived voice within its team increases (decreases), it should feel more (less) inclined to speak. To maintain the assumption that individuals perceive greater utility in sharing more consequential information (see Equations 12), an agent's current level of voice is linearly transformed such that $0 \leq \text{voice}'_{at} \leq 1$ and then multiplied by the absolute value of the perceived concordance of the information under consideration to determine its momentary expected utility of sharing:

$$EUS_{iatd} = \text{voice}'_{at} * \text{abs}(\text{concordance}_{iatd}), \quad \text{where:} \quad (14)$$

$$\text{voice}'_{at} = \frac{\text{voice}_{at} - (-1)}{1 - (-1)}$$

Lastly, agents compute the cumulative expected utility for sharing each piece of known information by summing together their momentary perceptions of utility across the total number of internal deliberation periods (n_d) that have occurred thus far during time point t :

$$EUS_{iat} = \sum_{i \in \text{known}_{at}} \sum_{d=1}^{d=n_d} EUS_{iatd} \quad (15)$$

The cumulative expected utility of sharing calculated for each piece of information at the end of a deliberation period can be interpreted as the degree to

which an agent believes it is useful to share that information piece with their team. Within each agent, information that has accumulated higher EUS_{iat} is perceived as being more useful to communicate than information that has accumulated lower EUS_{iat} . The overall magnitude of this cumulative belief determines whether an agent has identified a piece of known information it would like to share and, by extension, whether that agent may elect to speak during the current time point according to the following rule:

$$speak_{at} = \begin{cases} TRUE, & \text{if } EUS_{iat} > (W_s + warmth'_a) \text{ for any } i \\ FALSE, & \text{otherwise} \end{cases}, \quad \text{where:} \quad (16)$$

$$warmth'_a = \frac{warmth_a - (-1)}{1 - (-1)} \text{ and } 0 < W_s < \infty$$

The formulation of Equation 16 reflects that a team member who speaks is assumed to have identified a piece of information that they perceive is sufficiently useful to share with their team. In SIFT, this sufficiency criteria is operationalized as a boundary threshold determined by an agent's perceptions of its warmth. Specifically, SIFT proposes that individuals with higher (lower) warmth possess a higher (lower) internal criterion for deciding whether they have identified a piece of information they perceive is consequential enough to communicate. Given that SIFT assumes agents interpret the absolute magnitude of the expected utility of information (Equations 12, 14) and therefore EUS_{iat} will always be greater than 0, Equation 16 relies on a linear transformation of warmth (i.e., $0 \leq warmth'_a \leq 1$) to ensure that the internal criterion threshold an agent uses to evaluate whether it is prepared to speak is always positive. Additionally, Equation 16 incorporates a free parameter (W_s) to account for differences between the scaling of the values used to parameterize warmth and the expected utility of information in SIFT.

The preceding computations comprise a single deliberation period of the information sharing and selection process in SIFT. This information selection cycle iterates until at least one agent decides upon a piece of information it is willing to share with the team. At that point, two consistency checks are performed to further regulate the information sharing process modeled in SIFT.

First, the purely “bottom-up” information selection process modeled in SIFT can result in circumstances in which one or more agents may repeatedly identify the same piece of information to share across many successive time points. Although repeating the same information occurs during team knowledge-building activities (e.g., Knight, Kennedy, & McComb, 2016) and has been identified as contributing to collective learning and understanding (Grand et al., 2016), it seems implausible that motivated individuals working in organizational team contexts would persist in repeatedly communicating a single piece of information for extended durations of time. As such, a check is implemented in SIFT that can be used to limit agents from choosing to share a piece of information that has previously been communicated until a specified interval of time has elapsed (*interval_i*). If an agent identifies a piece of information that it perceives as sufficiently useful to share that has already been communicated by any agent within the prior *interval_i* time points, the information selection process will continue. The second consistency check implemented in the information selection and speaking process of SIFT controls what occurs if multiple agents are prepared to share eligible information at the same information selection period. In such cases, the agent with the highest voice at that time point is chosen to speak.

Of final note, the EUS_{ijat} of all information for all agents is set to zero at the start of information selection process during the first time period (i.e., for all a , $EUS_{ijad} = 0$ at $d = 0$ when $t = 1$). However, in subsequent time periods, the final cumulative expected utility of sharing computed for each piece of information at the end of the information selection process is assumed to carry over to the beginning of the next time period for all agents who did *not* speak at the prior time point (i.e., for all a that did not speak at $t-1$, $EUS_{ijad} > 0$ at $d = 0$ for some i when $t > 1$). This assumption reflects that once an agent has deliberated on what to share, that decision remains cognitively available in future steps unless actively re-evaluated. In other words, agents retain the outcome of their previous deliberation until either contextual changes (e.g., updated voice or team preference) or internal thresholds motivate a new deliberation. While the model does not include continuous updating based on others' behavior, this approach is loosely consistent with research suggesting that individuals form and hold tentative judgments about the utility of contributing, which can be re-engaged as the group context evolves (Bandura & Walters, 1977). Thus, maintaining prior deliberation outcomes allows the model to capture some of the inertia and efficiency seen in real-world teams—where individuals often consider contributing, defer action, and later return to their original reasoning rather than starting from scratch.

Once a single agent has identified a single eligible and sufficiently consequential piece of information to share, the sharing phase concludes with that agent communicating the selected information to its team (i_s ; $i_s \in I$). In SIFT, a sharing agent is proposed to communicate its selected piece of information to all

other agents in the team simultaneously. Furthermore, the sharing agent is proposed to communicate only its understanding of the verbal probability ($verbal_{pia}$) and relevance (x_{ijat}) of the selected information with other agents on the team.

Learning. The learning phase of SIFT characterizes how team members are proposed to update their individual knowledge pools following an information sharing event. In the present conceptualization of SIFT, all non-sharing agents are assumed to actively attend to all information sharing attempts and attempt to learn the verbal probability and relevance values communicated by the sharing agent. The learning phase begins with agents updating a cumulative count of the number of times the shared information piece has been communicated during the entire team knowledge-building process ($count_{it}$). Because SIFT presumes that all agents receive and attend to every piece of shared information, all agents have an equally accurate “memory” of which and how often information has been discussed in their team.

The next steps of the learning process describe how agents learn the verbal probability and relevance of shared information. SIFT implements two potential processes by which these informational attributes are learned based on whether the shared information is not yet known (i.e., acquiring new information) versus already known (i.e., revising known information) by the receiving agent. With respect to learning the verbal probability of shared information, agents for whom the shared information is new learn the verbal probability of the shared information piece exactly as it is communicated by the sharing agent. Upon learning the verbal probability estimate, these receiving agents then generate a unique numerical probability estimate for the shared information following the formulation in Equation

7. In contrast, agents for whom the shared information is already known retain their previous/already existing understanding of the verbal and numerical probability of the shared information. Note that because the verbal probability category assigned to each piece of known information in an agent’s knowledge pool during model initialization is reproduced exactly from the objective knowledge pool, the result of this procedure is that all agents—irrespective of whether they are acquiring new information or revising known information—will hold the same understanding of the shared information’s verbal probability following learning. However, agents’ interpretations of the numerical probability are likely to differ from one another.

The proposed formalization for how agents learn the relevance values of shared information in SIFT is as follows:

$$x_{i_s j a_{rcvr} t} = \begin{cases} x_{i_s j a_{spkr} t} + \sim N(0, decoding_x), & \text{if } i_s \notin \{known_{at}\} \\ lrn.rate * (x_{i_s j a_{spkr} t} - x_{i_s j a_{rcvr} t-1}) + x_{i_s j a_{rcvr} t-1}, & \text{if } i_s \in \{known_{at}\} \end{cases}, (17)$$

where: $0 \leq lrn.rate \leq 1$

Like the procedure used for determining the relevance of known information during model initialization, Equation 17 indicates that receiving agents which are acquiring new information they do not yet know during the learning phase may not learn/interpret the relevance of the shared information exactly as communicated by the speaking agent. Consistent with prior work which notes that the accuracy of learning during team knowledge-building may vary based on the source of the information (e.g., Grand et al., 2016), SIFT permits the potential degree of error expected when “acquiring” information from the objective knowledge pool during model initialization ($encoding_x$ in Equation 8) to differ from the degree of error expected when learning information through sharing ($decoding_x$ in Equation 17). By

comparison, Equation 17 reflects that agents which receive known information following a sharing attempt learn by reducing the discrepancy between their current understanding of the information's relevance and that communicated by the speaking agent by a specified percentage (*lrn.rate*). SIFT thus assumes that agents will tend to converge in their perceptions of informational relevance over time as information is shared (Dionne et al., 2010; McComb, 2007; Stasser, 1988). In the final step of the learning process implemented in SIFT, agents (re)compute their perceptions of the expected utility of the just shared information for each of the j decision options based on their new/revised interpretations of that information's certainty and relevance following Equation 9.

Acknowledgment. The final phase of SIFT represents how members of a team are proposed to acknowledge and subsequently update their understanding of the team and task environment based on the information communicated at the current time point (Step 22-25). Most notably, the acknowledgement phase characterizes how changes to agent voice and decision preferences are purported to occur in SIFT. This process begins when all receiving agents independently generate a reaction to the newly shared information. In the context of real-world teams, this "reaction" can be thought of as an observable or inferred response—ranging from explicit verbal feedback (e.g., agreement, challenge, or elaboration) to more implicit behavioral cues (e.g., nodding, ignoring, or adopting the suggestion without comment). Such reactions provide critical interpersonal signals about the value, relevance, or credibility of information shared (e.g., Edmonson, 1999; Mesmer-magnus & DeChurch, 2009) and

play a central role in knowledge building and social learning within teams. SIFT

formulates an agent's reaction to a piece of shared information as follows:

$$reaction_{a_{recv}t} = \begin{cases} 1, & \text{if } competence_a < warmth_a, \quad EUI_{i_s j_{\omega}at} > 0 \\ -1, & \text{if } competence_a < warmth_a, \quad EUI_{i_s j_{\omega}at} < 0 \\ 1, & \text{if } competence_a > warmth_a, \quad EUI_{i_s j_{\omega}at} > 0, \quad EUI_{i_s j_{\epsilon}at} > 0 \\ 0, & \text{if } competence_a > warmth_a, \quad EUI_{i_s j_{\omega}at} > 0, \quad EUI_{i_s j_{\epsilon}at} < 0 \\ 0, & \text{if } competence_a > warmth_a, \quad EUI_{i_s j_{\omega}at} < 0, \quad EUI_{i_s j_{\epsilon}at} > 0 \\ -1, & \text{if } competence_a > warmth_a, \quad EUI_{i_s j_{\omega}at} < 0, \quad EUI_{i_s j_{\epsilon}at} < 0 \end{cases}, (18)$$

where: $t > 1$

Equation 18 reflects that members of a team may perceive a piece of shared information positively, negatively, or neutrally (Dionne et al., 2010; McComb, 2007). The formulation of agent reactions in SIFT proposes that if an agent perceives its self-image as more warm than competent, its reaction is driven exclusively by whether the agent perceives the shared information to be supportive (i.e., concordant) versus unsupportive (i.e., discordant) of the decision option the agent believes is preferred by the team ($j_{\omega at}$; Piaget, 1975). In contrast, the reaction of an agent that perceives its self-image as more competent than warm takes into account whether the shared information is perceived to support both (positive reaction), neither (negative reaction), or is mixed towards (neutral reaction) the agent's preferred decision option ($j_{\epsilon at}$) and the option they believe is preferred by the team (Fiske et al., 2007; Leary & Kowalski, 1990). Note that agents in SIFT formulate their reactions to shared information prior to updating their beliefs about which decision option they prefer or believe the team prefers. This assumption reflects that team members often respond to new input based on initial affective, evaluative, or interpersonal cues before fully integrating it into their internal decision-making framework. In real teams, such reactions may be shaped by heuristic appraisals—such as perceived credibility of the

speaker, fit with prior beliefs, or social alignment—allowing for a quick interpersonal signal to be sent before more effortful belief revision occurs (Kahneman, 2011; van Ginkel & van Knippenberg, 2009). Structurally, this sequencing separates the social function of reaction from the cognitive function of belief updating, preserving the distinction between immediate interpersonal feedback and deeper shifts in team knowledge or preference alignment.

Once the reactions of all receiving agents are computed, the speaking agent which shared information during the current time point is proposed to react to how its most recent communication effort was received within the team. The current conceptualization of this operation in SIFT assumes that the speaking agent formulates an overall impression of the team’s reaction to its sharing attempt as an equally weighted sum of the reactions expressed by each receiving agent:

$$reaction_{a_{spkr}t} = \begin{cases} 1, & \text{if } \left(\sum_{a_{rcvr}} reaction_{a_{rcvr}t} \right) > 0 \\ 0, & \text{if } \left(\sum_{a_{rcvr}} reaction_{a_{rcvr}t} \right) = 0 \\ -1, & \text{if } \left(\sum_{a_{rcvr}} reaction_{a_{rcvr}t} \right) < 0 \end{cases} \quad (19)$$

The speaking agent’s interpretation of how its information sharing effort was received is subsequently used to update perceptions of its voice within the team:

$$voice_{a_{spkr}t} = \left(reaction_{a_{spkr}t} * \Delta voice \right) + voice_{a_{spkr}t-1}, \quad -1 < voice_{a_{spkr}t} < 1 \quad (20)$$

$$\text{where: } \Delta voice = \Delta voice_{max} * \frac{\varphi^{(1-voice_{a_{spkr}t-1})} - 1}{\varphi - 1}$$

Equation 20 represents that an individual's perception of their voice is generally expected to increase following a positive reaction, decrease following a negative reaction, and remain unchanged following a neutral reaction from their team.

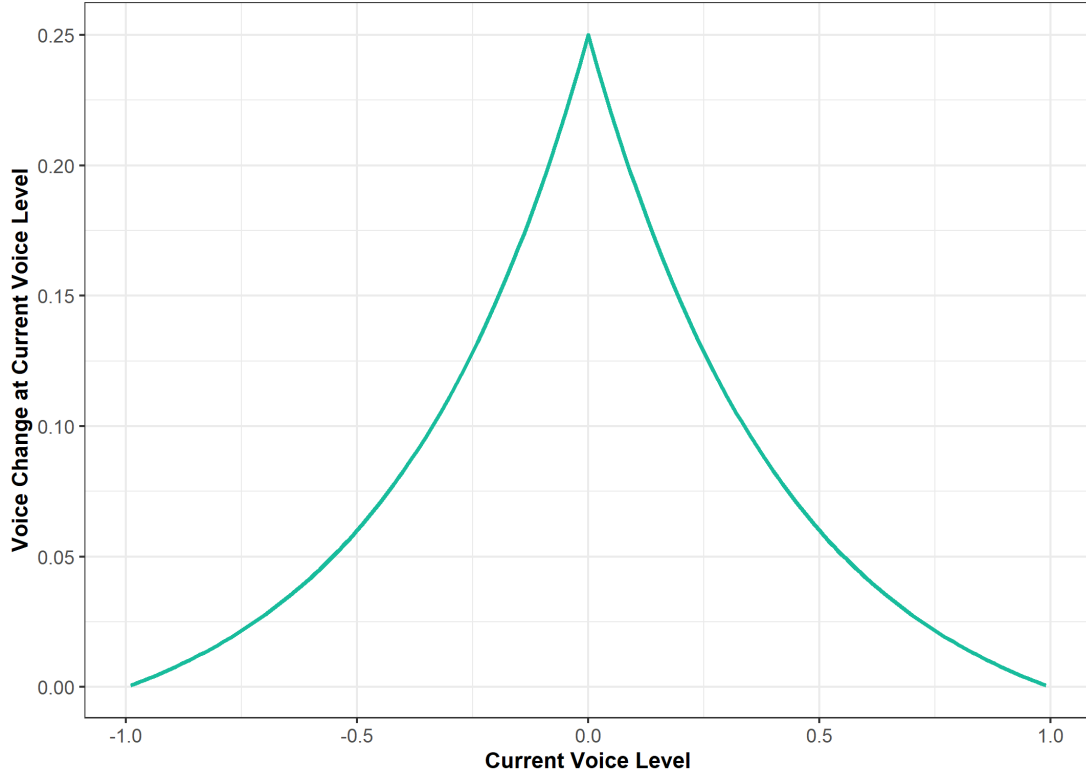
Equation 20 also incorporates the assumption that the amount by which an agent's voice is proposed to change ($\Delta voice$) depends on its current voice level (Edmondson, 1999). Figure 3 provides a visualization of the function used for determining how much an agent's voice changes after it shares information in SIFT. The specification of the voice change function indicates that voice is proposed to change more slowly for individuals who believe they currently possess either very low or very high voice in their team (e.g., individuals who already feel very (un)confident that something productive will come from their information sharing are less sensitive to how the team reacts to their contributions during team knowledge-building). However, the team's reaction is proposed to be more consequential when the individual's perceived voice lies in between these extremes (e.g., individuals who feel moderately confident that something productive may come from their information sharing are more sensitive to how the team reacts to their contributions during team knowledge-building).³ The degree to which an individual's voice change depends on their current voice (φ) as well as the maximum amount by which voice can change following an information sharing event ($\Delta voice_{max}$) are both free parameters in SIFT.

The final steps of the acknowledgement phase entail determining agents' beliefs about the decision option they believe is favored by their team based on which

³ The present operationalization of Equation 20 models voice changes as a locally independent process over time (i.e., no path dependency). This assumption can be changed in future model versions.

information has been shared and which decision option they personally prefer based on which information they considered during the current time point. Specifically, SIFT models that individuals formulate beliefs about which decision option their

Figure 3.
Visualization of Δ Voice Function in SIFT



Note. The vertical axis reflects the absolute potential magnitude of the voice change for an agent given a current level of voice (horizontal axis) based on Equation 20 when $\varphi = 10$ and $\Delta voice_{max} = 0.25$. The direction (+/-) of the actual voice change for a speaking agent during a given time point is determined by the aggregate team reaction and can be 0 if the aggregate team reaction is neutral (see Equation 19).

team prefers as:

$$j_{\omega at} = \arg \max_j (EU_{jat}), \quad \text{where:} \quad (21)$$

$$EU_{jat} = \sum_{i \in shared_t} count_{it} * EUI_{ijat}$$

This operationalization reflects that a member's inference regarding which decision option they believe is preferred by the team ($j_{\omega at}$) is formulated by considering the

aggregate EUI_{ijat} values for each decision option based only on the subset of information which has been shared within the team during its discussion up to the current time point ($shared_t \subset \{known_{at}\}$). Notably, this means that an agent's beliefs about its team's decision preference does not incorporate any information that an agent may know independently but has not yet communicated. Equation 21 further reflects that the EUI_{ijat} values associated with each piece of shared information are weighted by the total number of times a piece of information has been shared. This formulation captures that a person's evaluation of how their team perceives each decision option is influenced by the extent to which the team directs more time/attention towards sharing information perceived as highly consequential to each option (Festinger, 1954).

The acknowledgement phase in SIFT concludes with each agent updating its current perception for which decision option it perceives as most preferable ($j_{\epsilon at}$). The procedure to determine each agent's individual decision preference relies on a similar formulation to that used during model initialization in which an agent's initial preference is the option with the highest aggregate expected utility based on a subset of the information initially known by the agent (see Equations 10-11). However, once the team knowledge-building process is underway, SIFT models agents as basing their decision preferences on the set of information they considered communicating during their internal information selection deliberations and the piece of information that was communicated to the team during the current period. This operationalization thus treats the information that is sampled and shared/received during the sharing and learning phases of SIFT (respectively) as being actively recalled by the agent and

therefore available to incorporate into its decision preference at the current time period (cf., Thomas et al., 2008).

Table 2.

SIFT Model Assumptions

Assumptions
1. Agents begin with an initial voice expectation of their team.
2. Agents know information relevant to team task prior to team formation.
3. Agents evaluate information utility based on their individual preference.
4. Agents evaluate information based on the magnitude of utility, not valence.
5. Higher (lower) competence results in information sampling probabilities that are more uniform (proportional to the information piece's EU)
6. Agents evaluate utility only concerning their current belief about the team's preference.
7. Agents weigh voice and concordance equally when determining utility of sharing.
8. Agents' decision threshold is only affected by warmth.
9. Agents do not consider prior speaking turns.
10. Every time an agent speaks, multiple facets of information are shared.
11. All agents receive and learn shared information with some noise around relevance.
12. Agents track how often information is shared and weight utility accordingly.
13. Low-competence agents react based only on concordance with team preference.
14. High-warmth agents react based on both team and individual preference.
15. All agents' reactions are equally weighted.
16. Aggregate positive team reactions increase agents' voice, and negative decreases it.
17. Voice changes nonlinearly, adjusting less when already high or low.

Chapter 6: Proposed Questions

The conceptual and formal development of SIFT draws theoretical grounding from two key insights in the literature: 1) individuals often avoid sharing unique information or uncertain information in the presence of social evaluation concerns, and 2) perceived voice can influence whether individuals speak and change based on team feedback. A good theory should not only replicate these established patterns but also generate novel insights into underlying behavioral mechanisms. To that end, the SIFT model serves a dual purpose: it allows for validation checks (i.e., does the model replicate known phenomena?), while also supporting theory-building through the generation of testable propositions.

To evaluate the model's validity and generative value (i.e., does the model correspond to the real world; Rand & Rust, 2011), the simulation systematically manipulated informational distribution structure, certainty, and agent-level voice. Two expected findings were identified as markers of successful replication, and six theory-driven propositions were derived to guide future empirical work. These are summarized in Table 3, which organizes the insights into two categories aligned with the central research questions of the model: what information is shared and who shares information.

Surfacing Information: What Teams Choose to Share

The simulation study first explores how the content and valence of shared information differs across team conditions. While previous work has suggested that social norms may bias teams toward sharing positive or mutually known information, this model allows for direct observation of what agents choose to disclose. It

examines how often unique versus common information is contributed, and whether agents exhibit systematic patterns in the valence (positive, neutral, negative) of what they share. By comparing these patterns across voice, certainty, and distribution conditions, the model captures how team structure influences not only the volume, but also the nature of team communication.

During the process of model observation and development, information selection cycles emerged as an informative behavioral signal. The model records how long agents take to decide what and whether to share with their team on each round, and this measure appeared to vary meaningfully across conditions. Although not the primary focus of the initial model design, differences in information selection cycles offer insight into the internal decision-making processes proposed in SIFT and the potential role of cognitive demand in socially constrained environments. As such, the model tracks information selection cycles across certainty and voice conditions, as well as its progression over time. This allows for post-hoc exploration of how internal dynamics like uncertainty or perceived receptivity may shape the pace of interaction and reveal patterns of behavioral adaptation within teams. The model will also investigate factors that reduce information selection cycles, providing insights into how teams can streamline decision-making.

Shaping the Conversation: Who Takes the Floor

The model also investigates how individual voice and team feedback patterns evolve over the course of interaction. It explores whether early voice levels predict who speaks more over time and how team responses influence future behavior. This includes examining the types of feedback agents receive (positive, neutral, negative),

how often they change their voice levels in response, and whether these patterns differ across experimental conditions. By tracking these dynamics over repeated interactions, the model provides insight into how team members calibrate their behavior in response to both social affirmation and ambiguity.

Together, these two areas provide a framework for examining how agents operate under conditions of uncertainty, information distribution asymmetry, and varying voice perceptions. Rather than treating team decision-making as a static endpoint, the model emphasizes the micro-level processes that shape how group decisions unfold over time as a result of information sharing and knowledge emergence within teams.

Chapter 7: Methods

The conceptual and formal representation of SIFT is implemented as a computational model and a simulation study was performed to evaluate the theory's generative sufficiency and predictive validity through virtual experimentation (Epstein, 1999; Kozlowski et al., 2013). The simulation study involves systematically manipulating core parameters related to information distribution, certainty of information, and agents' initial voice perceptions to explore the research questions outlined above. In the present simulations with SIFT, simulated teams of agents iterate through the sharing, learning, and acknowledgement phases outlined in Table 1 for a fixed number of periods ($t_{max} = 500$). Based on preliminary testing, this duration allows the model mechanics to unfold long enough to examine the simulated team dynamics and explore whether the proposed model produced empirically plausible patterns and investigate novel insights.

Simulation Parameters

The computational model for SIFT contains a variety of parameters used to initialize the information and team environment. The parameters for the information environment determine how information is distributed across decision options and people as well as their corresponding attributes (e.g., certainty, relevance). The parameters for the team environment define the attributes of simulated team members as well as the mechanisms used to operationalize their sharing, learning, and acknowledgment behaviors. Table 3 presents the full set of model parameters defined in SIFT and the corresponding parameter values used in the reported simulation study. Where possible, prior empirical research were used to inform the values for

Table 3*Description of SIFT Model Parameters and Values Used in Simulation Study*

Parameter	Description	Value
Information Environment		
<i>hp_condition</i>	Type of hidden profile	<i>Manipulated:</i> [shared, unshared consensus, unshared conflict] ^a
<i>min_certainty</i>	Minimum information certainty ($p_{i_{lower}}$; Eq. 1)	<i>Manipulated:</i> [0, 0.34, 0.68]
<i>max_certainty</i>	Maximum information certainty ($p_{i_{upper}}$; Eq. 1)	<i>Manipulated:</i> [0.33, 0.67, 1]
<i>n_iterations</i>	Total number of information sharing events to simulate	500
<i>num_options</i>	Number of decision options (j)	3 ^a
<i>info_per_option</i>	Information pieces per option (k_j)	16 ^a
<i>pos_min_relevance</i>	Minimum relevance of positive/supportive information	.75
<i>pos_max_relevance</i>	Maximum relevance of positive/supportive information	1
<i>neutral_min_relevance</i>	Minimum relevance of neutral information	-.05
<i>neutral_max_relevance</i>	Maximum relevance of neutral information	.05
<i>neg_min_relevance</i>	Minimum relevance of negative/unsupportive information	-.5
<i>neg_max_relevance</i>	Maximum relevance of negative/unsupportive information	-.25
Team Environment		
<i>num_team</i>	Number of agents (n)	4 ^a
<i>min_voice</i>	Minimum initial agent voice ($voice_{lower}$; Eq. 4)	<i>Manipulated:</i> [-0.8, 0.7]
<i>max_voice</i>	Maximum initial agent voice ($voice_{upper}$; Eq. 4)	<i>Manipulated:</i> [-0.7, 0.8]
<i>min_warmth</i>	Minimum agent warmth perception ($warmth_{lower}$; Eq. 5)	-1
<i>max_warmth</i>	Maximum agent warmth perception ($warmth_{upper}$; Eq. 5)	1
<i>min_competence</i>	Minimum competence value perception ($competence_{lower}$; Eq. 6)	-1
<i>max_competence</i>	Maximum competence value perception ($competence_{upper}$; Eq. 6)	1
<i>alpha</i>	Relevance distortion for positive information (α ; Eq. 9)	0.88 ^b
<i>beta</i>	Relevance distortion for negative information (β ; Eq. 9)	0.88 ^b
<i>lambda</i>	Relevance distortion scaling for negative information (λ ; Eq. 9)	2.25 ^b
<i>gamma</i>	Certainty distortion for positive information (γ ; Eq. 9)	0.61 ^b
<i>delta</i>	Certainty distortion for negative information (δ ; Eq. 9)	0.69 ^b

<i>info_recall_alpha</i>	Incremental probability of recalling information based on load (μ ; Eq. 10)	3.33 ^c
<i>info_recall_beta</i>	Base probability of recalling information (τ ; Eq. 10)	0.373 ^c
<i>encoding_fidelity</i>	Degree of error learning from objective knowledge pool ($encoding_x$; Eq. 8)	0.05
<i>decoding_fidelity</i>	Degree of error learning shared information ($decoding_x$; Eq. 17)	0.10
<i>rel_learning_pct</i>	Learning rate for shared information ($lrn.rate$; Eq. 17)	0.05
<i>max_delta_voice</i>	Maximum voice change following team reaction ($\Delta voice_{max}$; Eq. 20)	0.25
<i>a_param</i>	Scaling parameter for voice change dependence on current voice (ϕ ; Eq. 20)	10
<i>threshold_scale</i>	Scaling parameter for information selection threshold (W_s ; Eq. 16)	3
<i>prior_share_interval</i>	Number of elapsed time points before information is eligible to be shared again	5

^a Parameter settings based on empirical studies by Stasser and Titus (1985)

^b Parameter settings based on empirical estimates reported in Tversky and Kahneman (1992)

^c Parameter settings based on empirical estimates reported in Stasser (1988)

model parameters used in the simulation study. Given the size of the parameter space, it was neither feasible nor beneficial to systematically manipulate all parameters. The simulation design and manipulations were thus specifically chosen to provide insight into the primary questions and model validation targets of interest. The following sections detail the three focal manipulations of the simulation study: (1) information distribution, (2) informational certainty, and (3) agents' initial voice perceptions.

Information distribution. The first manipulation involves how information is distributed across agents and decision options in a team and thus the structure of the information environment. Three information distribution structures reproduced from Stasser and Titus' (1985) paradigm examining the effects of hidden profiles on team decision-making are included in the simulation (Table 4). Across all conditions, a single optimal decision exists with the largest ratio of positive-to-negative information within the full knowledge pool (Option 1). The manipulation of information distribution within the team systematically varies which and how much

Table 4

Number and Type of Information Pieces about Each Option Received by Each Agent Before Discussion

Condition and information valence	Decision Option		
	1	2	3
<i>Shared</i>			
Positive	8	4	4
Neutral	4	8	8
Negative	4	4	4
<i>Unshared/consensus</i>			
Positive	2	4	1
Neutral	4	5	8
Negative	4	1	1
<i>Unshared/conflict</i>			
Positive	2	4 [4]	4 [4]
Neutral	4	6 [4]	4 [6]
Negative	4	0 [2]	2 [0]

Note. In the unshared/conflict condition, two members of a four-person group received configurations of information about Options 2 and 3 given by the numbers without brackets, whereas the other two members received configurations given by the numbers in brackets. Based on Stasser and Titus (1985).

of the knowledge pool individuals initially know and therefore whether their individual decision preferences differ from the optimal choice.

In the *shared* information distribution condition, each agent receives the entire knowledge pool and thus all agents began with the same set of information, providing a baseline where consensus is structurally easy to achieve. The *unshared/consensus* condition introduces informational asymmetry in which half of the knowledge pool is common/initially known by all members and the remaining half is distributed across agent (i.e., each agent knew six pieces of unique information that no other agent knew). However, the distribution of information across members in the unshared/consensus condition is such that all agents should initially prefer the same suboptimal decision (Option 2). By comparison, information in the *unshared/conflict* condition is also distributed such that half of the knowledge pool is shared/initially known by all members. The remaining information is then distributed across pairs of

members such that two members of the team initially prefer one suboptimal decision (Option 2) whereas the other two members initially prefer the other suboptimal decision (Option 3). This setup introduces not only asymmetry, but potential epistemic conflict, allowing for the observation of how disagreement might emerge or be resolved when team members start from different preferences. Taken together, these three conditions enable structured comparison of how information architecture influence information sharing behavior.

Informational certainty. The second manipulation targets the certainty of information. The goal of this manipulation is to explore how teams operating under environments characterized by different levels of epistemic confidence may affect their information sharing activities. As such, three qualitatively distinct information environments were created in which the range of certainty for the available information differed (see Equation 1). In the *high-certainty* condition, the certainty values of all information are randomly sampled between [0.68, 1], creating an environment with information of generally high reliability. In the *moderate* and *low-certainty* conditions, certainty values are sampled from [0.34, 0.67] and [0, 0.33], respectively, progressively increasing the unreliability of available information.

Initial voice. The third manipulation targets agents' initial voice levels to explore how perceptions of team social interactions may affect sharing behavior and social responsiveness. For purposes of the present simulation, two conditions were created to represent teams in which all agents began with similar voice levels (see Equation 4). In the *high-voice* condition, the initial voice level for all agents in the team are randomly initialized between [0.7, 0.8], reflecting groups where all

individuals enter with strong confidence that their inputs to team knowledge-building will be heard. In the *low-voice* condition, agent voice levels are initialized between [-0.8, -0.7], representing teams where all members begin with a sense of limited positive reception from team members.

Simulation Design

The final simulation design employs a fully crossed 3 (information distribution: shared, unshared/consensus, unshared/conflict) x 3 (certainty: high, moderate, low) x 2 (voice: high, low) factorial design, yielding a total of 18 experimental conditions. In addition, a set of “baseline” conditions were constructed to emulate information environments and decision-making teams comparable to those in prior empirical and simulation research (e.g., Grand et al., 2016; Stasser & Titus, 1985). Three additional conditions were thus constructed comprising teams of agents all with moderate initial voice levels (initial voice between [-0.1, 0.1]) operating in an environment of perfect certainty ($p_i = 1$ for all information) for each of the information distribution structures. In total, the simulation includes 2,520 teams (120 teams x 21 conditions), and over 5.04 million individual agent-rounds of interaction (4 agents x 500 rounds), capturing rich longitudinal interaction data and enabling both validation and discovery of emergent team dynamics.

All simulations are developed and executed in R (version 4.3.2) using custom-built agent-based modeling functions. Data processing and output analyses were conducted using the tidyverse, data.table, fs, dplyr, ggplot2, snowfall, relecuyer, reshape, and tictoc packages.

Simulation Measures

Several key outcome variables are recorded during the simulation to evaluate the proposed research questions and to capture emergent team dynamics over time (see Table 5). These measures are selected to represent core aspects of information sharing, information selection cycles, convergence, and social responsiveness in team decision-making contexts.

First, decision preferences are recorded for each agent at every simulation round, indicating which of the three available options the agent currently believes to be the best choice. These preferences are used to calculate team-level majority alignment, defined as the proportion of rounds in which at least three of the four agents share the same decision preference. This measure serves as a proxy for team consensus and is used to examine how information distribution and certainty affected alignment over time.

Second, the model records the content and characteristics of shared information. Each time an agent chooses to speak, the simulation logs which agent spoke, whether the shared information was positive, negative, or neutral in valence, and whether it was initially unique or shared. These data are used to assess how different conditions influenced the informational diversity and evaluative tone of team communication.

Third, the number of information selection cycles each agent engaged in before deciding what and whether to share is recorded at every round. While these steps do not represent real-time seconds, greater selection cycles reflect longer simulated cognitive processing before disclosure and thus capture the degree to which

Table 5
Simulation Outcome Variables and Operational Definitions

Measure	Type	Operational Definition
Agent Decision Preference	Primary Outcome	Categorical record of which option (1, 2, or 3) each agent preferred at each round.
Team Majority Alignment	Primary Outcome	Binary indicator (yes/no) of whether at least three agents shared the same decision preference.
Valence of Shared Information	Primary Outcome	Categorical variable indicating whether shared content was positive, negative, or neutral.
Uniqueness of Information	Primary Outcome	Binary indicator of whether a shared piece of information was unique or shared.
Information Selection Cycles	Process Measure	Number of iterations an agent took before sharing; proxies cognitive demand
Agent Voice Score	Social Dynamics	A continuous value representing an agent's perceived influence on the team at each round.
Voice Change Rate	Social Dynamics	Round-to-round change in voice score; reflects responsiveness to social feedback over time.

agents appear to hesitate or expend effort when deciding whether and what information to contribute. This measure is used to evaluate the effects of certainty and voice on cognitive demand and internal hesitation.

Fourth, the model tracks agent voice levels at every time step. Voice scores represent each agent's perceived influence in the team, and are updated based on social reactions received in response to shared information. In addition to recording absolute voice scores, the model calculates the rate of change in voice over time to capture how quickly or gradually agents stabilize in their perceptions of being heard. These measures are used to assess the social dynamics of responsiveness and feedback integration across team conditions.

Chapter 9: Results

This section reports on the behavioral and communicative patterns observed across simulated team decision-making environments for the SIFT model. Results are organized thematically to align with core research questions: 1) what information is shared and 2) who shares information.

Surfacing Information: What Teams Choose to Share

This section explores the content of team communication: what kinds of information agents choose to share over time. While the model systematically manipulates voice, certainty, and information distribution, the observed differences in what was shared are driven by information distribution. Across conditions, there were no consistent effects of voice or certainty on the type or amount of information shared—suggesting that what teams talked about is shaped more by the architecture of access than by individual-level attitudes or environmental ambiguity. Additionally, neutral information is rarely communicated in any condition, reflecting its low utility for influencing team decisions. As such, the following analyses focus on three meaningful dimensions of content: whether the information is common or unique, whether it is positive or negative in valence, and which decision option it pertained to.

Unique versus Common Information

A central insight from the hidden profile literature is that teams often fail to surface unique information, instead focusing on what is mutually known. Some explanations emphasize social or cognitive biases—such as common/shared information being easier to validate or preferred for building consensus (Stasser, 2003). However, SIFT does not assume agents have existing knowledge of whether a

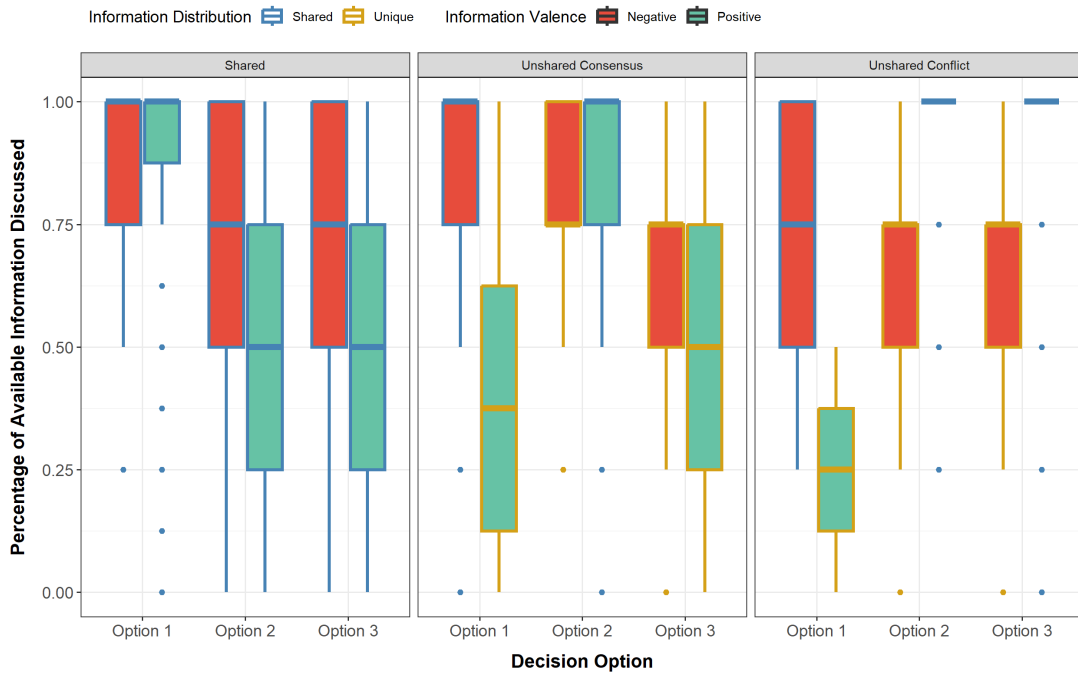
piece of information is common/shared or unique prior to group discussion nor does SIFT include any mechanisms by which agents can acknowledge whether they already knew a piece of information once it is communicated. As a result, sharing patterns in SIFT do not reflect an attentional or motivational preference for communicating commonly known/shared information per se, but rather emerge from agents' perceptions of their information (i.e., utility of information, personal decision preference) and team environments (i.e., voice, decision option believed to be preferred by the team).

Figure 4 replicates a well-established pattern from hidden profile research: across conditions, agents share more shared than unique information. This aligns with prior work showing that commonly held information is more likely to be discussed in group settings (e.g., Stasser & Titus, 1985). However, this volume-based finding does not tell the full story of how teams interact with the information available to them.

Figure 5 provides a more dynamic view by illustrating the likelihood that teams are discussing shared versus unique information at each point in time. Surprisingly, despite the greater volume of shared information ultimately communicated (Figure 4), the probability of a team discussing unique information is roughly equal to—or in some cases greater than—that of discussing shared information throughout the simulation. For example, in the unshared/consensus condition, a higher percentage of teams are discussing unique information at every time point. Even in unshared/conflict conditions, shared and unique content receive similar airtime by the midpoint of the task.

Figure 4.


Mean Percentage of Available Information for Decision Options Shared within Teams by Option, Valence, and Initial Distribution



Note. Data averaged across teams within condition. Results only include information for each option (see Table 3) that was communicated by an agent to others in the team.

Figure 5.

Percentage of Teams Discussing Shared Versus Unique Information at Each Time Point



Note. Vertical axis reflects the percentage of teams which communicated a shared or unique piece information at each simulation round. Figure plots LOESS curves estimated from simulated data within each information distribution condition.

This suggests that the disparity in total unique information shared (Figure 4) is not due to lower attention, salience, or attraction to unique content. Instead, teams are often engaged with unique information, but revisit the same pieces multiple times, limiting the diversity of what ultimately makes it into the conversation. In other words, it is not an attention problem—it is a breadth problem: teams aren't avoiding unique information, they are stuck on the same pieces of information.

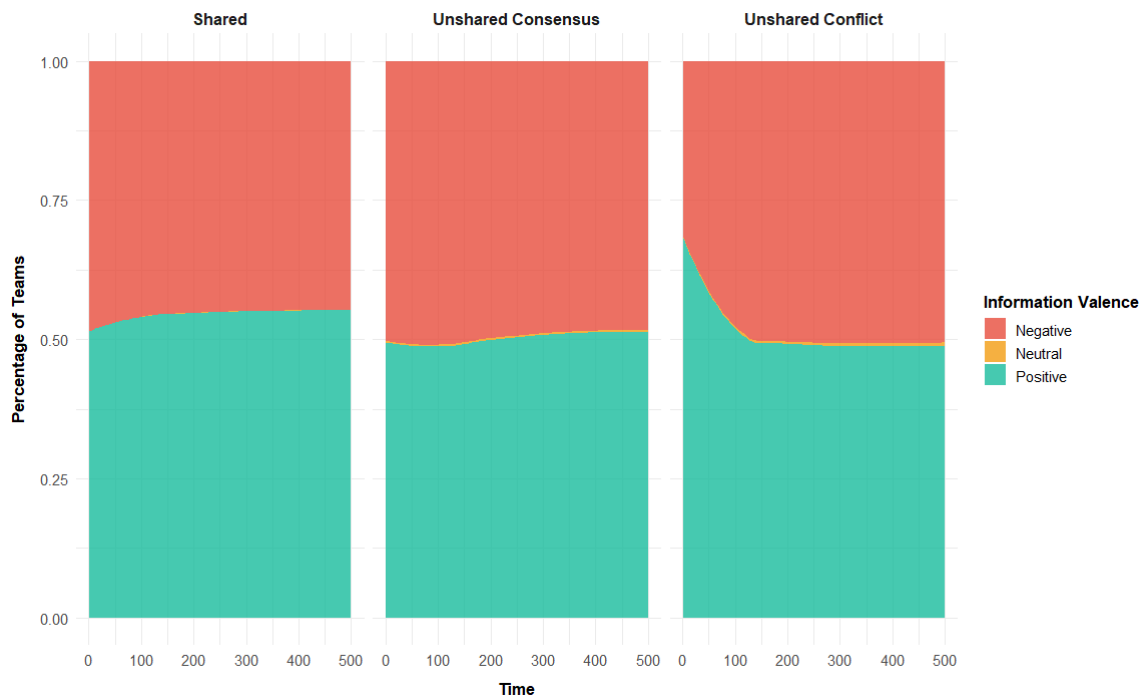
Valence of Shared Information

Given that more shared information is communicated than unique (see Figure 4), one possibility is that this disparity reflects a bias toward sharing positive information suggested in prior research, where individuals selectively disclose content that aligns with team preferences or feels socially safe (e.g., Wittenbaum et al., 2004). If agents favor sharing positive over negative information, and shared content is more often positive in valence, this could help explain the imbalance in what is ultimately communicated.

However, Figure 4 also complicates this explanation. While it may initially appear that more negative information is shared than positive—particularly among shared content—the effect is not consistent across conditions, nor is it large enough to suggest a strong valence bias. Instead, what Figure 5 reveals is that both positive and negative information are communicated at similar rates, and any imbalance in valence appears more reflective of availability within each condition than of a systematic preference for positivity or negativity.

This conclusion is reinforced by Figure 6 which shows the proportion of teams discussing each valence type across time. Across all three distribution

Figure 6.
Information Valence Discussed over Time



Note. Vertical axis reflects the percentage of teams which communicated a positive, neutral, or negative piece information at each simulation round. Figure plots LOESS curves estimated from simulated data within each information distribution condition.

conditions, teams are just as likely to discuss negative information as they are to share positive information. This suggests that agents do not strongly discriminate based on valence; what got shared has less to do with its evaluative tone and more to do with its perceived utility or structural availability.

In sum, while prior literature has proposed that social dynamics might drive individuals to prefer affirming content, these results suggest that valence is not a primary filter for information selection in SIFT. This leaves open the question of what does influence which specific pieces of information are prioritized. The next section examines whether the decision option associated with the information played a central role in shaping team discussion patterns.

Decision Option

The final lens through which to examine what teams shared involves the decision option the information supported or opposed. While neither information type or valence fully explains what information is communicated, it is possible that team discussions are influenced by the option that information is about—and that differences in discussion content across options may explain why some information surfaced while other content remain unshared.

As shown in Figure 4, some decision options receive more overall contributions than others, though this varies by condition. These differences become clearer when interpreted alongside Figure 5 (which shows the likelihood of discussing unique versus shared information over time) and Figure 7 (which shows the likelihood a team is discussing a given option).

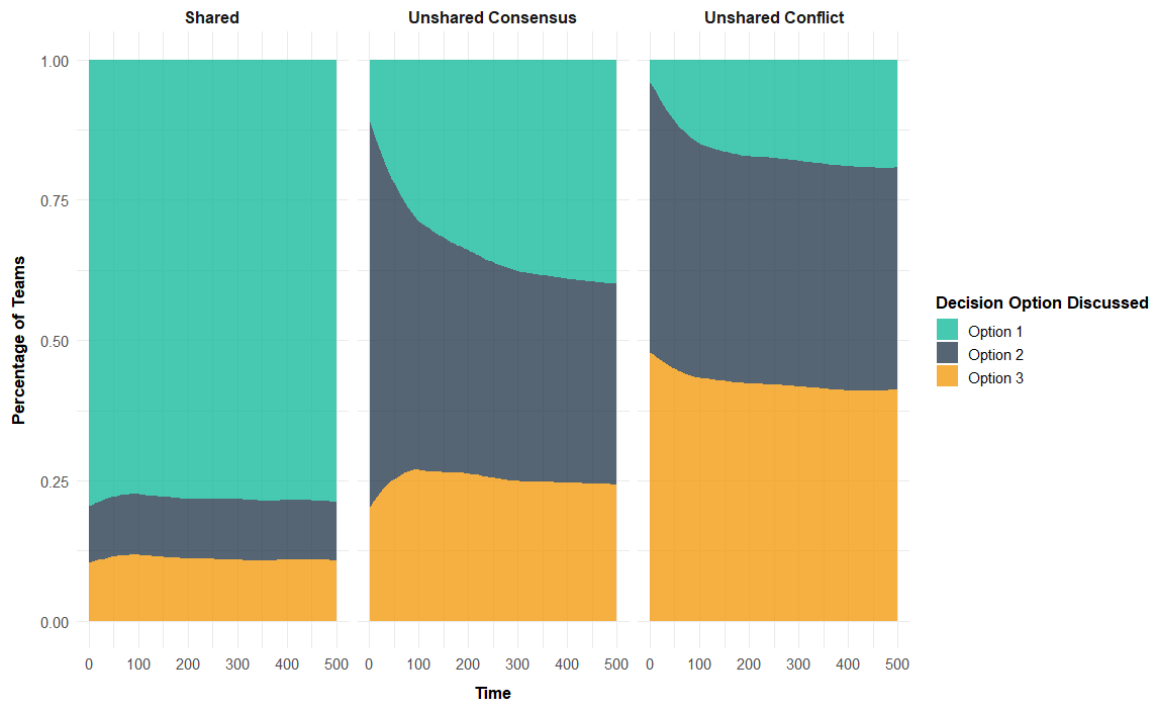
In the *unshared/consensus* condition, teams begin the simulation with information aligned around Option 2, leading all agents to initially prefer it based on their private information. However, the majority of information available about Option 2 in this condition is both unique and negatively valenced. Once surfaced, this content has the potential to undermine confidence in the initial preference. Figure 4 illustrates this pattern: for Option 2, most of the unique information shared is negative, while very little is shared or positive. At the same time, Figure 5 shows how these teams consistently discuss unique information throughout the simulation, and Figure 7 reveals that discussion of Option 2 steadily declines. These trends suggest that the content being shared—negative unique information about Option 2—prompts teams to reevaluate their initial choices.

As attention to Option 2 decreases, the probability teams are discussing Option 1 increased. This option is associated with a greater proportion of positive unique information (Figure 4), and as teams engage with this content (Figure 5), it gains traction. Over time, as shown in Figure 7, discussions of Option 1 and Option 2 became more balanced. This gradual shift implies that Option 1 is drawing more attention as teams accumulate new insights—but because there is more unique information to surface, the shift takes longer to unfold. Rather than signaling avoidance of unique information, this pattern reveals how the structural difficulty of surfacing unshared content shapes the pace and focus of team communication.

In the *unshared/conflict* condition, teams start with more divergent preferences: two agents favor Option 2 and two agents Option 3. Both Options 2 and 3 are tied to unique information that is largely negative, which means that as these options are discussed, their weaknesses become more apparent. Figure 5 shows that teams spend time discussing both shared and unique information over time, and Figure 7 captures how attention shifts away from Option 1 early on. This pattern likely reflects that, after surfacing enough of the negative information tied to Option 2 and 3, some teams begin revisiting Option 1 as a viable alternative.

However, unlike the *unshared/consensus* condition—where Option 1 only has to overtake one competitor—Option 1 in the *unshared/conflict* condition must compete against two partially attractive alternatives. This creates a more complex evaluative landscape in which teams fluctuated between Options 2 and 3, requiring Option 1 to overcome two distinct informational hurdles. As a result, attention in

Figure 7.
Percentage of Teams Discussing Decision Options over Time



Note. Vertical axis reflects the percentage of teams which communicated information most consequential to decision option 1, 2, or 3 at each simulation round. Figure plots LOESS curves estimated from simulated data within each information distribution condition.

conflict teams are more fragmented around Option 2 and 3, even though many teams still return to reevaluate Option 1 overtime.

Together, the findings from this section challenge common assumptions about what drives communication patterns in teams. While SIFT replicates the broad empirical pattern that common/shared information tends to be communicated more frequently than unique information, this cannot be attributed to motivational preference or selective attention. Instead, the model demonstrates that teams frequently attend to unique content but face structural constraints in surfacing it—especially when such information is unevenly distributed or cognitively effortful to locate. Similarly, the valence of information does not appear to systematically bias sharing decisions, and teams show no clear preference for positive over negative

content. Instead, the topic of discussion (i.e., the decision option being evaluated) emerges as a more proximal influence on communication dynamics. The interplay between information distributions, information valence, and team focus shape what is repeatedly shared, ignored, and gradually gained traction over time. These results suggest that the patterns of what gets shared are not just a product of cognitive or social preference, but of the deeper architecture of information and the evolving nature of team-level attention.

Shaping the Conversation: Who Takes the Floor

While the previous section focuses on the content of team communication, this section turns to the agents themselves—specifically, how patterns of participation and voice evolve over time. Understanding who shares information is essential to explaining not just what gets said, but how team knowledge emerges and is distributed. The model allows us to observe how individual differences in voice, feedback, and traits such as competence and warmth contribute to enduring participation patterns and inequalities.

Voice and Speaking Patterns

Participation in group discussions is frequently uneven, with a smaller number of members often dominating the conversation. This phenomenon has been well documented in both early and contemporary group research. For instance, Stephan and Mishler (1952) found that a minority of participants tend to account for most of the speaking time. Similarly, Dabbs and Ruback (1987) demonstrated that those who speak early are more likely to continue speaking, reinforcing their conversational

presence. These patterns tend to hold even in small groups (Parker, 1988) and tend to lead to disparities in information sharing (Stasser & Taylor, 1991).

The present section examines how similar dynamics unfold in the SIFT model, where participation is shaped not only by initial voice perceptions but by the accumulation of feedback and perceived receptivity over time. It explores how internal perceptions of voice and early speaking behavior interact to produce entrenched communication hierarchies.

These dynamics manifest clearly in the simulated teams: the knowledge-building process in SIFT results in marked asymmetries in speaking behavior across agents (Figure 8). This is especially noticeable in low-voice environments, where agents begin the simulation with weaker expectations of being heard. In these conditions, an agent that emerges early as a dominant speaker tends to grow and maintain the status as the highest volume speaker in the team over time. Importantly, this agent does not engage in more information sharing because it has better or more useful information, but rather as a result of an emergent self-reinforcing feedback loop between early participation, feedback, and voice.


The mechanics of this process can be gleaned from Figure 9, which demonstrates how agent voice tends to change over time based on their relative speaking frequency. In sum, Figure 9 demonstrates that an agent who speaks earlier tends to experience rapid increases in voice relative to their teammates in environments in which all agents begin with low voice levels. Consistent with the information selection process in SIFT (see Step 13 in Table 1 and Equation 14), agents with higher levels of voice tend to require fewer deliberations to select a piece

of information they wish to share, thus making it more likely they will speak relative to agents with lower levels of voice. Consequently, speaking raises the voice of an agent, which increases the likelihood of future speaking. However, the zero-sum nature of speaking turns in team discussion means that other agents have less opportunity to speak and improve their voice, thereby resulting in sharply unequal participation rates – a dynamic that mirrors real-world reinforcement processes in group influence (Parker et al., 2010).

While the reinforcement cycle is most pronounced in low-voice environments, voice trajectories in high-voice conditions follow a markedly different pattern. When all agents begin with strong expectations of being heard, early differences in participation are less likely to cascade into dominance. Instead, voice levels remain more stable and closely cluster across agents throughout the simulation, resulting in more equitable sharing patterns. These environments facilitate faster initial speaking and reduce the likelihood that a single agent monopolizes feedback and attention.

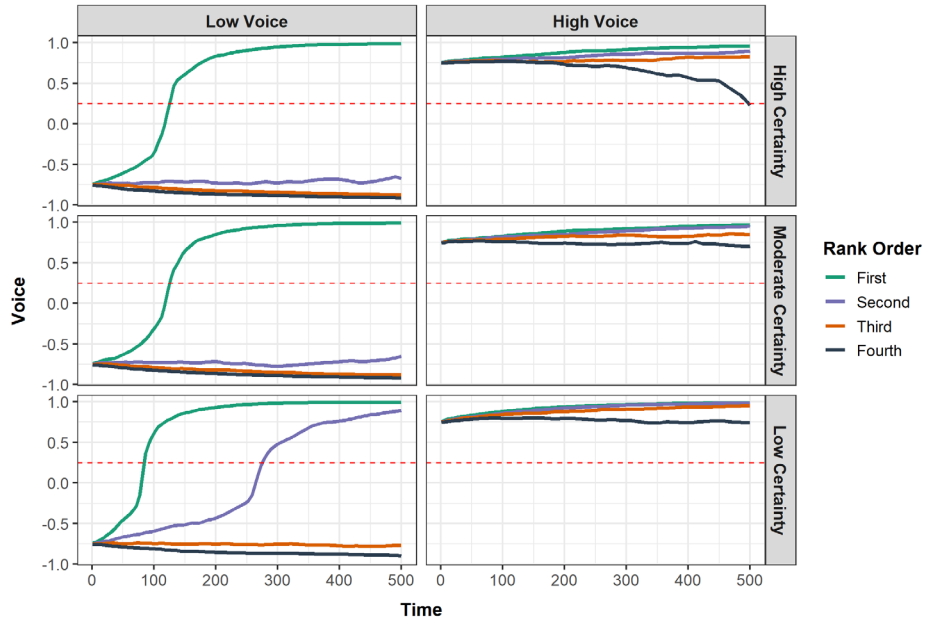
One exception worth noting appears to be in the low-voice, low-certainty condition in Figure 9, where voice growth trajectories diverge more dramatically. This interaction highlights how hesitation can be compounded when agents are both unsure how their contributions will be received (low voice) and unsure whether their information is reliable (low certainty). In such environments, some agents delay sharing for many rounds, while others—after receiving even minimal feedback—begin speaking repeatedly. These dynamics amplify early gaps in perceived receptivity and create particularly strong disparities in who ends up shaping the discussion.

Figure 8.
Distribution of Speaking Participation within Teams



Note. Rank order was determined by summing the cumulative number of times agents had spoken up to each time point (First = agent who had spoken the most by a given time point; Fourth = agent who had spoken the least by a given time point). Lines thus represent the median percentage of cumulative number of speaking turns by an agent of the given rank order at a given time point across all teams.

Figure 9.
Changes in Agent Voice over Time by Agent Speaking Frequency



Note. Rank order was determined by summing the cumulative number of times agents had spoken up to each time point (First = agent who had spoken the most by a given time point; Fourth = agent who

had spoken the least by a given time point). Lines thus represent the median percentage of cumulative number of speaking turns by an agent of the given rank order at a given time point across all teams.

Reactions and Feedback

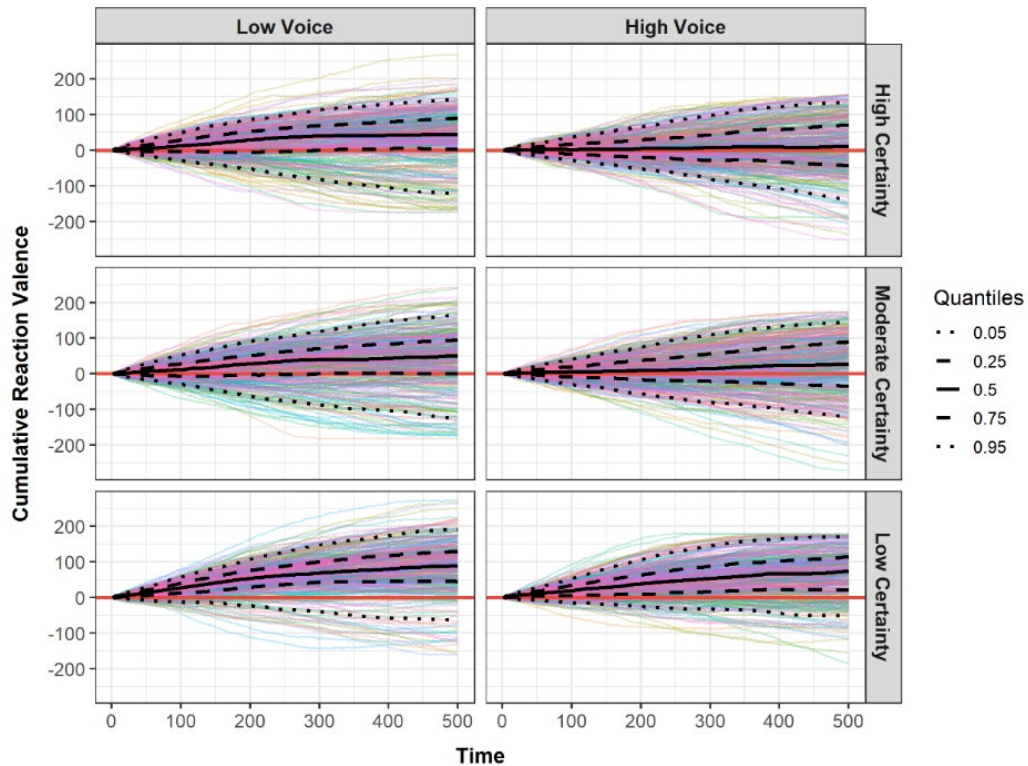
In SIFT, agents that share information receive feedback in the form of reactions from their teammates. These reactions—positive, neutral, or negative—shape how an agent’s perceived voice may change. Crucially, agents update their voice not simply based on whether they contributed but based on how others responded to their contribution. This design allows for social reinforcement or discouragement to emerge over time and influence future participation behavior.

Figure 10 plots the cumulative reaction valence across teams, offering a longitudinal view of the feedback in each team. While early values are necessarily smaller due to limited number of interactions, the figure shows that by later rounds, reaction trajectories vary meaningfully across teams. Some teams exhibit predominantly positive cumulative feedback, while others trend toward neutral or even negative. Importantly, the median reaction valence is greater than zero, indicating that more than half of the teams experienced a net-positive feedback pattern over time.

This suggests that, although feedback is not uniformly affirming, many teams are relatively aligned in their perceptions of the team’s preferred decision, leading to a higher rate of positive responses to shared information. Because agents can only react positively when a contribution supports the team preference (and individual preference for agents with high competence perceptions), this feedback pattern reflects a broader convergence in perceived alignment, rather than simply a socially generous environment. Over time, this alignment contributed to reinforcement loops:

Figure 10.

Cumulative Sum of Reactions to Shared Information by Teams over Time



Note. Lines indicate the cumulative sum of team reactions to shared information across time (i.e., values greater than zero reflect instances where teams have expressed more positive than negative feedback to shared information). Each colored line corresponds to the cumulative reaction valence for a single simulated team. The red horizontal line denotes a net reaction valence of zero.

agents who spoke early often receive positive reactions, which increase their voice and make them more likely to speak again, deepening the participation disparities observed in earlier analyses.

Agent Traits and Early Speakers


Finally, how agent-level traits—specifically, competence and warmth—shape early participation behavior and patterns of dominance are examined. These traits influence how readily agents choose to contribute, though in distinct ways. Warmth acts as an internal moderator regulating the information selection process in SIFT (see Step 15 in Table 4 and Equation 16); specifically, agents who care more about being perceived as warm require a higher level of expected utility before they are willing to

speak. This reflects a motivation to avoid sharing socially disruptive or poorly received contributions.

As shown in Figure 11, agents who contribute less frequently tend to hold higher warmth, suggesting that self-image concerns play a role in participation restraint. This pattern holds across voice and certainty conditions, with no clear evidence that the effect of warmth varies substantially by environment. Together, these findings suggest that internal social motivations for self-image— independent of structural constraints— can suppress early participation, contributing to longer-term silence in team discussions.

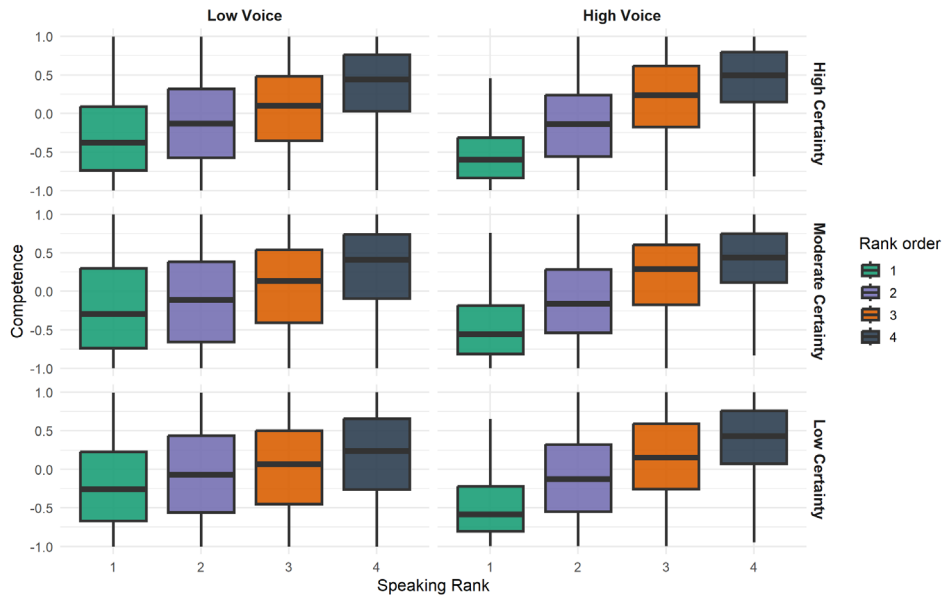
Agent competence also affects the likelihood of engaging in information sharing such that lower competence agents also tend to participate more in information sharing (Figure 12). In contrast to warmth, though, competence does not influence speaking participation rates by affecting the internal threshold for speaking but rather by affecting how many pieces of information agents are likely to consider sharing. In SIFT, the extent to which an agent is likely to consider sharing a piece of information is based on their perception of how consequential/informative a piece of information is perceived to be to the agent's preferred decision (see Step 9 in Table 4 and Equation 12). However, agents with higher competence are less sensitive to these differences in the expected utility of information and thus are more willing to consider sharing information that may not be as directly informative to their own preference (see Figure 2). In contrast, lower competence agents are more sensitive to differences in the expected utility of information and thus tend to only consider

Figure 11.
Warmth Self-Perceptions of Agents Based on Speaking Frequency



Note. Rank order was determined by identifying the agent with the highest cumulative speaking frequency at the end of a model run (First = agent who had spoken the most by end of team discussion; Fourth = agent who had spoken the least by end of team discussion).

Figure 12.
Competence Self-Perceptions of Agents Based on Speaking Frequency



Note. Rank order was determined by identifying the agent with the highest cumulative speaking frequency at the end of a model run (First = agent who had spoken the most by end of team discussion; Fourth = agent who had spoken the least by end of team discussion).

sharing information that is closely related to its own preference. This does not reflect a motivational shift toward team-oriented behavior per se, but rather a cognitive tendency to take on “informational risk” in pursuit of a potentially more informative contributions.

Taken together, SIFT represents two distinct pathways through which individual differences shape participation in team discussions. Warmth regulates the internal threshold for sharing, making high-warmth individuals more hesitant to contribute unless their information is seen as strongly beneficial or socially safe. In contrast, competence affects the breadth of information sampling, with lower-competence individuals more likely to focus narrowly on high-utility items, leading them to find shareable information more quickly. These dynamics help explain how individuals with lower warmth and lower competence—those who may be perceived as less affiliative or insightful—can nonetheless come to dominate team conversations. Not because they offer better content, but because they face fewer internal and cognitive barriers to speaking early and often.

Information Selection Cycles

While individual traits such as warmth and competence shape the likelihood of participation, contextual factors often determine how easy or difficult it is to act on that inclination. In SIFT, the decision to speak is inextricably tied to the information selection process—an internal deliberation in which agents evaluate the utility of their knowledge, anticipate team preferences, and weigh the risks of disclosure. How long this process takes depends not just on who the agent is, but on the environment they

are operating in: uncertainty, perceived voice, and information structure all modulate the cognitive demand involved in selecting what to share.


These information selection cycles—the number of deliberations before a decision to speak—offer a behavioral window into the cognitive demand required for contribution. Drawing from cognitive load theory (Sweller, 2011), longer cycles reflect increased mental effort, especially under ambiguous or socially constrained conditions. While much of the literature on team cognition emphasizes outcomes like consensus or accuracy, this measure allows us to examine a subtler question: how hard is it to participate, and how does that difficulty shift based on context?

The next section explores how selection cycles vary across conditions of certainty and perceived voice, providing insight into how the environment shapes—not just who shares—but how effortful sharing becomes. Information distribution does not meaningfully impact selection cycles and is therefore excluded from the analyses presented below.

As shown in Figure 13, the number of information selection cycles differ meaningfully by certainty condition across the entire simulation. Agents operating in high-certainty environments consistently deliberate the least, followed by those in moderate and then low-certainty conditions. These differences are evident from the earliest rounds and remain stable over time, suggesting a persistent effect of informational clarity on internal decision-making effort. Environments with uncertain information impose higher intrinsic processing demands on agents, evidenced by more information selection cycles when identifying a piece of information worth

Figure 13.

Information Selection Cycles across Certainty Conditions over Time




Note. Lines represent the median number of selection cycles that agents in each team required before sharing information at each round. Results are grouped by certainty condition and averaged across teams.

sharing. Conversely, when information is more certain, agents can more quickly evaluate content and reach the internal utility threshold required for contribution.

In addition to informational clarity, perceived voice also plays a central role in shaping how many information selection cycles agents accumulate before contributing. As shown in Figure 14, agents in low-voice conditions consistently take longer to select information to share compared to those in high-voice conditions, and this difference persists across all rounds of the simulation. Although deliberation times decrease early in the interactions as agents acclimate to the team, a stable gap remains between voice conditions throughout. This suggests that low voice environments impose a persistent psychological cost on participation. Because agents in these settings begin with lower expectations that their contributions are valued,

Figure 14.
Information Selection Cycles across Voice Conditions over Time



Note. Lines show the median number of selection cycles by team per round under high-voice and low-voice conditions. Higher values reflect greater internal hesitation or cognitive demand prior to contributing.

they must search longer for information that crosses their internal utility threshold.

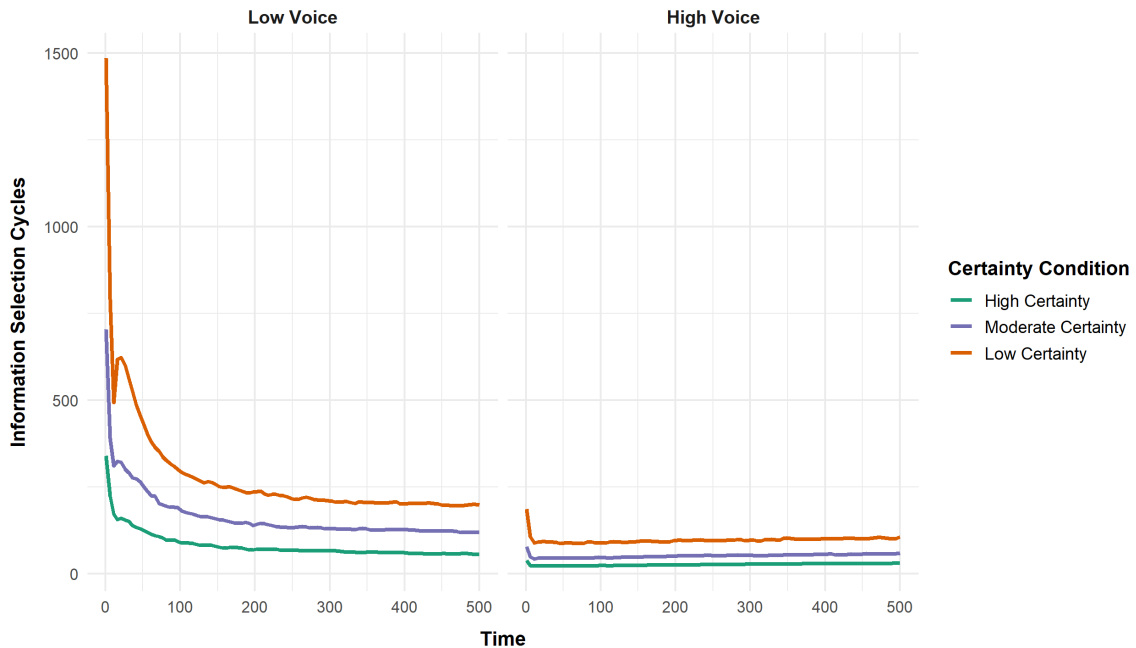
The result is slower and more effortful engagement, even when relevant knowledge is readily available.

When examined jointly, voice and certainty together shape the cognitive demands of information selection. The highest levels of deliberation occurs in the low-voice, low-certainty condition (Figure 15), while the lowest levels are observed in the high-voice, high-certainty condition. These effects persist throughout the simulation and suggest that both uncertainty and low perceived voice impose cumulative effort costs that slow down the decision to contribute. Although deliberation times generally decline as agents adjust to the environment and team, differences across conditions remain stable, indicating these contextual features exert

independent and additive effects on internal decision-making. There is also some visual evidence of an

Figure 15.

Information Selection Cycles across Voice and Certainty Conditions over Time



Note. Lines depict the median number of information selection cycles by round across combined voice and certainty conditions. Teams in the low-voice, low-certainty condition exhibited the longest deliberation times.

interaction: the gap between certainty levels is more pronounced in low-voice conditions, suggesting that uncertainty may impose a heavier burden when agents already lack confidence in their social standing. While exploratory, this pattern points to the value of considering how social and informational constraints may jointly affect the cognitive accessibility of team participation.

Together, these patterns suggest that information selection cycles capture an important but often overlooked aspect of team cognition: the effort required to participate. Across all rounds, both informational certainty and perceived voice systematically shape how long agents deliberate before sharing, indicating that internal decision-making is highly sensitive to the social and informational

environment. These findings extend the model's utility beyond identifying who contributed and what is shared—they highlight how hard it is to contribute in the first place, and how this difficulty evolves as the team interacts. By treating information selection cycles as a behavioral trace of cognitive demand, this section provides a complementary lens for understanding participation dynamics—one that emphasizes not just the visibility of contributions, but the underlying friction that can delay or discourage them.

Chapter 10: Discussion

While prior work illuminates many static predictors of team knowledge-building, this simulation aims to explore the evolving patterns that emerge when individuals make repeated decisions about whether—and what—to share with their team. By formalizing key psychological mechanisms—information distribution, information certainty, and voice perceptions, the model is designed to reproduce well-documented information sharing behaviors and generate new insights into how communication inequalities and information diversity unfold over time (Epstein, 1999; Kozlowski & Ilgen, 2006; Vancouver & Weinhardt, 2012). In this way, the simulation serves both a confirmatory and generative role, offering a testbed for classic findings while surfacing new explanations and trajectories for empirical research. Table 6 summarizes the key replicated findings and new propositions.

Alignment with Prior Research

The model reproduces several foundational findings from the team knowledge literature. First, consistent with decades of research on the hidden profile paradigm (e.g., Stasser & Titus, 1985; Lu et al., 2012), the model replicates the robust finding that teams tend to communicate more common information than unique information, particularly under asymmetrical information distributions. These teams also engage in broader, less focused discussion as a result. The model extends these findings by showing that unique information is not neglected due to attention or preference (Wittenbaum et al., 2004), but because it is less accessible and less frequently held, leading to reduced diversity in what got surfaced—even when agents attempt to share it.

Table 6
SIFT Model Expected Findings and Propositions

Expected Findings
1. Agents share more pieces of common than unique information.
2. Positive team reactions increase agent voice and participation rates.

New Propositions
<i>Dynamic Attention to Information</i>
1. Teams are equally, if not more, likely to discuss unique or common information.
2. Teams are equally likely to discuss positive or negative information.
3. Time is crucial in allowing for attention shifts on decision options to occur.
<i>Social and Contextual Influences</i>
4. Degree of disparity in participation is heavily influenced by voice.
5. Competence and warmth operate on two distinct pathways.
6. External environments and social contexts impact the level of cognitive demand.

Second, the role of voice as a driver of participation is strongly supported (Morrison, 2023; Sherf et al., 2021). Especially in low-voice conditions, disparities in early speaking behavior snowball into lasting asymmetries. Agents who speak early receive more positive reactions from their teammates, which subsequently increases their perceived voice and leads to a higher likelihood of speaking again in the future. This feedback loop may result in large within-team variability in speaking participation among agents that is not necessarily due to competence or insight, but due to social position reinforced by cumulative feedback (Fast et al., 2012).

Third, voice and certainty interact to shape information selection cycles, with lower voice and lower certainty conditions leading to more prolonged deliberation before speaking. This supports prior findings around social hesitation and cognitive effort in ambiguous or socially constrained environments (e.g., Edmondson, 1999; Fast et al., 2014) and offers a dynamic view into how these factors stabilize or amplify over time.

Finally, the model provides direct visibility into reaction mechanisms that are hypothesized but rarely tracked in real-time (Morrison, 2023; Tangirala &

Ramanujam, 2008). Positive reactions early in team interaction are strongly associated with faster voice growth, reinforcing patterns of dominance. Teams with broadly positive feedback patterns experience greater divergence in voice levels over time, suggesting that even encouraging social climates can unintentionally widen influence gaps if early contributors are disproportionately reinforced.

Generative Insights and Predictions for Future Research

Beyond reproducing known effects, the simulation offers several testable predictions that could guide future empirical work. One novel finding is the emergence of dominant speakers with low competence traits—agents who contribute disproportionately despite limited strategic judgement. This pattern appears to result from the model’s decision threshold and sampling structure: lower-competence agents speak earlier and more often because they sample narrowly and reach the expected utility threshold quickly. Over time, this leads to a self-reinforcing voice advantage, even in the absence of informational superiority.

This suggests that early participation, not informational quality, may drive influence in team discussions—particularly under low-voice conditions. Future research might examine whether this occurs in real teams by tracking who speaks first and how voice accumulates over time. Similarly, the model implies that high-competence individuals may be slower to speak not due to lack of insight, but because their sampling strategies delay threshold crossing. In settings that reward rapid input over reflective contribution, this could lead to systematic underutilization of highly competent but cautious voices.

The model also offers important predictions about the temporal trajectory of team interaction, particularly around how information selection cycles and voice dynamics evolve across repeated rounds. Specifically, the model suggests that both information selection time and voice change tend to stabilize over time for most teams, particularly after early rounds of interaction. However, this trend is not universal—some teams exhibit increased deliberation later in the simulation, highlighting that stabilization is a common but not guaranteed outcome of team dynamics in the model. If this dynamic holds in real-world settings, it would imply that much of team’s communicative trajectory is shaped very early on, with later rounds involving less exploration, less adjustment, fewer shifts in voice, and less cognitive demand. Researchers could empirically test this by tracking volatility in team behavior over time, and by examining whether neutral feedback correlates with declines in agent responsiveness, idea generation or participation diversity. Understanding these inflection points in team development could provide critical insight into when interventions—such as feedback training, facilitation prompts, or role-switching—are most likely to be effective (DeChurch & Mesmer-Magnus, 2010; Kozlowski & Ilgen, 2006; Edmondson, 1999; Bunderson & Reagans, 2011).

Finally, the simulation shows that team structure often outweighs psychological intervention (Stasser & Stewart, 1992; Wittenbaum et al., 2004). Voice and certainty manipulations matter—but they could not compensate for structural asymmetries in information access. Shared-information teams are better able to maintain balanced participation and diverse input, even in low-voice or low-certainty environments. By contrast, teams with fragmented knowledge bases struggle to

converge, regardless of confidence or perceived receptivity. This reinforces the idea that team behavior is often constrained by what information is available, not just who holds it or how empowered they feel.

Implications for Team Design and Management

These findings carry several implications for researchers and practitioners. First, the model underscores the importance of early interaction dynamics in shaping long-term participation. Facilitators or leaders may wish to monitor who speaks early and whether certain individuals are dominating conversation—not because they are wrong to contribute, but because they may crowd out alternative voices before other agents reach their decision to share. Rotational speaking turns, or feedback calibration systems, may help counteract this early lock-in.

Second, in uncertain or fragmented environments, the model suggests that providing shared access to high-quality information is often more effective for promoting team convergence than interventions aimed at individual confidence or communication skill. When agents operate from a common knowledge base, coordination becomes more efficient, misunderstandings are reduced, and alignment emerges more naturally. In contrast, even well-intentioned, confident team members struggle to find common ground when key pieces of information are distributed asymmetrically. While promoting psychological safety or encouraging voice remains valuable, these efforts may not overcome the deeper cognitive friction caused by structural information gaps. The model thus supports a simple insight: if teams are expected to make timely and accurate decisions, structural clarity should be prioritized over individual adjustment.

Finally, the model illustrates how dysfunctional team behaviors—like dominance by one member or widespread disengagement—can arise not from flawed individuals, but from the cumulative effects of individually rational decisions. Agents follow simple rules: speak when utility is high, respond to feedback, and adjust voice accordingly. But when one agent speaks early and receives positive feedback, their voice snowballs, while others facing higher thresholds or less frequent reactions, contribute less. Over time, these asymmetries compound into familiar patterns of inequality. This challenges the tendency to attribute group failures to traits like arrogance or passivity and instead points to emergent system dynamics as key drivers of behavior. Recognizing these dynamics opens the door to interventions that target interaction structure, not just individual skill or intent, and highlights the utility of simulation models for revealing these underlying processes.

Model Boundaries

As with any simulation, this model operates under simplified assumptions that limit the scope of its application. Agents are treated as boundedly rational, utility-driven decision-makers who interact repeatedly in close task-focused teams. The model does not capture emotional contagion, strategic deception, power hierarchies, or the nuanced social dynamics often present in real-world teams. As such, its predictions are likely best suited to analytical or problem-solving teams where information quality, decision convergence, and feedback loops are central—such as project teams in engineering, scientific collaboration, healthcare diagnostics, or intelligence analysis.

The model is less applicable to creative, brainstorming, or conflict-laden teams, where ambiguity may be valued, norms are fluid, or alignment is not the primary goal. Likewise, teams with strong formal leadership, status dynamics, or distributed accountability may exhibit very different communication patterns than those modeled here. Finally, while the model incorporates interpersonal traits like voice, warmth, and competence, these are operationalized in simplified and abstract terms; in real teams, such traits interact with context, history, and identity in ways that go beyond what a controlled simulation can fully represent.

Despite these limitations, the model offers a valuable scaffold for exploring how information, perception, and participation interact over time. By articulating specific, testable mechanisms, it provides a bridge between theory and application—and a starting point for designing better team environments where communication and coordination are not left to chance.

Conclusion

This thesis develops and tests a computational model of team interaction, SIFT, to investigate how individuals decide whether and what to share with their teammates over time. Rather than treating communication as a static act, the model focuses on the emergence of knowledge through interaction—how individuals adapt their contributions based on the structure of information, perceived receptivity, and the responses they receive from others.

Several key insights emerged. First, information distribution structure plays a central role in shaping what was shared. Teams with shared access to information have broader and more diverse content contributions, while teams in unshared conditions, especially conflict, tend to repeat the same unique information across rounds. Importantly, the model reveals that agents are not intentionally avoiding unique information—they simply encounter common information more often due to structural access. Thus, what appears as a redundancy bias is not attentional, but a byproduct of availability.

Second, voice shapes not just who spoke, but how participation patterns evolved. Teams with higher initial voice levels show more balanced speaking distributions, while low-voice teams often demonstrate one agent dominate early and persistently. However, voice effects are deeply intertwined with feedback loops: once an agent begins speaking and receives positive reactions, their perceived voice—and participation—grow rapidly, while others fall behind. This underscores the importance of early-stage dynamics in shaping downstream information flows and group structure.

Third, the model reveals that voice and certainty are the key drivers of information selection cycles and cognitive demand over time. Agents with lower perceived voice deliberate longer before choosing to speak, reflecting greater internal hesitation about whether their contributions are valued. In early rounds, when voice levels are still fluid and agents are learning how their inputs are received, information selection cycles are highest. As voice begins to stabilize, particularly for dominant speakers who quickly accumulate positive feedback, selection cycles drop off sharply. This pattern suggests that cognitive demand is concentrated during early periods of social ambiguity, when agents are actively monitoring reactions and calibrating their willingness to contribute. Over time, as voice hierarchies solidify, speaking becomes more habitual and less cognitively demanding for some agents, while others may disengage. This has important implications for understanding when teams are most reflective, hesitant, or open to influence—and reinforces the idea that voice dynamics shape not just participation, but also cognitive labor of decision-making.

As teams continue to operate in increasingly complex, uncertain environments—from healthcare to policy to virtual collaboration, the ability to engineer environments for knowledge emergence will be critical. This model offers a starting point for that design: not by prescribing rigid structures, but by revealing the conditions under which ideas are heard, beliefs evolve, and collective intelligence takes form.

References

- Abele, A. E., & Wojciszke, B. (2007). Agency and communion from the perspective of self versus others. *Journal of Personality and Social Psychology*, *93*, 751–763. <https://doi.org/10.1037/0022-3514.93.5.751>
- Argote, L., Gruenfeld, D., & Naquin, C. (2001). Group learning in organizations. In M. E. Turner (Ed.), *Groups at work: Advances in theory and research*: 369–411. Hillsdale, NJ: Lawrence Erlbaum.
<https://doi.org/10.4324/9781315805986-17>
- Aryee, S., Walumbwa, F. O., Mondejar, R., & Chu, C. W. (2017). Core self-evaluations and employee voice behavior: Test of a dual-motivational pathway. *Journal of management*, *43*(3), 946-966.
<https://doi.org/10.1177/0149206314546192>
- Bandura, A., & Walters, R. H. (1977). *Social learning theory* (Vol. 1, pp. 141-154). Englewood Cliffs, NJ: Prentice hall.
- Barrett, T., Dowle, M., Srinivasan, A., Gorecki, J., Chirico, M., Hocking, T., Schwendinger, B., & Krylov, I. (2025). *data.table: Extension of data.frame* (R package version 1.17.99).
- Bussemeyer, J.R. & Townsend, J.T. (1993). Decision field theory: A dynamic cognition approach to decision-making. *Psychological Review*, *100*, 432-459.
<https://doi.org/10.1037/0033-295X.100.3.432>
- Cannon-Bowers, J. A., Salas, E., & Converse, S. (1993). Shared mental models in expert team decision making.

- Chi, M. T. H., Glaser, R., & Farr, M. J. (1988). *The nature of expertise*. Hillsdale, NJ: Erlbaum
- Courtright, S. H., Thurgood, G. R., Stewart, G. L., & Pierotti, A. J. (2015). Structural interdependence in teams: An integrative framework and meta-analysis. *Journal of Applied Psychology, 100*(6), 1825.
<https://doi.org/10.1037/apl0000027>
- Cronin, M. A., & Weingart, L. R. (2007). Representational gaps, information processing, and conflict in functionally diverse teams. *Academy of management review, 32*(3), 761-773.
<https://doi.org/10.5465/AMR.2007.25275511>
- Curcuruto, M., Strauss, K., Axtell, C., & Griffin, M. A. (2020). Voicing for safety in the workplace: A proactive goal-regulation perspective. *Safety science, 131*, 104902. <https://doi.org/10.1016/j.ssci.2020.104902>
- Davis, J.P., Eisenhardt, K.M., & Bingham, C.B. (2007). Developing theory through simulation methods. *Academy of Management Review, 32*, 480-499.
<https://doi.org/10.5465/amr.2007.24351453>
- DeChurch, L. A., & Mesmer-Magnus, J. R. (2010). Measuring shared team mental models: A meta-analysis. *Group dynamics: Theory, research, and practice, 14*(1), 1. <https://doi.org/10.1037/a0017455>
- 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*(6), 1035-1049. <https://doi.org/10.1016/j.leaqua.2010.10.007>

- Durso, F. T., Sethumadhavan, A., & Crutchfield, J. (2008). Linking task analysis to information relevance. *Human factors*, 50(5), 755-762.
<https://doi.org/10.1518/001872008X312369>
- Edwards, W. (1954). The theory of decision making. *Psychological Bulletin*, 51, 380-417. <https://doi.org/10.1037/h0053870>
- Edmondson, A. (1999). Psychological safety and learning behavior in work teams. *Administrative science quarterly*, 44(2), 350-383.
<https://doi.org/10.2307/2666999>
- Edmondson, A. C., & Bransby, D. P. (2023). Psychological safety comes of age: Observed themes in an established literature. *Annual Review of Organizational Psychology and Organizational Behavior*, 10, 55-78.
<https://doi.org/10.1146/annurev-orgpsych-120920-055217>
- Endsley, M. R. (1995). Measurement of situation awareness in dynamic systems. *Human factors*, 37(1), 65-84.
<https://doi.org/10.1518/001872095779049499>
- Epstein, J.M. (1999). Agent-based computational models and generative social science. *Complexity*, 4, 41-60. [https://doi.org/10.1002/\(SICI\)1099-0526\(199905/06\)4:5<41::AID-CPLX9>3.0.CO;2-F](https://doi.org/10.1002/(SICI)1099-0526(199905/06)4:5<41::AID-CPLX9>3.0.CO;2-F)
- Espinosa, J. A., & Clark, M. A. (2014). Team knowledge representation: A network perspective. *Human factors*, 56(2), 333-348.
<https://doi.org/10.1177/0018720813494093>

- Fast, N. J., Sivanathan, N., Mayer, N. D., & Galinsky, A. D. (2012). Power and overconfident decision-making. *Organizational behavior and human decision processes*, 117(2), 249-260. <https://doi.org/10.1016/j.obhdp.2011.11.009>
- Festinger, L. (1954). A theory of social comparison processes. *Human relations*, 7(2), 117-140. <https://doi.org/10.1177/001872675400700202>
- Fiore, S.M., Rosen, M.A., Smith-Jentsh, 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. <https://doi.org/10.1177/0018720810369807>
- Fiske, S. T., Cuddy, A. J. C., Glick, P., & Xu, J. (2002). A model of (often mixed) stereotype content: Competence and warmth respectively follow from perceived status and competition. *Journal of Personality and Social Psychology*, 82, 878–902. <https://doi.org/10.1037/0022-3514.82.6.878>
- Fiske, S. T., Cuddy, A. J. C., & Glick, P. (2007). Universal dimensions of social cognition: Warmth and competence. *Trends in Cognitive Sciences*, 11(2), 77–83. <https://doi.org/10.1016/j.tics.2006.11.005>
- Gigerenzer, G., & Gaissmaier, W. (2011). Heuristic decision making. *Annual review of psychology*, 62(2011), 451-482. <https://doi.org/10.1146/annurev-psych-120709-145346>
- Gigone, D., & Hastie, R. (1993). The common knowledge effect: A theory of social comparison processes. *Human relations*, 7(2), 117-140. <https://doi.org/10.1037/0022-3514.65.5.959>

- Grand, J.A., Braun, M.T., & Kuljanin, G. (2024). Hello world! Building computational models to represent social and organizational theory. *Organizational Research Methods*, 27(3), 619–652. <https://doi.org/10.1177/10944281241261913>
- Grand, J. A., Braun, M. T., Kuljanin, G., Kozlowski, S. W., & Chao, G. T. (2016). The dynamics of team cognition: A process-oriented theory of knowledge emergence in teams. *Journal of Applied Psychology*, 101(10), 1353. <https://doi.org/10.1037/apl0000136>
- 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. <https://doi.org/10.5465/amr.2007.26586485>
- Hastie, R. (2001). Problems for judgment and decision-making. *Annual Review of Psychology*, 52, 633-683. <https://doi.org/10.1146/annurev.psych.52.1.653>
- Hester, J., Wickham, H., & Csárdi, G. (2025). *fs: Cross-Platform File System Operations Based on 'libuv'* (R package version 1.6.6). <https://fs.r-lib.org> (fs.r-lib.org)
- Hinsz, V. B., Tindale, R. S., & Vollrath, D. A. (1997). The emerging conceptualization of groups as information processors. *Psychological bulletin*, 121(1), 43. <https://doi.org/10.1037/0033-2909.121.1.43>
- Hood, A. C., Bachrach, D. G., Zivnuska, S., & Bendoly, E. (2016). Mediating effects of psychological safety in the relationship between team affectivity and transactive memory systems. *Journal of Organizational Behavior*, 37(3), 416-435. <https://doi.org/10.1002/job.2050>

- Huang, S., & Cummings, J. N. (2011). When critical knowledge is most critical: Centralization in knowledge-intensive teams. *Small Group Research*, 42(6), 669-699. <https://doi.org/10.1177/1046496411410073>
- Hussain, I., Shu, R., Tangirala, S., & Ekkirala, S. (2019). The voice bystander effect: How information redundancy inhibits employee voice. *Academy of Management Journal*, 62(3), 828-849. <https://doi.org/10.5465/amj.2017.0245>
- Izrailev, S. (2025). *tictoc: Functions for Timing R Scripts, as Well as Implementations of "Stack" and "StackList" Structures* (R package version 1.2.1). <https://cran.r-project.org/package=tictoc>.
- Janis, I. L. (1973). Groupthink and group dynamics: A social psychological analysis of defective policy decisions. *Policy Studies Journal*, 2(1), 19. <https://doi.org/10.1111/j.1541-0072.1973.tb00117.x>
- Johnson-Laird, P. N. (2004). The history of mental models. *Psychology of reasoning: Theoretical and historical perspectives*, 8, 179-212. <https://doi.org/10.4324/9780203506936-10>
- Kahneman, D. (2011). *Thinking, fast and slow*. New York, NY: Farrar, Straus and Giroux.
- Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47, 263-292. <https://doi.org/10.2307/1914185>
- Kameda, T., Ohtsubo, Y., & Takezawa, M. (1997). Centrality in sociocognitive networks and social influence: An illustration in a group decision making context. *Journal of Personality and Social Psychology*, 73, 296-309. <https://doi.org/10.1037/0022-3514.73.2.296>

- Kent, S. (1964). Words of estimative probability. *Studies in intelligence*, 8(4), 49-65.
- Klimoski, R., & Mohammed, S. (1994). Team mental model: Construct or metaphor?. *Journal of management*, 20(2), 403-437.
<https://doi.org/10.1177/014920639402000206>
- Knaus, J. (2023). *snowfall: Easier Cluster Computing (Based on 'snow')* (R package version 1.84-6.3). <https://cran.r-project.org/package=snowfall> (CRAN R-
[Universe](https://cran.r-project.org/package=snowfall)).
- Knight, A.P., Kennedy, D.M., & McComb, S.A. (2016). Using recurrence analysis to examine group dynamics. *Group Dynamics: Theory, Research, and Practice*, 20, 223-241. <https://doi.org/10.1037/gdn0000046>
- 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(4), 581-615.
<https://doi.org/10.1177/1094428113493119>
- Kozlowski, S. W., & Klein, K. J. (2000). *A multilevel approach to theory and research in organizations: Contextual, temporal, and emergent processes*.
- Kozlowski, S. W., & Ilgen, D. R. (2006). Enhancing the effectiveness of work groups and teams. *Psychological science in the public interest*, 7(3), 77-124.
<https://doi.org/10.1111/j.1529-1006.2006.00030.x>
- Kuljanin, G., Braun, M.T., Grand, J.A., Olenick, J., Chao, G.T., & Kozlowski, S.W.J. (2024). Advancing leadership and organizational science with computational process theories. *The Leadership Quarterly*, 35, e101797.
<https://doi.org/10.1016/j.leaqua.2024.101797>

- Larson, J. R., Jr., Christensen, C., Abbott, A. S., & Franz, T. M. (1996). Diagnosing groups: Charting the flow of information in medical decision making teams. *Journal of Personality and Social Psychology*, 71, 315-330. <https://doi.org/10.1037/0022-3514.71.2.315>
- Larson Jr, J. R., Christensen, C., Franz, T. M., & Abbott, A. S. (1998). Diagnosing groups: The pooling, management, and impact of shared and unshared case information in team-based medical decision making. *Journal of Personality and Social Psychology*, 75(1), 93. <https://doi.org/10.1037/0022-3514.75.1.93>
- Lin, S. H. J., & Johnson, R. E. (2015). A suggestion to improve a day keeps your depletion away: Examining promotive and prohibitive voice behaviors within a regulatory focus and ego depletion framework. *Journal of Applied Psychology*, 100(5), 1381. <https://doi.org/10.1037/apl0000018>
- Lipshitz, R., & Strauss, O. (1997). Coping with uncertainty: A naturalistic decision-making analysis. *Organizational Behavior and Human Decision Processes*, 69, 149-163. <https://doi.org/10.1006/obhd.1997.2679>
- Liu, W., Tangirala, S., Lam, W., Chen, Z., Jia, R. T., & Huang, X. (2015). How and when peers' positive mood influences employees' voice. *Journal of applied psychology*, 100(3), 976. <https://doi.org/10.1037/a0038066>
- Love, B. C., Medin, D. L., & Gureckis, T. M. (2004). SUSTAIN: a network model of category learning. *Psychological review*, 111(2), 309. <https://doi.org/10.1037/0033-295X.111.2.309>

- 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*(1), 54-75. <https://doi.org/10.1177/1088868311417243>
- Mathieu, J. E., Hollenbeck J.R., Knippenberg D.V., Ilgen D. R. (2017). A century of work teams in the journal of applied psychology. *Journal of Applied Psychology, 102*, 452–467. <https://doi.org/10.1037/apl0000128>
- Mathieu, J. E., Gallagher, P. T., Domingo, M. A., & Klock, E. A. (2019). Embracing complexity: Reviewing the past decade of team effectiveness research. *Annual Review of Organizational Psychology and Organizational Behavior, 6*, 17-46. <https://doi.org/10.1146/annurev-orgpsych-012218-015106>
- Mathieu, J., Maynard, M. T., Rapp, T., & Gilson, L. (2008). Team effectiveness 1997-2007: A review of recent advancements and a glimpse into the future. *Journal of Management, 34*(3), 410-476. <https://doi.org/10.1177/0149206308316061>
- Marks, M. A., Mathieu, J. E., & Zaccaro, S. J. (2001). A temporally based framework and taxonomy of team processes. *Academy of Management Review, 26*(3), 356-376. <https://doi.org/10.5465/amr.2001.4845785>
- McClellan, E. J., Martin, S. R., Emich, K. J., & Woodruff, C. T. (2018). The social consequences of voice: An examination of voice type and gender on status and subsequent leader emergence. *Academy of Management Journal, 61*(5), 1869-1891. <https://doi.org/10.5465/amj.2016.0148>
- 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.
[https://doi.org/10.1016/S1475-9144\(07\)06005-5](https://doi.org/10.1016/S1475-9144(07)06005-5)

Medin, D. L., & Schaffer, M. M. (1978). Context theory of classification learning. *Psychological review*, 85(3), 207. <https://doi.org/10.1037/0033-295X.85.3.207>

Mesmer-Magnus, J. R., & DeChurch, L. A. (2009). Information sharing and team performance: a meta-analysis. *Journal of applied psychology*, 94(2), 535. <https://doi.org/10.1037/a0013773>

Mesmer-Magnus, J. R., DeChurch, L. A., Jimenez-Rodriguez, M., Wildman, J., & Shuffler, M. (2011). A meta-analytic investigation of virtuality and information sharing in teams. *Organizational Behavior and Human Decision Processes*, 115(2), 214-225. <https://doi.org/10.1016/j.obhdp.2011.03.002>

Mohammed, S., Klimoski, R., & Rentsch, J. (2000). The measurement of team mental models: We have no shared schema. *Organizational Research Methods*, 3, 123–165. <https://doi.org/10.1177/109442810032001>

Mohammed, S., Rico, R., & Alipour, K. K. (2021). Team cognition at a crossroad: Toward conceptual integration and network configurations. *Academy of Management Annals*, 15(2), 455-501. <https://doi.org/10.5465/annals.2018.0159>

Morrison, E. W. (2023). Employee voice and silence: Taking stock a decade later. *Annual Review of Organizational Psychology and Organizational Behavior*, 10, 79-107. <https://doi.org/10.1146/annurev-orgpsych-120920-054654>

- Ng, T. W., & Feldman, D. C. (2015). Felt obligations to reciprocate to an employer, preferences for mobility across employers, and gender: Three-way interaction effects on subsequent voice behavior. *Journal of Vocational Behavior, 90*, 36-45. <https://doi.org/10.1016/j.jvb.2015.07.005>
- Ng, T. W., Lucianetti, L., Hsu, D. Y., Yim, F. H., & Sorensen, K. L. (2021). You Speak, I Speak: The Social-Cognitive Mechanisms of Voice Contagion. *Journal of Management Studies, 58*(6), 1569-1608. <https://doi.org/10.1111/joms.12698>
- Oppenheimer, D.M., & Kelso, E. (2015). Information processing as a paradigm for decision making. *Annual Review of Psychology, 66*, 277-294. <https://doi.org/10.1146/annurev-psych-010814-015148>
- Parker, S. K., Bindl, U. K., & Strauss, K. (2010). Making things happen: A model of proactive motivation. *Journal of Management, 36*(4), 827–856. <https://doi.org/10.1177/0149206310363732>
- Payne, J. W., Bettman, J. R., & Johnson, E. J. (1993). *The adaptive decision maker*. Cambridge university press. <https://doi.org/10.1017/CBO9781139173933>
- Piaget, J. (1975). L'équilibration des structure cognitives: Problème central du développement [Equilibration of cognitive structures: The central problem of intellectual development]. Presses universitaires de France.
- Postmes, T., Spears, R., & Cihangir, S. (2001). Quality of decision making and group norms. *Journal of Personality and Social Psychology, 80*, 918-930. <https://doi.org/10.1037/0022-3514.80.6.918>

- Rand, W., & Rust, R.T. (2011). Agent-based modeling in marketing: Guidelines for rigor. *International Journal of Research in Marketing*, 28, 181-193.
<https://doi.org/10.1016/j.ijresmar.2011.04.002>
- Röllmann, L. F., Weiss, M., & Zacher, H. (2021). Does voice benefit or harm occupational well-being? The role of job insecurity. *British Journal of Management*, 32(3), 708-724. <https://doi.org/10.1111/1467-8551.12471>
- Rudman, L. A., & Glick, P. (1999). Feminized management and backlash toward agentic women: The hidden costs to women of a kinder, gentler image of middle managers. *Journal of Personality and Social Psychology*, 77, 1004–1010. <https://doi.org/10.1037/0022-3514.77.5.1004>
- Russell, R. D. (1928). A comparison of two methods of learning. *The Journal of Educational Research*, 235-238.
<https://doi.org/10.1080/00220671.1928.10879882>
- Schamber, L., Eisenberg, M. B., & Nilan, M. S. (1990). A re-examination of relevance: toward a dynamic, situational definition*. *Information processing & management*, 26(6), 755-776. [https://doi.org/10.1016/0306-4573\(90\)90050-C](https://doi.org/10.1016/0306-4573(90)90050-C)
- Schoemaker, P.J.H. (1982). The expected utility model: Its variants, purposes, evidence, and limitations. *Journal of Economic Literature*, 20, 529-563.
- Sherf, E. N., Parke, M. R., & Isaakyan, S. (2021). Distinguishing voice and silence at work: Unique relationships with perceived impact, psychological safety, and burnout. *Academy of Management Journal*, 64(1), 114-148.
<https://doi.org/10.5465/amj.2018.1428>

- Stasser, G. (1999). The uncertain role of unshared information in collective choice. In *Shared cognition in organizations* (pp. 49-70). Psychology Press.
- Stasser, G., & Stewart, D. (1992). Discovery of hidden profiles by decision-making groups: Solving a problem versus making a judgment. *Journal of personality and social psychology*, 63(3), 426. <https://doi.org/10.1037/0022-3514.63.3.426>
- Stasser, G., Stewart, D., & Wittenbaum, G. (1995). Expert roles and information exchange during discussion: The importance of knowing who knows what. *Journal of Experimental Social Psychology*, 31, 244-265. <https://doi.org/10.1006/jesp.1995.1012>
- Stasser, G., & Titus, W. (1985). Pooling of unshared information in group decision making: Biased information sampling during discussion. *Journal of personality and social psychology*, 48(6), 1467. <https://doi.org/10.1037/0022-3514.48.6.1467>
- Stasser, G., & Titus, W. (2003). Hidden profiles: A brief history. *Psychological Inquiry*, 14(3/4), 304-313. <https://doi.org/10.1080/1047840X.2003.9682897>
- Sevcikova, H., & Rossini, A. J. (2023). *rlecuyer: R Interface to RNG with Multiple Streams* (R package version 0.3-8). <http://www.iro.umontreal.ca/~lecuyer/myftp/papers/streams00.pdf>.
- Starzyk, A., Sonnentag, S., & Albrecht, A. G. (2018). The affective relevance of suggestion-focused and problem-focused voice: A diary study on voice in meetings. *Journal of Occupational and Organizational Psychology*, 91(2), 340-361. <https://doi.org/10.1111/joop.12199>

- Steiner, I. D. (1972). *Group process and productivity*. New York: Academic press.
- Stephan, F. F., & Mishler, E. G. (1952). The distribution of participation in small groups: An exponential approximation. *American Sociological Review*, 17, 598-608. <https://doi.org/10.2307/2088227>
- Stewart, D. D., & Stasser, G. (1995). Expert role assignment and information sampling during collective recall and decision making. *Journal of Personality and Social Psychology*, 69(4), 619. <https://doi.org/10.1037/0022-3514.69.4.619>
- 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(2), 248. <https://doi.org/10.1037/0096-1523.31.2.248>
- Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction* (Vol. 1, No. 1, pp. 9-11). Cambridge: MIT press.
- Sweller, J. (2011). Cognitive load theory. *Psychology of learning and motivation*, 55, 37-76. Academic Press. <https://doi.org/10.1016/B978-0-12-387691-1.00002-8>
- Taiyi Yan, T., Tangirala, S., Vadera, A. K., & Ekkirala, S. (2022). How employees learn to speak up from their leaders: Gender congruity effects in the development of voice self-efficacy. *Journal of Applied Psychology*, 107(4), 650. <https://doi.org/10.1037/apl0000892>
- Tangirala, S., & Ramanujam, R. (2008). Exploring nonlinearity in employee voice: The effects of personal control and organizational identification. *Academy of*

Management Journal, 51(6), 1189-1203.

<https://doi.org/10.5465/amj.2008.35732719>

Thomas, R.P., Dougherty, M.R., Sprenger, A.M., & Harbison, J.I. (2008). Diagnostic hypothesis generation and human judgment. *Psychological Review*, 115, 155-185. <https://doi.org/10.1037/0033-295X.115.1.155>

Tucker, S., & Turner, N. (2015). Sometimes it hurts when supervisors don't listen: The antecedents and consequences of safety voice among young workers. *Journal of occupational health psychology*, 20(1), 72. <https://doi.org/10.1037/a0037756>

Tversky, A., & Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, 5, 297-323. <https://doi.org/10.1007/BF00122574>

Vancouver, J.B., & Weinhardt, J.M. (2012). Modeling the mind and milieu: Computational modeling for micro-level organizational researchers. *Organizational Research Methods*, 15, 602-623. <https://doi.org/10.1177/1094428112449655>

van Ginkel, W. P., & Van Knippenberg, D. (2008). Group information elaboration and group decision making: the role of shared task representations. *Organizational behavior and human decision processes*, 105(1), 82-97. <https://doi.org/10.1016/j.obhdp.2007.08.005>

Venkataramani, V., Zhou, L., Wang, M., Liao, H., & Shi, J. (2016). Social networks and employee voice: The influence of team members' and team leaders' social

- network positions on employee voice. *Organizational Behavior and Human Decision Processes*, 132, 37-48. <https://doi.org/10.1016/j.obhdp.2015.12.001>
- von Neumann, J., & Morgenstern, O. (1947). *Theory of games and economic behavior* (2nd edition). Princeton, NJ: Princeton University Press.
- Wang, S., & Noe, R. A. (2010). Knowledge sharing: A review and directions for future research. *Human resource management review*, 20(2), 115-131. <https://doi.org/10.1016/j.hrmr.2009.10.001>
- Weber, E.U., & Johnson, E.J. (2009). Mindful judgment and decision making. *Annual Review of Psychology*, 60, 53-85. <https://doi.org/10.1146/annurev.psych.60.110707.163633>
- Wegner, D. M. (1987). Transactive memory: A contemporary analysis of the group mind. In *Theories of group behavior* (pp. 185-208). New York, NY: Springer New York. https://doi.org/10.1007/978-1-4612-4634-3_9
- Weiss, M., & Morrison, E. W. (2019). Speaking up and moving up: How voice can enhance employees' social status. *Journal of Organizational Behavior*, 40(1), 5-19. <https://doi.org/10.1002/job.2262>
- Welsh, D. T., Outlaw, R., Newton, D. W., & Baer, M. D. (2022). The social aftershocks of voice: An investigation of employees' affective and interpersonal reactions after speaking up. *Academy of Management journal*, 65(6), 2034-2057. <https://doi.org/10.5465/amj.2019.1187>
- Wickham, H. (2007). Reshaping data with the reshape package. *Journal of Statistical Software*, 21(12), 1–20. <https://www.jstatsoft.org/v21/i12/> (CRAN)

- Wickham, H. (2016). *ggplot2: Elegant Graphics for Data Analysis*. Springer-Verlag New York. <https://ggplot2.tidyverse.org>.
- Wickham, H., Averick, M., Bryan, J., Chang, W., D'Agostino McGowan, L., François, R., ... Yutani, H. (2019). *Welcome to the tidyverse*. *Journal of Open Source Software*, 4(43), 1686. <https://doi.org/10.21105/joss.01686>
- Wickham, H., François, R., Henry, L., Müller, K., & Vaughan, D. (2025). *dplyr: A Grammar of Data Manipulation* (R package version 1.1.4). <https://dplyr.tidyverse.org> (dplyr.tidyverse.org)
- Wintle, B. C., Fraser, H., Wills, B. C., Nicholson, A. E., & Fidler, F. (2019). Verbal probabilities: Very likely to be somewhat more confusing than numbers. *PLoS One*, 14(4), e0213522. <https://doi.org/10.1371/journal.pone.0213522>
- Wittenbaum, G. M., Hubbell, A., & Zuckennan, C. (1999). Mutual enhancement: Toward an understanding of the collective preference for shared information. *Journal of Personality and Social Psychology*, 77, 967-978. <https://doi.org/10.1037/0022-3514.77.5.967>
- Wittenbaum, G. M., & Park, E. S. (2001). The collective preference for shared information. *Current Directions in Psychological Science*, 10, 70-73. <https://doi.org/10.1111/1467-8721.00118>
- 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(3), 286-310. <https://doi.org/10.1080/0363452042000299894>

- Xiao, Y., Zhang, H., & Basadur, T. M. (2016). Does information sharing always improve team decision making? An examination of the hidden profile condition in new product development. *Journal of Business Research*, 69(2), 587-595. <https://doi.org/10.1016/j.jbusres.2015.05.014>
- Xie, X. Y., Ling, C. D., Mo, S. J., & Luan, K. (2015). Linking colleague support to employees' promotive voice: A moderated mediation model. *PloS one*, 10(7), e0132123. <https://doi.org/10.1371/journal.pone.0132123>
- Xu, Y., & Chen, Z. (2006). Relevance judgment: What do information users consider beyond topicality?. *Journal of the American Society for Information Science and Technology*, 57(7), 961-973. <https://doi.org/10.1002/asi.20361>