An Agent-Based Model of Collective Decision-Making: How Information Sharing Strategies Scale With Information Overload

Dirk-Jan van Veen, *Graduate Student Member, IEEE*, Ravi S. Kudesia, and Hans R. Heinimann, *Member, IEEE*

***Abstract*— Organizations rely on teams for complex decision-making. By bringing diverse information together and utilizing information sharing strategies, teams can make intelligent decisions. However, as organizations face increasing information overload, it has become unclear whether such strategies remain adequate or whether bounds on human rationality will prevail. We develop an agent-based model that simulates information sharing in teams, where critical information is distributed across its members. We tested how robust various information sharing strategies are to information overload and bounds on rationality in terms of the speed and accuracy of collective decision-making. Our results suggest distinct strategies depending on whether speed or accuracy is imperative and, more broadly, shed light on how intelligence is best attained in collective decision-making.**

***Index Terms*— Agent-based model, collective decision-making, distributed cognition, hidden profile, organization theory.**

I. Introduction

WHY do people form organizations, rather than solve their problems through crowds, markets, or communities? The answer to this—the foundational question in organization theory—concerns information: organizations are uniquely suited to integrate diverse information from across multiple individuals to produce intelligent collective decisions [1]. Organizations can bring together individuals with different areas of functional expertise and use mechanisms, such as teams, routines, and hierarchy, to integrate their diverse information when making decisions. In fact, Simon [2, pp. 18 and 19] even canonically defined the organization in terms of "the pattern of communications and relations among a group of human beings, including the processes for making and implementing decisions." As such, the integration of information in organizations continues to be a topic of intensive research and immense practical importance (e.g., [3]–[5]).

Manuscript received May 8, 2019; revised November 9, 2019 and February 28, 2020; accepted March 11, 2020. Date of publication May 6, 2020; date of current version June 10, 2020. This work was supported by the Future Resilient Systems (FRS) Lab, Singapore-ETH Centre (SEC), which was established collaboratively between ETH Zurich and Singapore's National Research Foundation, through its Campus for Research Excellence And Technological Enterprise (CREATE) Program, under Grant FI 370074011.
(Corresponding author: Dirk-Jan Van Veen.)

Dirk-Jan van Veen and Hans R. Heinimann are with ETH Zurich, Singapore-ETH Centre, Future Resilient Systems, Singapore 138602 (e-mail: dvanveen@frs.ethz.ch).

Ravi S. Kudesia is with the Fox School of Business, Temple University, Philadelphia, PA 19122 USA (e-mail: rskudesia@temple.edu).

Digital Object Identifier 10.1109/TCSS.2020.2986161

Organization theorists, particularly those in the Carnegie tradition following Simon [6], [7], have emphasized the importance of studying human information processing in the context of the organization in which decisions are made [8]. For instance, many organizational decisions can be made in a routine matter by individuals, but certain types of decisions are problematic because they are both nonroutine and complex—and, therefore, exceed the capacity of any one person to solve [9], [10]. Because these problematic decisions are nonroutine, they cannot be solved by individuals relying on past experience or standard operating procedures [1]. Because they are complex, these problematic decisions require knowledge of the interdependencies among diverse information sets. As such, they are best solved through interactions among multiple individuals in a team, each of whom has access to different information [11]. Indeed, this is why teams are the favored mechanism for complex strategic decisions in organizations, such as in research and development teams and top management teams: team members can search their respective information sets for the most diagnostic information for the current decision, share this diverse information through communication, and, thereby, collectively integrate the information into an intelligent decision [12]–[14]. Recent meta-analyses accordingly find that teams benefit from deep-level diversity—for example, when team members possess unique information because they have different functional expertise in marketing, finance, and so forth—and that this diversity is more useful for teams making decisions about complex problems with interdependent information [15], [16]. In contrast to surface-level diversity of demographic characteristics, which offers little benefit for decision-making, this deep-level diversity of informational expertise confers an advantage for teams faced with complex problems [16].

Despite the advantage of possessing more diverse information, teams are not always effective at actually integrating this information during decision-making [4]. In fact, an influential paradigm in collective decision-making, known as the hidden profile paradigm, has repeatedly shown that when teams have deep-level diversity in their information, they tend to share the information they all have in common, rather than communicating their unique individual information—and that doing so makes them eight times less likely to arrive at intelligent collective decisions [17]–[19]. Given that teams are the locus of complex problem solving in organizations, their inefficiencies in sharing the most vital information constitute a fundamental flaw in the very reason people form organizations [20].

Two literatures speak to this flaw. One focuses on identifying the bounds on rationality that impair decision-making, and another focuses on formal interventions that can improve collective decision-making. The first literature emphasizes how individuals are subject to bounds on their rationality and how the quality of their information sharing determines the intelligence of their collective decisions [7]. One consequence of this bounded rationality is that individuals are most able to search for information in their local neighborhood, that is, around their current train of thought [21], [22]. People are less able to search their memories and environments for information that is distant. Another important consequence of bounded rationality in organizations concerns communication. Not only do people have limitations that cause them to anchor on their personal information, which they receive before discussion [23]–[25], but they frequently have good social reasons to weigh their individual information more heavily than information shared by others, as other members of a team may have different motives and incentives [11], [20], [26]. Taken in sum, the intelligence of collective decision-making, therefore, relies on an interplay between epistemic motivations to search for information deeply and social motivations in how much one trusts the self, relative to others [27].

Regarding the second literature, a large body of work has examined information sharing strategies in teams that influence collective decision-making and how formal interventions can enhance their intelligence. In general, teams make better decisions to the extent that members share their individual information [4], [19]. Doing so requires that team members are individually and collectively motivated to search the broader information space, rather than searching only around their local neighborhood [22], [27]. Formal interventions that prompt team members to disagree with each other, such as devil's advocacy and dialectical inquiry, facilitate a more thorough search of this information space by counteracting the tendency toward groupthink and emphasizing shared information [28]–[33]. Indeed, deep-level diversity especially enhances decision-making when teams use formal interventions [4].

In this study, we examine how bounded rationality and information sharing strategies impact collective decision-making, specifically in the light of information overload. Information overload has become an incredibly salient issue for organizations, particularly with the adoption of new information technologies [34]–[38]. These new technologies make access to information both faster and cheaper, producing opportunities and challenges alike for organizations [38]. Namely, although these technologies can provide information that aids decision-making [39], [40], ensuring the relevance of this information for decision-makers remains a challenge [40]—particularly given the prevalence of unreliable and low-quality information sources [34], [41]. As a consequence of the increasing scope and reach of these information technologies [42], [43], today's organizations face an abundance of information when making decisions, leading both to fatigue and poorer decision quality [44]. Following Simon's [45, pp. 40 and 41] prediction about information technologies that "a wealth of information creates a poverty of attention and a need to allocate that attention efficiently among the overabundance of information sources that might consume it," organizations have started treating attention as a scarce currency and wondered how it might be better allocated [46]. The technology firm Intel, for instance, calculated that employees lose approximately 8 hours each week to information overload, which translates into an annual cost of $1 billion for a firm of its size [47].

Beyond its practical relevance for organizations, information overload also introduces important theoretical considerations. First, the more the information available for making a decision, the greater the information processing demands placed on individual team members [19], [48]. Information overload would, therefore, seem to exacerbate the bounded rationality of individual team members. However, it is possible that, under certain conditions, communication within the team could help overcome this bounded rationality. For instance, perhaps formal interventions could help team members search their growing information space more effectively. Given that there has been relatively little research on information overload in teams, it remains unclear how teams can be made more robust to information overload [49]. Second, and relatedly, information overload speaks directly to the temporal dynamic inherent within collective decision-making. Common information tends to be shared first in teams and only with time does the individual information come out, depending on the strategies team members use to share information (such as by emphasizing disagreement over agreement) [50]. Therefore, the effect of different information sharing strategies seems to manifest especially over time. However, the research on collective decision-making has not adequately explored this temporal dimension, focusing mostly on accuracy rather than on how information overload impacts the tradeoff between speed and accuracy [51]–[53]. It is possible that the most effective information sharing strategy at low levels of information load may not perform as well when the information load increases. It is further possible that the greatest accuracy in decision-making is reached using information sharing strategies that do not scale well with information overload and that organizations that value speed should utilize a different strategy. However, neither the speed–accuracy tradeoff nor its specific relation to information sharing strategies has been explicitly studied.

To explore how the speed and accuracy of collective decision-making scale with information overload, we utilize an agent-based model that computationally simulates multiple team members and their interactions with each other and their environment [54]. An agent-based model is the appropriate method for this research question for four reasons. First, it is difficult to experimentally manipulate different aspects of bounded rationality in laboratory settings [55]. One can instead study naturally occurring individual differences in bounded rationality [56], but this approach precludes strong causal claims. This is, indeed, why scholars have embraced agent-based models to explore bounded rationality [57]–[59]. Second, using formal interventions to examine the effects of information sharing strategies in laboratory settings has limitations. For instance, these interventions can prompt subjects

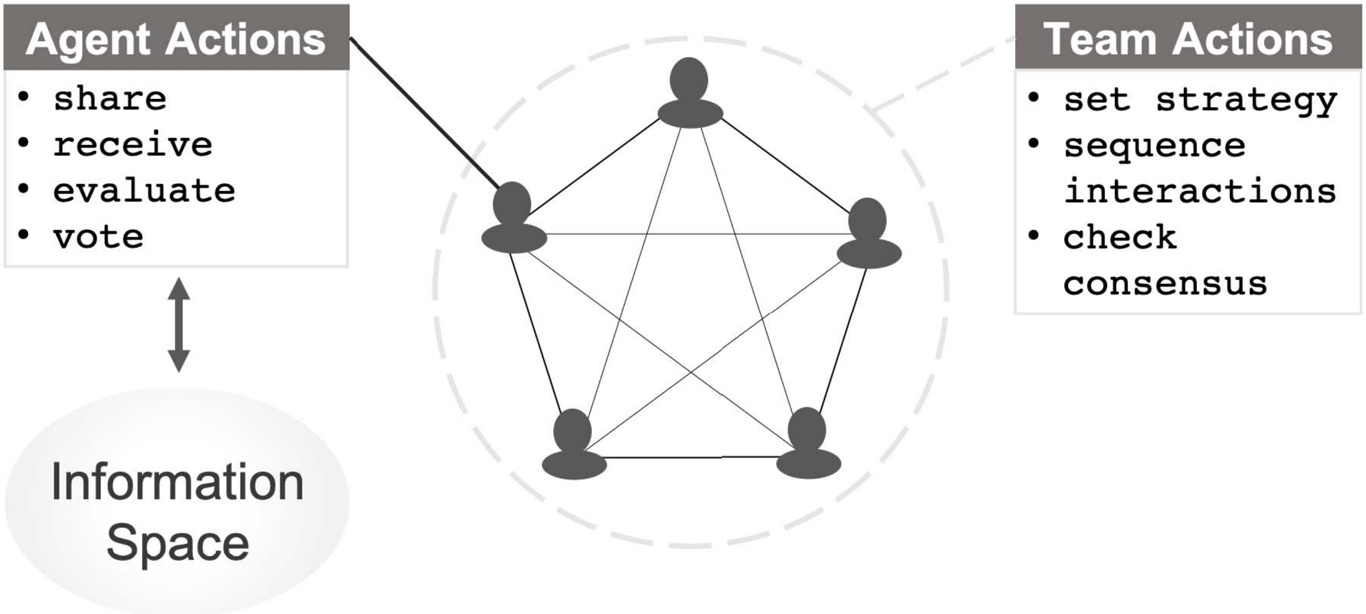

Fig. 1. Conceptual model of the collective decision-making setting. A fully connected network of agents that share information in order to make more intelligent collective decisions. Together, the agents must reach a consensus on the best decision. Agents are responsive to team-level processes that influence what information they share.

to switch away from decision-making and have secondary conversations about how to improve their communication processes [31]. The observed effects of formal interventions in laboratory settings can, therefore, conflate actual improvements from the formal interventions with those from the secondary conversations. Since agents in an agent-based model are programed, they do not suffer from unintended secondary behaviors. Third, because agent-based models simulate the behavior, they provide an opportunity to run experiments that simultaneously explore behavior across a wide parameter space that would, otherwise, require an intractably large number of human subjects to examine [60]. In this way, we can replicate prior work showing the importance of features, such as team diversity and information load [4], [19], and extend it in the light of our specific research questions. Fourth, and most generally, agent-based models are valuable for research questions in which the outcome of interest is emergent from the interactions of several individuals, as is the case with collective decision-making [54], [57].

In what follows, we first describe the architecture of our agent-based model and the collective decision-making task agents are concerned with. Then, to explore the effect of bounded rationality in agents, we lay out a basic model that explores how four information sharing strategies scale with information overload in terms of the speed and accuracy of decision-making, sequentially introducing two key features of bounded rationality. Finally, we conclude with a discussion of our results and recommendations for future research and practice.

II. Agent-Based Computational Model

A. Purpose

Our model aims to evaluate how various information sharing strategies affect the speed and accuracy with which a team of agents makes decisions. Each agent in the simulation can take up, store, and evaluate information on its own. Agents can also influence the opinion of other agents by sharing information, thereby improving the intelligence of team-level decisions. Following Simon [61] and Shannon and Weaver [62], the information in our model has meaning for a particular decision situation; it can be stored in and retrieved from memory, shared with others through communication, and subjectively evaluated in terms of how much it reduces uncertainty about the decision to be made.

The implementation of the model follows an "agent-and-information modeling" approach [63], which sets apart the physical space and the information space. In many agent-based models, agents act based on simple if-then rules. The agent-and-information modeling paradigm allows for more complex motivations in action selection. Under the agent-and-information modeling paradigm, agents select actions based on a set of information where information may be conflicting and selection may be stochastic. This requires that we model the information each agent has access to the information space. Actions performed by one agent in the physical space, namely, sharing new information, can affect the information space of another agent. The physical space and information space can, thus, evolve both independently and interactively.

B. Conceptual Framework

Fig. 1 illustrates our conceptual framework, which consists of a fully connected team of agents that are able to communicate peer-to-peer. Each agent has an information space, which is a repository of information akin to a person's memory. The actions of agents follow typical patterns of conversation. The actions `receive` and `share` emulate turn-taking in conversation [64]–[66] and facilitate the flow of information in and out of memory. The action `evaluate` emulates the information appraisal process [61], [67] and `vote` keeps track of consensus formation [66], [68]. The team-level information space is a repository that keeps track of the votes of individual agents and previously shared information and represents a collective memory. Three functions characterize team behavior: `set strategy` specifies a set of instructions that tell agents how to select information to share with other agents and `sequence interactions` refers to the order in which agents are sharing or receiving information, while `check consensus` aggregates the votes of agents and stops the simulation when consensus has been achieved.

C. Scope

Organization theorists, particularly in the Carnegie tradition, emphasize that decisions are to be judged from within a particular "definition of the situation" [1]. As such, we consider a closed system of information, where all the available information is distributed across the team members at the start of the simulation, and this information, when taken in sum, is adequate to indicate one decision alternative over the others. No additional information is added at a later stage. For the purpose of this study, the complete information set accurately represents their definition of the decision situation, which allows us to examine the speed with which information sharing leads to accurate collective decisions. Doing so leaves open the possibility that the team may have improperly defined the situation with inadequate information gathering or incomplete development of alternatives. These questions, which relate to problem formulation, fall beyond the scope of our research question about information sharing [20]. The closed system of information modeled in this study also closely mirrors the hidden profile paradigm [17], which is widely used to explore team decision-making [19], [69].

As Simon noted, decision-making must be examined in terms of bounded rationality and the organizational context in which decisions are made, both of which are addressed in our model [6], [7]. In terms of bounded rationality, we capture both the epistemic and social aspects of bounded rationality that impact collective decision-making. These aspects influence both the search for information and the evaluation of information. First, as noted previously, when searching for relevant information, people are often limited to their local neighborhood: the information nearest to the information from which they are currently operating [21], [22], [58]. The local nature of search can limit their ability to identify and share their most relevant information and may be especially limiting with high information load, as their individual information space grows in size. To this end, we create, for each agent, an associative memory and vary the extent to which search is constrained to its local neighborhood. Second, also noted previously, when evaluating information, people often anchor on their initial individual information [23]–[25], thereby discounting information they receive subsequently from other team members [11], [20], [26]. To this end, we give agents the capacity to update their decision preference based on additional information (i.e., Bayesian belief updating [70]–[72]). However, we vary the degree to which they weigh their individual information over the information they receive from others to explore the consequences of anchoring (see Algorithm 1).

Turning from the bounded rationality of agents to the context of decision-making, in our model, the organizational environment consists of the hidden profile decision task, the other agents, and the distribution of information across them. The hidden profile constrains the information that an agent initially has access to, which can be augmented as other agents share additional information. This reflects the core insight that organizations structure the flow of information into teams [2], [73]. It further aligns with the finding within the hidden profile paradigm that the allocation of information within teams—i.e., the team's deep-level diversity—is an important driver of intelligent decision-making [4], [19], [74]. We, therefore, vary both the diversity of information within teams and the degree of information overload, which is itself influenced by organizational factors, such as the use of information technology and organizational culture [35]–[37], [75].

D. Decision Task

We examine the team decision task shown in Fig. 2, which is modeled after the seminal hidden profile paradigm [17], [19], a multicriteria decision scheme in which a team of individuals must pick one option from a set of decision alternatives $\{\mathbf{O_1}, \mathbf{O_2}, \ldots\}$ based on a set of information attributes $\{\mathbf{a_1}, \mathbf{a_2}, \ldots\}$. Each attribute has one attribute score (between 0 and 1, higher is better) for each alternative. Attribute scores indicate the degree of support for an alternative and are drawn from a uniform distribution $\mathbf{x_i} = [x_{i1}, x_{i2}, \ldots], \quad x_{ij} \sim \mathrm{U}\{0, 1\}$, such that each attribute is independent, and then normalized $\mathbf{a_i} = (\mathbf{x_i}/(\| \mathbf{x_i} \|))$ so each attribute is weighted equally ($\sum \mathbf{a_i} = 1$). The correct decision is the one with the highest cumulative attribute score $\mathbf{O_{correct}} = \max\{\sum_{i=1}^{n} x_{i1}, \sum_{i=1}^{n} x_{i2}, \ldots\}$.

The hidden profile paradigm captures important team dynamics in organizational decision-making, as when a top management team is composed of chief executive, operations, marketing, finance, and technology officers. Each of these team members possesses a different set of information. To make intelligent strategic decisions, team members must share and integrate this information, despite their bounded rationality in searching for and evaluating information [20].

E. Information Distribution

Each attribute in Fig 2 has a label "common" or "individual." Following [48], the label "common" indicates that an attribute is known by all agents, whereas "individual" means that this attribute is only known by a single agent. At the start of the simulation, common information and individual information are distributed to agents in equal amounts, such that half of an agent's information is common information and half is individual. Since we generate attribute scores randomly, it may occur that a majority of agents already agree on the optimal decision, based on their initial information. We exclude these cases from our analysis because no interaction is needed. In the hidden profiles that we analyze, the incomplete set of information favors a suboptimal decision alternative, whereas all information combined reveals the

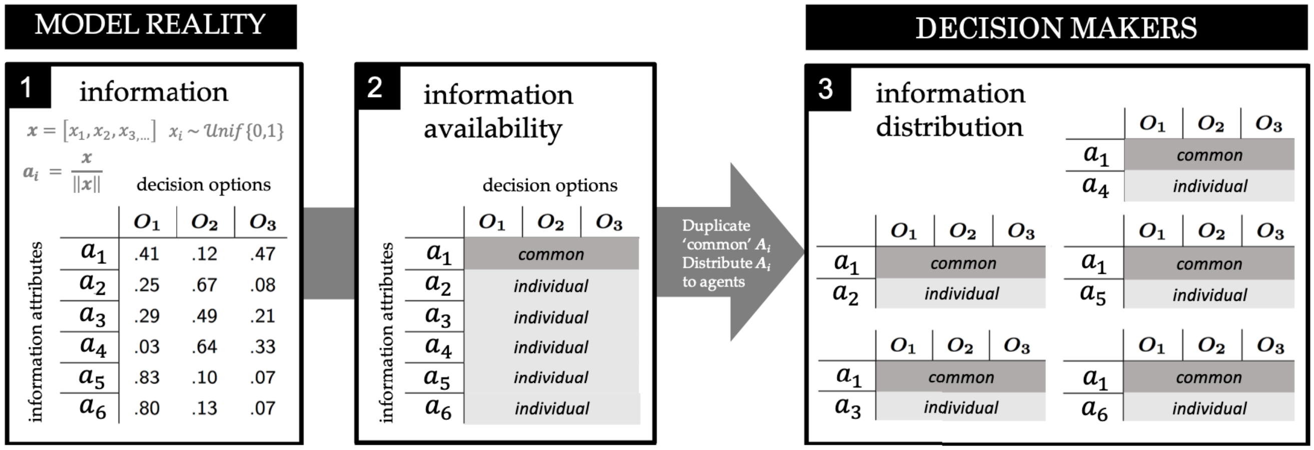

Fig. 2. Creation of decision situations used in this simulation. Box 1: for each decision alternative $\mathbf{O_i}$, a number between 0 and 1 is drawn from a uniform distribution. These numbers are normalized by row to form one information attribute $\mathbf{a_i}$. Box 2: each information attribute is assigned a label "common" or "individual," meaning that this attribute is either known by all agents or by just a single agent. Box 3: each agent has a memory in which it can store information. Information is distributed equally among agents, such that each agent has an equal number of common and unique information attributes. We vary total information attributes in steps of 10 (5 agents $\times$ 1 common + 5 $\times$ individual) such that there is never an imbalance in the number of attributes each agent receives initially.

Algorithm 1 Update Decision Preference

Input: A, W, a, w
A ▷ attributes in memory of agent, e.g. [[0.31, 0.42, 0.27], [0.23, 0.07, 0.70], …]
W ▷ weights of attributes in memory, e.g. [1, 1, …]
a ▷ newly shared attribute, e.g. [0.63, 0.32, 0.05]
w ▷ weight of new attribute, e.g. 0.5
Output: pref, A, W
pref ▷ agent's preferred decision

function UPDATE_DECISION_PREFERENCE(A, W, a, w)
  $A \leftarrow A.append(a)$
  $W \leftarrow W.append(w)$
  AW $\leftarrow A * W$
  AWsum $\leftarrow$ sum(AW)
  pref $\leftarrow$ max(AWsum)
  return pref, A, W

optimal decision alternative, relative to the team's definition of the situation. As a result, the optimal decision alternative is hidden to the team as a whole but can be discovered if agents share their individual information with each other.

F. Agent Behavior

Agents associate certain information with other information. Hence, we represent their information space as a memory network in which each node represents an attribute and each edge represents an association between two attributes (see Fig. 3). This network resembles an agent's personal view of the decision situation. The network is constructed as follows.

We first create a random graph with n_{common} nodes and d edges per node, where all nodes represent common attributes. For each agent, we rewire 20% of edges, such that each agent's network of "common" information is similar, but different. This represents how agents have only partially overlapping understandings of decision situations. After rewiring, nodes representing "individual" information $n_{\text{individual}}$ are added to each agent's network and connected randomly to d other nodes, which can be both common nodes and other individual nodes. As the network grows larger, the likelihood of duplicate connections decreases.

The topology and connectedness of the network affect how efficiently agents can search for information. Because the topology of such memory networks at the neurological level remains unknown [76], [77], random topologies are often used to model such associative relationships [77], [78]. We, therefore, use a random topology although since our agents' memories consist of only 2–20 nodes, the topology is less important than the level of connectedness. The number of edges per node d represents the cognitive capacity of agents to associate attributes. As the number of attributes increases, the network becomes sparser and less dense.

When an agent takes a turn to share information, it must search for an attribute to share. Agents search for information to share using a hill-climbing heuristic. Initially, the search starts at a random node in the agent's memory, which gives

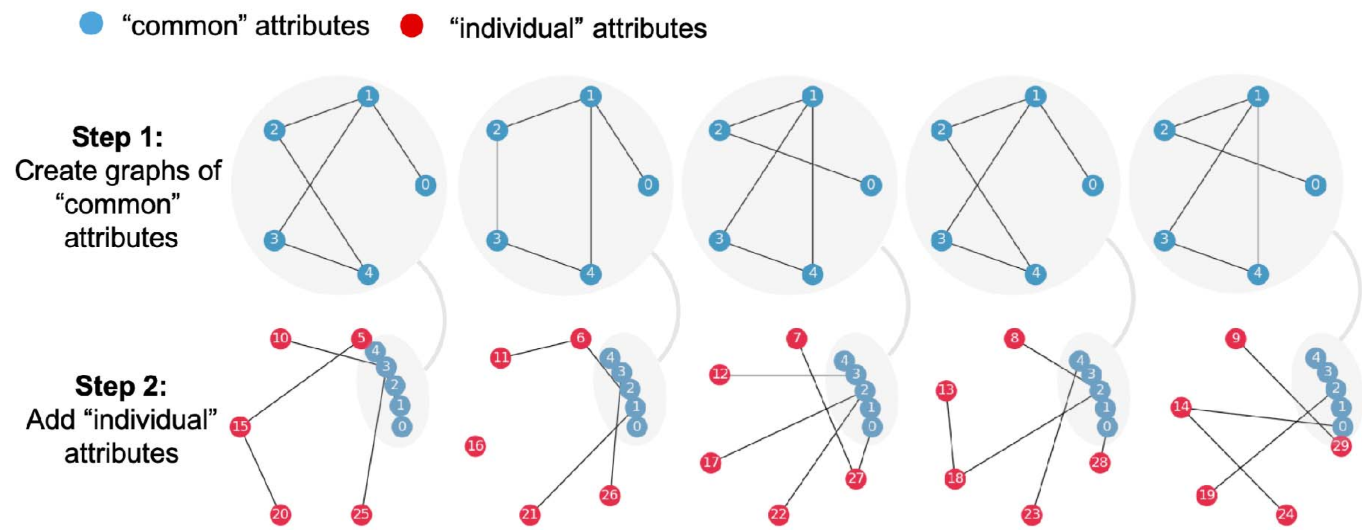

Fig. 3. Associative memories of five agents in a hidden profile with 50 attributes where half the information is "common" and the other half "individual."

it a neighborhood. The neighborhood N includes the node itself and its neighbors (connected nodes). Each node in the neighborhood is associated with a particular attribute. Agents then pick the best attribute in the neighborhood. Next, the agents move to the node that is associated with the best attribute and stay if the current node is associated with the "best" attribute. At the next turn to share information, the agent starts at said node.

Agents follow an information sharing strategy to select which attribute is the "best." The strategy is determined at the team level, through the `set strategy` action, prior to discussion. We simulate teams that adopt one of the following four information sharing strategies.

1) *Agreement:* The agent shares the attribute that most strongly supports the decision alternative that the team voted for in the previous round of communication. The focus is on the team-level opinion, and sharing information that maximally agrees with and supports it, which is found to impair decision-making, as with the well-known groupthink phenomenon [29], [79], [80].
2) *Advocacy:* The agent shares the attribute that most strongly supports the agent's own belief about the correct decision alternative. This represents the natural tendency of team members to argue according to their decision preference, commonly seen in free discussion groups [81]–[83].
3) *Disagreement:* The agent shares the attribute that most strongly supports a different decision alternative than the alternative supported by the last shared attribute. The focus is on the information-level frame and sharing contrarian information, mimicking formal interventions, such as devil's advocacy and dialectical inquiry [28], [30], [32], [33], [84]–[86].
4) *Random:* The agent shares an attribute from its memory at random. This strategy acts as a benchmark. It reflects the principle that collective decision-making and search can benefit by the infusion of random variability, which team members can selectively retain [87]–[89].

When sharing information at Random, an agent selects an attribute by drawing with equal probability one attribute $\mathbf{a_i}$ from the attribute set that comprises its local neighborhood.

In the case of Advocacy, Agreement, and Disagreement, the information sharing strategy determines which decision alternative the agent wants to argue for. Next, it searches within its neighborhood for the strongest attribute score $x_{\max}$ supporting that decision alternative

$$x_{\max} = \max\left(x_{ij} \in [\mathbf{a_i}, \ldots, \mathbf{a_n}]^{\mathrm{N}}_{\mathrm{agent}}\right). \tag{1}$$

Finally, the agent shares the attribute that contains the most supportive attribute score

$$\mathbf{a_{share}} = \mathbf{a_i} \ni x_{\max}. \tag{2}$$

Note that this search process is not optimal. The optimal attribute is the one that has the highest attribute score relative to other attribute scores $x_{ij} \in \mathbf{a_i}$. For instance, it could be smarter for an agent to share an attribute with values [0.3, 0.4, 0.3] than [0.15, 0.43, 0.42] (assuming alternative C is the incumbent alternative), yet the latter would be selected under this local neighborhood search.

When agents share information, they share it with all agents at once, in a broadcasting fashion. They share only complete attributes, comprising attribute scores for all alternatives. Information is received without distortion, mimicking a copy-paste operation of $\mathbf{a_i}$ from one agent's memory into another agent's memory. We investigate the robustness of our results when noise is added to this transfer of attributes at a later stage but do not consider biases, whereby an agent may misremember information due to motivated information processing or favor information from certain agents over others due to relationship history.

An agent's information appraisal involves the functions `evaluate` and `vote`. Agents evaluate their preferred alternative using the Bayesian updating: agents sum the attribute scores of all the attributes available in their memory $\sum_{n=1}^{i} x_{ij}$ (each attribute score x_{ij} can be interpreted as a probability that alternative j is the optimal solution). The alternative with the highest cumulative attribute score gains the agent's preference. Hence, agents can change their preference when they receive new attributes that tilt the cumulative attribute score in favor of another alternative (updating). Agents vote for their preferred alternative. We investigate how anchoring on their initial individual information influences this process later.

G. Team Behavior

Team-level dynamics are governed by three functions. Prior to the discussion, the function `set strategy` selects one of the four information sharing strategies for the team. Once a strategy has been chosen, it applies to all agents and does not change during the simulation.

The function `sequence interactions` determines the sequence in which agents share information using random assignment. All agents are equally likely to claim the turn to share information.

The function `check consensus` regulates when agents stop sharing information. This occurs either when the correct alternative receives the majority vote or when the agents share the same attribute for γ times consecutively, indicating that the team is stuck on a suboptimal decision.

H. Sensitivity Analysis

The quality of team decision-making is measured in terms of both accuracy and speed. Accuracy is defined as

$$\text{accuracy} = \frac{\text{discussions with correct decisions}}{\text{total discussions}}. \tag{3}$$

For discussions that end with a correct decision, we report the number of interactions until consensus, where fewer interactions represent faster speed. Because the decision task is randomly generated, the required number of interactions varies per experiment. Therefore, we continue repeating simulation experiments until the median number of interactions remains stable for 2000 consecutive experiments, which we report as the speed of decision-making

$$\text{speed} = \text{median(number of interactions)}. \tag{4}$$

Accuracy and speed are calculated for simulation experiments that span the range of parameter values, which are described and justified in Table I.

I. Implementation Details

The model described earlier was implemented in Python [100] and run on a computing cluster using Docker [101]. NumPy [102] was used for mathematical operations.

We calibrated the stopping threshold γ by setting this parameter at 20, 30, and 50 interactions and found this does not change results qualitatively, which we tested by plotting the simulation results and doing a side-by-side comparison. The curves follow the same shape, and we only observe negligible shifts in accuracy. We suspect that when $\gamma \leq 10$, decision accuracy will be affected more substantially because the conversation would be cut off prematurely.

To ensure that the simulation is in a steady state when we report results, we require that the median number of interactions must at least be stable 2000 times in a row, which is checked every 100 steps. This includes steps where the team does not solve the problem. The lowest reported accuracy rate is ~10%, so the minimum number of solved problems per run is 200. The median tends to stabilize between 2200 and 10 000 steps for Advocacy, Agreement, and Disagreement and up to 20 000 steps for Random. These numbers should be multiplied by the accuracy rate to get an estimate of the sample size. By the time the median stabilizes, the average has not yet stabilized, indicating that there is still some variability left and that we do not report results from a wastefully large sample size.

Proper implementation of the model was achieved through unit testing, a standard procedure in software development to validate that each unit of the software performs as intended.

III. Model Application

We designed a number of multifactorial experiments. Two factors—information sharing strategy and information load—are varied in all experiments. The information sharing strategy entails four factors: Agreement, Advocacy, Disagreement, and Random (see Section II-A5), whereas information load covers a range between 10 and 100 attributes. We discretized the information load into ten classes, which results in a $4 \times 10 \times n$ layout with $40n$ possible combinations, which we all covered in our experiments.

We performed ten runs per cell, which yielded 400 simulation runs for each tertiary variable. For every combination, we report the percentage of correct decisions and the median number of interactions it took to reach a correct decision. The median number of interactions is an integer number. To increase the resolution of decision speed from 0 to 1 decimal place, we perform ten runs per cell. Less than ten runs would result in rounding errors at the 0.1 significance level. Thus, we report here the averaged median of ten simulation runs. For reasons of clarity, we specified the terminology that we use to describe batches of simulation runs in Table II.

A. Descriptive Analysis

Fig. 4 shows the speed–accuracy tradeoff for the four information sharing strategies, under different levels of bounded rationality in the search for information (i.e., in the degree to which search is constrained to the local neighborhood). We consider the strategies first. In terms of decision accuracy, we can rank the strategies as follows: Random > Disagreement > Advocacy > Agreement. In terms of decision speed, we rank Advocacy > Agreement > Disagreement > Random.

Furthermore, we observe that decision accuracy decreases with increasing information load in a nonlinear fashion. The number of interactions needed to reach consensus varies per strategy but is generally low and increases gradually with increasing information load in most cases.

B. Effect of Information Sharing Strategies

We find that Random scanning of the information space is thorough (since all information is considered) but not efficient (as there is no selection structure for information sharing). Because there is a generous limit on how often agents are

TABLE I

PARAMETERS USED IN THE SIMULATION: NOTATION, DESCRIPTION, VALUE, AND JUSTIFICATION

Parameter	Description	Input Values*	Justification
v	agents	5	Five has been found to be an optimal size for team decision-making [90]–[92], given that conversations typically do not work with much more than four individuals [93]. Top management teams also have around 5 members [94]. An odd-number of agents also prevents decision draws.
$\mathbf{O}$	decision alternatives	3	Considering a limited number of possible alternatives is a common heuristic solution to make complex decisions where possible alternatives are manifold [95]. Our model does not concern the process of narrowing down alternatives, but sharing information to decide which alternative to select.
$\mathbf{a}$	information attributes	10-100	A comprehensive meta-analysis of 65 hidden profile studies used between 6 and 76 attributes [19]. We encompass and extend this range a bit further to better observe trends. We simulate information load in multiples of 10 attributes. This way it is always possible to split information in exactly equal proportions of common and individual information.
S	information sharing strategy	*Random, Advocacy, Agreement, Disagreement*	See Agent Behavior section for further elaboration and justification for these information sharing strategies [4], [28]–[33] .
r	ratio $\frac{individual\ information}{total\ information}$	0.25, 0.5, 0.75	A comprehensive meta-analysis of 65 hidden profile studies used ratios between 0.30 and 0.75, with an average ratio of 0.50 [19]. We use this baseline of .50 (half common information and half individual information) and explore ratios of .25 to simulate homogeneity and .75 to simulate diversity.
γ	stopping threshold (repetitiveness)	30	Number of consecutive turns after which team discussion is terminated if no new unique information has been shared. After 30 turns, it was extremely unlikely, for all levels of information load, that agents will share new information that will change the beliefs of another agent. This value was set empirically: increasing or lowering this value by 10 does not meaningfully change results.
d	edges per node in memory network	1, 2	Since memory networks contain 2-20 nodes, a value of $d \geq 2$ nearly approximates a connected network, $d = 1$ allows for sparsity and disconnectedness, which represents a limited ability to search for information. $d = 0$ means no nodes are connected and is not meaningful.
w	weight of discussed information	1.0 0.75, 0.5	The extent to which people anchor on initial information varies per task [24]. Relevant research on information recall has shown that initial information is 2-2.5 times more likely to be recalled [96] and research on priming shows that anchoring produces outcomes that deviate by a factor of 1 to 3.5 [97]. We select readily interpretable values in our model, such that $w = 1.0$ means no anchoring and that all information weighted equally, whereas $w = 0.5$ entails that initial individual information is weighted twice as much as the subsequent information shared by others.
ϵ	noise level (subjective interpretation)	0.0 0.05, 0.30	Information shared can be miscommunicated, misperceived, or misremembered—introducing noise that may harm decision-making [98], [99]. We added noise ϵ to all attribute scores agents received x_{ij} of $\mathbf{a_i}$, such that $x'_{ij} = x_{ij} + \epsilon_j, \quad \epsilon_j \sim \mathrm{U}\ (-\epsilon, \epsilon)$ and $\mathbf{a'_i} = [x'_{i1}, x'_{i2}, ...]$. Experiments with $\epsilon = 0, \epsilon = 0.05$ and $\epsilon = 0.30$ (recall $0 < x_{ij} < 1, \sum \mathbf{x_i} = 1$) showed that our results are fairly robust to noise. Only for $\epsilon = 0.30$—a value so large that it represents complete misunderstanding—can we observe a pronounced effect (e.g., the number of interactions required for a solution are increased by about 1 interaction). For simplicity and ease of interpretability, we therefore report our results with a null value for noise.

* Black values are default input values. Grey input values were used in the selected simulation experiments that we report.

allowed to reshare previously shared attributes before the discussion is considered stagnant, there is ample opportunity to scan the entire information space for all levels of information load, keeping accuracy high. Decision time increases linearly with total information load because each additional attribute has an equal chance of being selected. Random, therefore, demonstrates the starkest speed–accuracy tradeoff in relation to information load.

For the other strategies in Fig. 4, we instead observe a limited slowdown in decision speed as information load increases, which is due to strict selection. Agents try to maximize their impact on the other agents' decision preferences by sharing the attribute that is most supportive of a particular decision alternative. As a result, the team misses out on attributes that also support the desired alternative but to a lesser extent. Thus, sharing the most supportive attribute leads to a hit-or-miss dynamic: the attribute either changes the other agents' preferences, leading to a quick solution, or it does not change their preferences, so the agent keeps sharing the same attribute in vain.

TABLE II

TERMINOLOGY OF SIMULATION GRANULARITY

Unit	Attempts to solve decision task	Outcomes
interactions	-	1 interaction
step	1	Solved in X interactions OR Not Solved
run	$>$ 2000 steps	datapoint - median interactions, accuracy
simulation experiment	$S \times I \times D\|R\|W \times 10$ runs	Figure

Note: S, I, R, D, W signify the entire input space of strategies, information load, "individual" / total attributes, edges per node in memory network, and weight of novel information, respectively.

As the decision situation contains more attributes, the relative importance of one attribute decreases and, with it, its ability to change another agent's preference. Consequently, we see

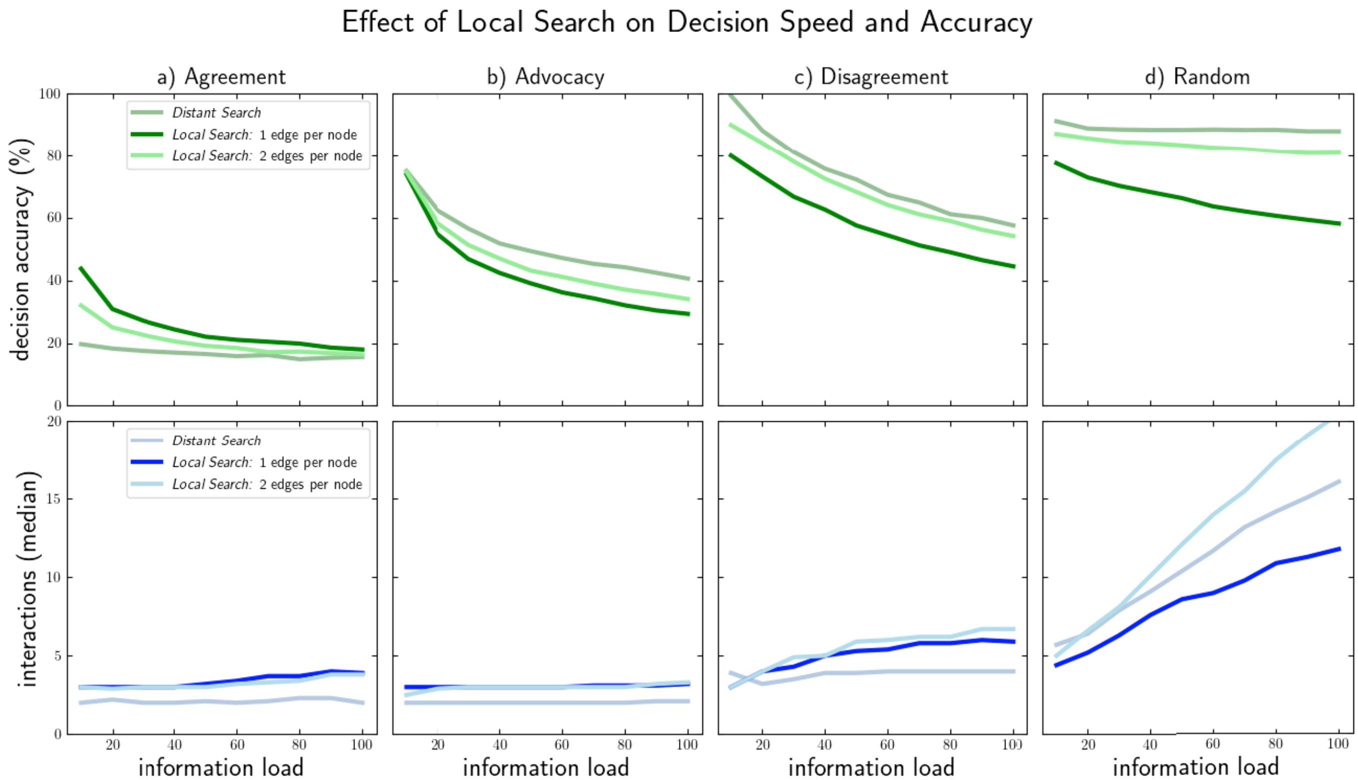

Fig. 4. Speed–accuracy tradeoffs for four information sharing strategies and three different levels of cognitive limitation. Green depicts decision-making accuracy and is expressed as the percentage of correctly solved problems. Blue depicts the time needed to reach consensus, expressed in the median number of interactions (based on repeated runs with different hidden profile problems). Agents share information that (a) agrees with the current team consensus, (b) advocates for their own belief about the correct decision alternative, (c) supports a different decision alternative than the last decision alternative that was argued for, or (d) at random.

a decrease in accuracy as the information load increases. Furthermore, each additional attribute causes a smaller increase in difficulty relative to the previous. Consequently, we see an exponential decrease in accuracy.

For Agreement [see Fig. 3(c)], we observe utter failure in decision-making. A team that starts out supporting an incorrect alternative and shares information that confirms incorrect alternative is not going to change its preference. Agreement, therefore, only succeeds if the attribute that supports the current team opinion most inadvertently supports the correct alternative even more, leading some agents to change their preference.

Conversely, Disagreement [see Fig. 3(d)] instructs agents to share contrarian information, meaning that the attribute to be shared can support every decision, except the last supported one. As a result, a wider variety of information attributes is brought up in discussion, lifting accuracy.

Advocacy, sharing information in line with one's own decision preference, holds the middle ground between agreement and disagreement. Because different agents advocate different decision alternatives, there is a greater variety in discussed attributes than in the Agreement strategy. However, since agents are tied to their own decision preferences, the total variety of information shared is less than in the Disagreement strategy.

C. Effect of Local Search on Decision Speed and Accuracy

Fig. 3 shows that, in terms of decision accuracy, Distant Search > Local Search (two edges per node) > Local Search (one edge per node), consistent with the principle that bounded rationality in search for information will impair decision-making. Distant Search indicates that agents have immediate access to all information in their memory, no matter how distant the information might be from their current position. Consequently, agents are able to fully search their information space and, thereby, select the best attribute to share with other agents given their information sharing strategy. Their most influential attribute (the global optimum) is readily available, leading, generally, to high accuracy and fast decisions. When agents can only pick the best attribute from a limited set of associated attributes in their local neighborhood N using a hill-climbing process, agents will share a variety of suboptimal attributes that are aligned with their strategy, before finding a local optimum attribute. As a result, decision accuracy drops for Advocacy, Disagreement, and Random. We see a reversal of this trend for the Agreement strategy, because Agreement is a losing strategy. When a losing strategy is badly executed, results improve.

We also observe that the accuracy of Local Search at two edges per node is higher than at one edge per node. Since a memory network consists of only 2–20 nodes, two edges per node means that agents can usually find the global optimum attribute within 1–3 hill-climbing steps. Associative memories with $d > 2$ return a near-identical result to $d = 2$. The main difference between Local Search at two edges per node and Distant Search is that in Distant Search, the agent finds the global optimum immediately, whereas in Local Search (two edges per node), it requires a few hill-climbing

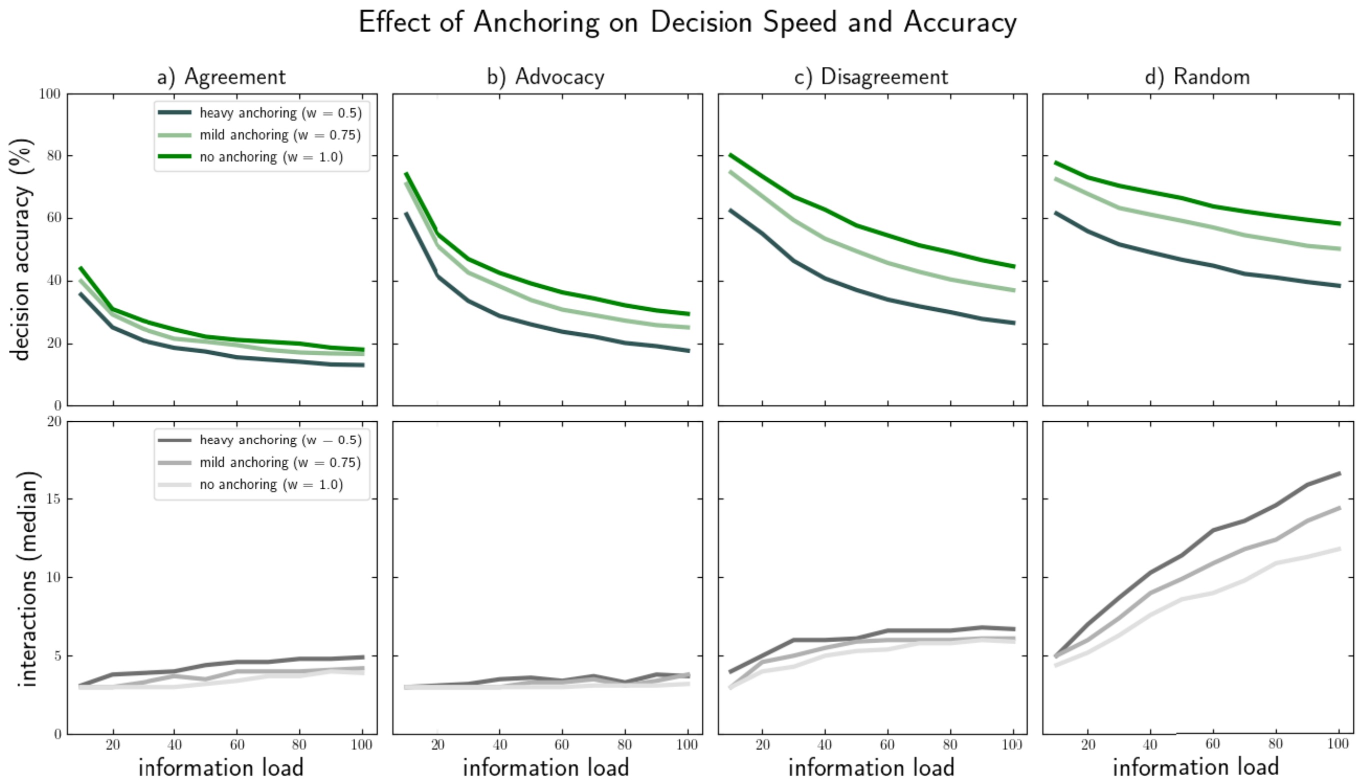

Fig. 5. Speed–accuracy tradeoffs for four information sharing strategies and three different levels of anchoring on initial information. Green depicts decision-making accuracy and is expressed as the percentage of correctly solved problems. Gray depicts the time needed to reach consensus, expressed in the median number of interactions (based on repeated runs with different hidden profile problems). Agents share information that (a) agrees with the current team consensus, (b) advocates for their own belief about the correct decision alternative, (c) supports a different decision alternative than the last decision alternative that was argued for, or (d) at random.

steps. With Local Search at one edge per node, connections between attributes follow a more tortuous path or may even be disconnected.

In terms of decision speed, we observe that Local Search at both two edges per node and one edge per node performs similar for Advocacy, Agreement, and Disagreement. Decision speed slows down in all cases relative to Distant Search. This suggests that Local Search at both two edges per node and one edge per node takes an equal number of turns to find a local optimum, but, in the case of two edges per node, the local optimum often coincides with the global optimum. The picture changes for Random where the one edge per node yields faster decision speeds. This can be explained by the sparsity of the network. If each node is, on average, connected to only one other node, it may occur that sections of the network are disconnected, which effectively shrinks the information space that can be explored. Random exploration in a restricted space is faster but also lowers accuracy if crucial attributes cannot be discovered.

D. Effect of Anchoring on Decision Speed and Accuracy

To account for the ways in which agents place a higher value on their initial individual information, which exerts an influence on their evaluation of information [26], [32], [103], we let agents weight attributes differently. Specifically, we let them discount attributes shared by other agents with weight w, where $w = 1$ represents equal weighting and $w = 0.5$ means that new information from others is weighted half as much as their initial information.

Fig. 5 shows the effects of anchoring on decision accuracy and speed. Not only does anchoring decrease decision accuracy but also decreases decision speed. When new information is discounted, more information is needed to sway an agent's opinion, and thus, more interactions are required.

Furthermore, anchoring decreases decision speed much more for Random than for the other strategies, while also substantially diminishing the accuracy advantage offered by Random in the absence of anchoring. The main difference between Random and the other strategies is that the other strategies direct search to select attributes for sharing. Conditional on the overall accuracy of the strategy—be it less accurate like Agreement or more accurate like Disagreement—these strategies enhance the strength of the attributes that are shared. Strong arguments can withstand discounting better than weak arguments.

Two opposing dynamics underlie these results. The higher the information load, the more the initial information an agent receives, and thus, the stronger its anchor, which suppresses decision quality. Simultaneously, a high information load produces larger memories to navigate, which means that more hill-climbing steps are needed to find a local optimum. In the process of taking these additional hill-climbing steps, more diverse attributes are shared, which counteracts the effect of

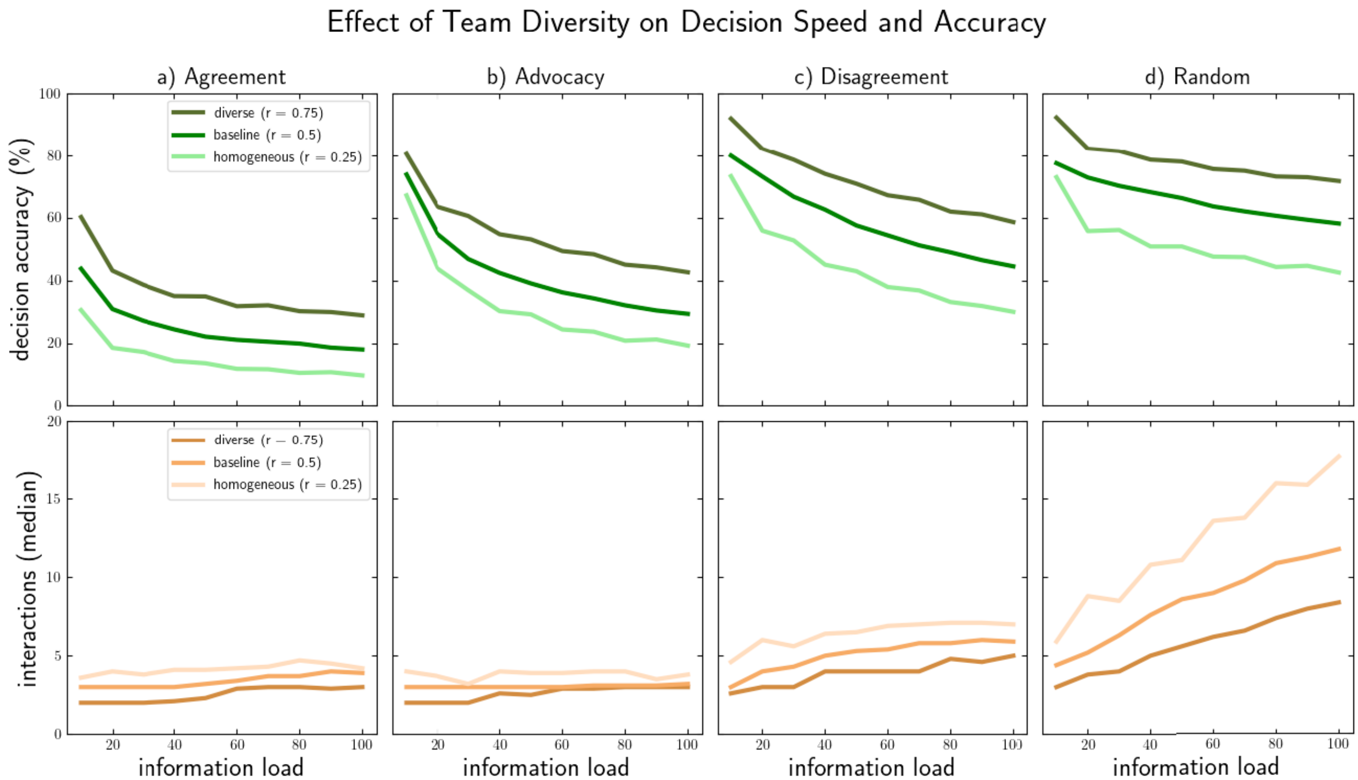

Fig. 6. Speed–accuracy tradeoffs for four information sharing strategies and three different ratios of individual/total information. Green depicts decision-making accuracy and is expressed as the percentage of correctly solved problems. Orange depicts the time needed to reach consensus, expressed in the median number of interactions (based on repeated runs with different hidden profile problems). Agents share information that (a) agrees with the current team consensus, (b) advocates for their own belief about the correct decision alternative, (c) supports a different decision alternative than the last decision alternative that was argued for, or (d) at random.

anchoring. These dynamics are relatively invariant to information load, as we can see from the equal drop in decision accuracy for both low and high information load.

E. Effect of Team Diversity on Decision Speed and Accuracy

Deep-level diversity in the information team members possesses typically enhances the intelligence of collective decision-making because it increases the pool of unique information available [15], [16]. To replicate this finding, we varied the proportion of individuals to total information. Fig. 6 shows that in teams in which members have more individual than common information, both decision speed and accuracy increase substantially.[1] The reverse is also true: in relatively homogeneous teams, decision-making takes longer and is less accurate. This effect is due to simple probability. Only individual information can change the opinion of other agents; sharing common information is a waste of time. When there is more individual information, agents are more likely to share individual information.

[1]The jagged lines for $r = 0.25$ and $r = 0.75$ are a result of rounding. Theoretically, if there is 25% common information, we should add 2.5 common attributes per ten attributes. However, splitting attributes is not possible. Therefore, we add three common attributes for the first increase in information load and two common attributes for the next increase. For example, when there are 30 attributes and 25% are common, meaning that 8/30 attributes are common. As a result, three agents receive two common attributes and two agents receive one common attribute. When there are 40 attributes and 25% are common, we have 10/40 common attributes, in this case, every agent gets two common attributes.

In addition, the strength of the hidden profile can play a reinforcing role. When there is a lot of common information, team members tend to prefer the same suboptimal choice. Thus, more members need to be converted. When team members have a lot of individual information, team members tend to prefer different alternatives, making it easier to form a majority [19].

F. Interaction Between Anchoring and Team Diversity on Decision Speed and Accuracy

Fig. 7 displays an interaction effect between anchoring and team diversity, exploring how bounded rationality manifests across different organizational environments. The impact of anchoring on decision accuracy increases with information load for diverse teams (dark green lines diverge), whereas the impact of anchoring decreases with information load for homogeneous teams (light green lines converge). This effect can be explained by the ratio of novel information to initial information.

When there is little information, say ten attributes, and agents have 25% individual information, then 3/10 attributes are individual. Any attribute that can change the team's opinion must be one of the three individual attributes. Changing the team's opinion is especially difficult when agents are anchored to their initial common information. Even if all three individual attributes are shared, that may not be enough to change the opinion of the

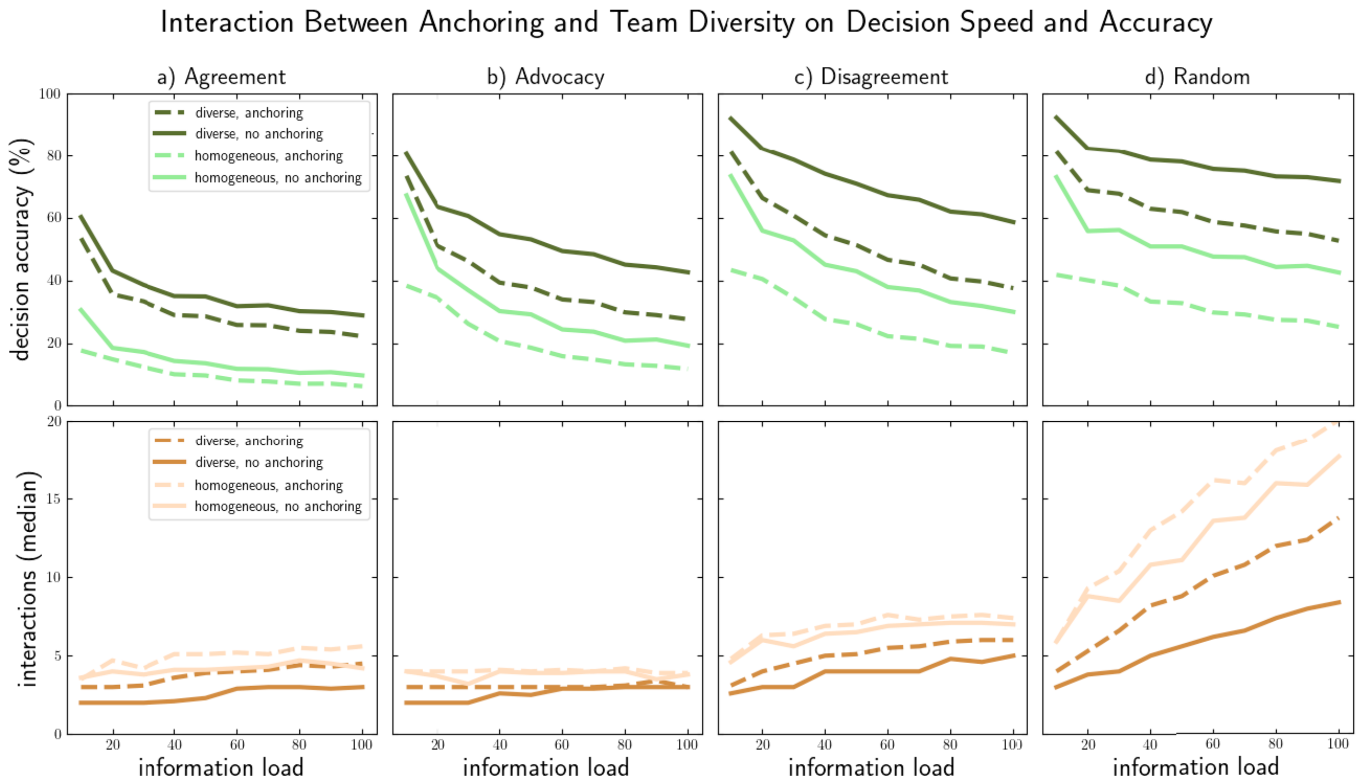

Fig. 7. Interaction between anchoring and team diversity. Diverse teams (dark colors, $r = 0.75$) suffer more from anchoring as information load increases, whereas homogeneous teams (light colors, $r = 0.25$) suffer less from anchoring as information load increases. Fixed parameters: $v = 1.0$—one edge per node in associative memory, $w = 1.0$—no anchoring bias, and $w = 0.5$—new information is weighted half as much as initial information.

agents whose initial information weighs double. As information load increases, more individual attributes from other agents may become available to compensate for the anchoring.

The reverse is true when agents have 75% initial information and there are only three common attributes. For the two agents who start with only individual information, even information that is common for the other three agents can change their opinion for the better. Moreover, when there is a low information load, it easy for agents to bring out all individual information. The combined amount of seven or three individual attributes has a good chance to outweigh the two anchored attributes. As information load increases, it becomes more difficult for teams to bring out all individual attributes, which decreases the team's ability to compensate for anchoring.

In sum, for homogeneous teams, more information means more opportunity to discover individual information, which can be used to convince anchored individuals to change their preferences. For diverse teams, who already have enough individual information to outweigh anchors, more information makes it more difficult to do so. Because some individual information weighs stronger than other individual information, it is more likely to be shared, yet this unshared individual information might make a difference in convincing anchored team members. This finding squares with meta-analytic work that notes how information coverage is more important than discussion focus [19].

G. Strength of Effects and Prediction Model

We conducted an analysis of covariance (ANCOVA) on decision accuracy (the percentage of total decisions that resulted in correct solutions) to determine the strength of effects (see Table III). Strategy accounted for 51% of variability, information load for 14%, the ratio of individual/total attributes for 20%, anchoring for 5%, and edges per node in memory for 2%, and we observed 8% residual variation. We note that individual-level factors cause much less variation than team-level factors, showing that collective decision-making has emergent properties beyond those of the individual team members.

The underlying statistical model is as follows:

$$a = S + I^e + R + W + D \pm \varepsilon \tag{5}$$

where a is decision accuracy, and capital letters signify the entire input space of Strategies S, information load I, ratio of individual to total attributes R, edges per node in memory network D, and weight of novel information W. ε is the error term. Since information load is a continuous variable, its effect on decision accuracy might be nonlinear, which we account for using a transformation with a power function. By minimizing the standard-error ($\varepsilon = 6.5$ for $e = 1$, and $\varepsilon = 6.1$ for $e = 0.1$), we obtained an exponent for information load. We assume linearity for R, W, and D since we only have three values for these continuous variables.

TABLE III
RESULTS OF ANCOVA ANALYSIS

factor	range	strength of effect	$\sum \varepsilon^2$	*df*	*F*	*p*-value
Strategy *	4 strategies	51%	1.909980e+06	3	15196	< 0.000001
Information load	$10 < I < 100$	14%	5.159021e+05	1	12313	< 0.000001
Individual / total information	$r = 0.25, 0.50, 0.75$	20%	7.383109e+05	1	17622	< 0.000001
Anchoring weight	$w = 0.5, 0.75, 1.0$	5%	2.099998e+05	1	5012	< 0.000001
Edges per node in memory	$d = 1, 2$	2%	8.457493e+04	1	2018	< 0.000001
Residual		8%	3.013158e+05	7192	-	< 0.000001

* One of the assumptions of ANCOVA is that each cell should contain the same number of cases. However, not all our simulations contain the same number of data points because we stopped simulating once the median number of interactions stabilized, resulting in a different number of steps for each strategy and level of information load. Therefore, we down-sampled the experiment data, such that each experiment contains 1,000 steps for statistical analysis.

IV. DISCUSSION

For decades, researchers interested in teams and their role in organizational decision-making have studied the effect of information sharing on decision outcomes. There is consensus that teams tend to share common information rather than individual information, which hampers their ability to make good decisions [4], [19]. Information sharing strategies can help counteract this tendency and, thereby, allow teams to better integrate their diverse information [4], [28]–[33]. Our research extends the current body of knowledge by investigating how information sharing strategies scale with increasing information load in teams of boundedly rational agents. To this end, we developed an agent-based simulation of a hidden profile decision task. Our simulation experiment both replicates prior findings from the literature and extends it to generate new insights.

A. Literature Comparison

Regarding decision quality, the ranking of information strategies (Disagreement > Advocacy > Agreement) confirms existing evidence [28], [29], [79], [104]–[106]. Extending this evidence, we add that Disagreement can increase decision accuracy by about 10%–50% compared with Advocacy. Agreement decreases decision accuracy by about 50%–70% compared with Advocacy. Intriguingly, Random information sharing had the best overall accuracy, which stresses the importance of fully searching and sharing the information space. Its usefulness for accuracy, however, was substantially impacted by bounded rationality, both in terms of local search and anchoring. To the extent that a perfect random strategy may not be cognitively feasible for humans to enact, its usefulness as a strategy may be more modest than our results suggest.

The results show that decision accuracy decreases exponentially with increasing information load in almost all cases. In general, decisions with fewer information attributes are easier to resolve than decisions with many. Each additional attribute has a smaller impact on the complexity of the decision situation than the previous one. In order to attempt to answer the question "what is advisable for teams faced with information overload?," it may be instructive to look at the case where accuracy does not drop exponentially: Random information sharing with distant search [see Fig. 4(d)]. In this strategy, agents are not pushed to search for information in any specific direction. Their search is neither confined by a local neighborhood search nor by an objective to share information supporting a specific alternative. Thus, if teams seek to increase their decision accuracy, they should particularly work to eliminate biases in information search.

In our experiment, we do not observe any simultaneous increase in both decision speed and accuracy, confirming evidence of a speed–accuracy tradeoff in decision-making [51]–[53]. We extend this evidence by noting how decision speed changes relatively little with increased information load for all strategies except for random, which demonstrated the starkest speed–accuracy tradeoff. This finding reveals the importance of having a strategy to direct the search for information. These strategies, in essence, produce a "quick solution or no solution" dynamic, whereby information load shapes whether a solution will be reached at all, more than it shapes the speed with which the decision is reached.

Our Anchoring simulation concurs with an overwhelming consensus that anchoring deteriorates decision quality, including accuracy [23]–[25]. Greitemeyer *et al.* [32] noted that individuals can be so anchored to their prediscussion preferences that it undermines the importance of information sharing. Xiao *et al.* [103] contend that when information is shared, it is not necessarily also used. They argue that the assumption that shared information is also used only holds for teams where information is equally distributed—a condition that our simulation model satisfies. Indeed, our simulation supports these findings if we decrease the weight w of novel information further. An important consideration that our simulation uniquely surfaces is how diverse teams become less able to compensate for anchoring as information load increases. Given that organizations assemble diverse teams for precisely this reason, and subject them to an abundance of information, this finding merits further unpacking.

B. Limitations and Future Research Directions

The current simulation model entails several simplifying assumptions, which stem from translating human behaviors into computer code. Nonetheless, idealized models are still valuable for building intuition about relationships between key variables and to make approximate order-of-magnitude calculations [107]. Details of individual-level cognition often have little impact on macrobehavior of a multimember cognitive system [108]. For instance, it may not matter exactly which contrarian information is shared as long as that contrarian information is shared. Still, there are model choices that deserve discussion.

The simulation follows a precise sequence of actions (share information, process information, vote, and repeat/end), which limits agent autonomy. While agents are autonomous in choosing what information to share from their memory network, how to interpret information using anchoring, and what vote to cast, they are bound to this sequence of actions. We see this forced sequence of actions as a natural constraint of being engaged in a discussion. However, incorporating further autonomy, for instance, in turn-taking dynamics, is something we consider for future work.

We model a system in which all information needed to make a "correct" decision is present. This decision, however, is only correct based on the team's definition of the decision situation (i.e., the available information and alternatives). Thus, we discard the possibility of new information being introduced in the middle of an ongoing conversation. In practice, information overload often means that the influx of new information is greater than the capacity to process it [35], [36]. Because our model considers a narrow time window of a single conversation, information overload in our model means that agents may not have the cognitive bandwidth to retrieve and share crucial information with other agents. In reality, information that could not be processed might be discarded altogether, instead of weakly considered by one person.

We assumed that a number of features are standard for all agents. All agents use the same strategy, all agents have the same amount of information, all information is equally relevant, and all agents take equal turns at sharing information. Making agents more heterogeneous might be more realistic, and the effects would certainly be interesting to study. However, heterogeneous agents would also produce more volatile simulation results that are harder to interpret. The current setup allows us to better estimate the focal individual- and team-level factors of local search, anchoring, and deep-level diversity. Still, nothing in the principles or design of the simulation prevents future incorporation of such heterogeneity.

We assume that all agents are connected and all agents receive all information that is sent. Even for a round-table discussion, these are perhaps optimistic assumptions. Some members might temporarily tune out, or there might be social factors that permanently disrupt the communication between two members. Certainly, this would be detrimental to decision speed, and if none of the other team members recovers the lost communication in later rounds, it would also affect accuracy.

The hill-climbing process that agents use to search for information to share is successful in the sense that it ensures that agents share a larger percentage of total information in simple problems than in complex problems. It also ensures that agents still share more diverse information when there is more information than when there is little. Thus, while the basic characteristics of information coverage in discussions are present in our model, how closely the information coverage of our agents matches that of actual humans is unknown.

Real teams also likely do not cast a vote as frequently as our simulated agents do. We had agents vote every round because tallying preferences allowed us to practically keep track of team consensus. Although real teams may not cast a vote after each argument, team members still likely keep track of who agrees with whom and maintain an implicit vote count. While the majority is the most preferred decision rule for teams [109], there may also be cases in which either collectives need to reach unanimity, such as in parliamentary procedures that require unanimous consent, or cases where people may never vote visibly and must infer the preferences of others, which could lead to a potential false consensus effect [110]. Such alternate arrangements for collective decision-making are certainly possible and could be studied in future work.

Our model excludes teams that begin with the right choice. As a consequence, the results may exaggerate the downsides of agreement and the benefits of disagreement or advocacy techniques, albeit without affecting the rank ordering of the strategies. Given widespread acknowledgment that agreement is a suboptimal decision approach [29], we see this as a minor issue. It is a necessary consequence of our decision to forego the limiting assumption in hidden profile research that private information is necessarily superior to public information in favor of generating information randomly in our simulation.

In sum, to improve our understanding of multiagent cognitive systems, future simulation-based work may benefit from relaxing some crude assumptions about human cognition inherent to the simulation presented here. Specifically, it would be interesting to investigate the effects of heterogeneity—such as the unequal distribution of information, unequal cognitive capacity, and unequal turn-taking—or simulate different decision-making contexts, such as where agents and information can more fluidly enter and exit throughout the duration of the decision-making process. Ideally, this would go hand-in-hand with giving agents greater autonomy over their actions, where the choice of action is based outcome of the previous action, such as a Markov process. It is also possible to further extend our model based on contextual factors related to human and organizational behavior that we did not consider. For instance, decision accuracy decreases with time pressure [111], which could undermine the value of a random strategy for teams, which is the slowest strategy, or, conversely, the value of disagreement may be more valuable in individualistic countries than our results reveal and less valuable in more collectivist countries [112].

V. Conclusion

As the amount of available information continues to proliferate, many of the most complex problems will be directed to teams for the solution. How can we ensure that teams are up to this task? Our work suggests that there is not only value in assembling teams with diverse information sets but also ensuring that they interact using strategies to guide their search through this information space. For organizations that value accurate decisions, they should search through the information space as thoroughly as possible, by attempting to mimic random sharing and surfacing any and all relevant information. For organizations that value speed in decision-making, they should leverage the wisdom of disagreement, which yields the highest accuracy with minimum cost to decision speed and greater robustness to bounded rationality. In this way, a richer understanding of how collectives cope with information overload may bring our collectives closer to intelligence in decision-making.

REFERENCES

[1] J. G. March and H. A. Simon, *Organizations*. Hoboken, NJ, USA: Wiley, 1958.

[2] H. A. Simon, *Administrative Behavior: A Study of Decision-Making Processes in Administrative Organisations*, 4th ed. New York, NY, USA: MacMillan, 1947.

[3] F. A. Csaszar and J. P. Eggers, "Organizational decision making: An information aggregation view," *Manage. Sci.*, vol. 59, no. 10, pp. 2257–2277, Oct. 2013.

[4] J. Mesmer-Magnus and L. Dechurch, "Information sharing and team performance: A meta-analysis," *IEEE Eng. Manag. Rev.*, vol. 40, no. 1, pp. 119–136, Mar. 2012.

[5] R. Westrum, "The study of information flow: A personal journey," *Saf. Sci.*, vol. 67, pp. 58–63, Aug. 2014.

[6] H. A. Simon, *Models of Man; Social and Rational*. New York, NY, USA: Wiley, 1957.

[7] H. A. Simon, "Bounded rationality and organizational learning," *Reflections SoL J.*, vol. 1, no. 2, pp. 17–27, Dec. 1999.

[8] G. Gavetti, D. Levinthal, and W. Ocasio, "Perspective–Neo-Carnegie: The Carnegie School's past, present, and reconstructing for the future," *Org. Sci.*, vol. 18, no. 3, pp. 523–536, Jun. 2007.

[9] R. S. Kudesia, "Mindfulness as metacognitive practice," *Acad. Manage. Rev.*, vol. 44, no. 2, pp. 405–423, Apr. 2019.

[10] D. Levinthal and C. Rerup, "Crossing an apparent chasm: Bridging mindful and less-mindful perspectives on organizational learning," *Org. Sci.*, vol. 17, no. 4, pp. 502–513, Aug. 2006.

[11] J. A. Nickerson and T. R. Zenger, "A knowledge-based theory of the firm, the problem-solving perspective," *Org. Sci.*, vol. 15, no. 6, pp. 617–632, Dec. 2004.

[12] R. K. Sawyer, "Group flow," in *Group Genius: The Creative Power of Collaboration*. New York, NY, USA: Basic Books, 2007.

[13] R. S. Tindale, T. Kameda, and V. B. Hinsz, "Group decision making," in *Sage Handbook of Social Psychology*, M. A. Hogg and J. Cooper, Eds. Newbury Park, CA, USA: Sage, 2003, pp. 381–403.

[14] S. Wuchty, B. F. Jones, and B. Uzzi, "The increasing dominance of teams in production of knowledge," *Science*, vol. 316, no. 5827, pp. 1036–1039, May 2007.

[15] U. R. Hülsheger, N. Anderson, and J. F. Salgado, "Team-level predictors of innovation at work: A comprehensive meta-analysis spanning three decades of research," *J. Appl. Psychol.*, vol. 94, no. 5, pp. 1128–1145, 2009.

[16] J. Wang, G. H. Cheng, T. Chen, and K. Leung, "Team creativity/innovation in culturally diverse teams: A meta-analysis," *J. Organizational Behav.*, vol. 40, no. 6, pp. 693–708, Jul. 2019.

[17] G. Stasser and W. Titus, "Pooling of unshared information in group decision making: Biased information sampling during discussion," *J. Personality Social Psychol.*, vol. 48, no. 6, pp. 1467–1478, 1985.

[18] G. Stasser and Z. Birchmeier, "Group creativity and collective choice," in *Group Creativity: Innovation Through Collaboration*. 2003, pp. 85–109.

[19] L. Lu, Y. C. Yuan, and P. L. McLeod, "Twenty-five years of hidden profiles in group decision making: A meta-analysis," *Personality Social Psychol. Rev.*, vol. 16, no. 1, pp. 54–75, Feb. 2012.

[20] M. Baer, K. T. Dirks, and J. A. Nickerson, "Microfoundations of strategic problem formulation," *Strategic Manage. J.*, vol. 34, no. 2, pp. 197–214, Feb. 2013.

[21] R. M. Cyert and J. G. March, *A Behavioral Theory of the Firm*. Upper Saddle River, NJ, USA: Prentice-Hall, 1963.

[22] G. Gavetti and D. Levinthal, "Looking forward and looking backward: Cognitive and experiential search," *Administ. Sci. Quart.*, vol. 45, no. 1, p. 113, Mar. 2000.

[23] E. Dane, "Reconsidering the trade-off between expertise and flexibility: A cognitive entrenchment perspective," *Acad. Manage. Rev.*, vol. 35, no. 4, pp. 579–603, Oct. 2010.

[24] A. Furnham and H. C. Boo, "A literature review of the anchoring effect," *J. Socio-Econ.*, vol. 40, no. 1, pp. 35–42, Feb. 2011.

[25] S. M. Smith, "The constraining effects of initial ideas," in *Group Creativity: Innovation Through Collaboration*. Oxford, U.K.: Oxford Univ. Press, 2003, pp. 15–31.

[26] G. M. Wittenbaum, A. B. Hollingshead, and I. C. Botero, "From cooperative to motivated information sharing in groups: Moving beyond the hidden profile paradigm," *Commun. Monographs*, vol. 71, no. 3, pp. 286–310, Sep. 2004.

[27] B. A. Nijstad and C. K. W. De Dreu, "Motivated information processing in organizational teams: Progress, puzzles, and prospects," *Res. Org. Behav.*, vol. 32, pp. 87–111, Jan. 2012.

[28] D. M. Schweiger, W. R. Sandberg, and P. L. Rechner, "Experiential effects of dialectical inquiry, devil's advocacy and consensus approaches to strategic decision making," *Acad. Manage. J.*, vol. 32, no. 4, pp. 745–772, Dec. 1989.

[29] R. A. Cosier and C. R. Schwenk, "Agreement and thinking alike: Ingredients for poor decisions," *Acad. Manage. Perspect.*, vol. 4, no. 1, pp. 69–74, Jan. 1990.

[30] S. Schulz-Hardt, M. Jochims, and D. Frey, "Productive conflict in group decision making: Genuine and contrived dissent as strategies to counteract biased information seeking," *Org. Behav. Human Decis. Processes*, vol. 88, no. 2, pp. 563–586, Jul. 2002.

[31] G. A. Okhuysen and K. M. Eisenhardt, "Integrating knowledge in groups: How formal interventions enable flexibility," *Org. Sci.*, vol. 13, no. 4, pp. 370–386, Aug. 2002.

[32] T. Greitemeyer, S. Schulz-Hardt, F. C. Brodbeck, and D. Frey, "Information sampling and group decision making: The effects of an advocacy decision procedure and task experience," *J. Exp. Psychol., Appl.*, vol. 12, no. 1, pp. 31–42, 2006.

[33] M.-H. Tsai and C. Bendersky, "The pursuit of information sharing: Expressing task conflicts as debates vs. disagreements increases perceived receptivity to dissenting opinions in groups," *Org. Sci.*, vol. 27, no. 1, pp. 141–156, Dec. 2015.

[34] P. G. Roetzel, "Information overload in the information age: A review of the literature from business administration, business psychology, and related disciplines with a bibliometric approach and framework development," *Bus. Res.*, vol. 12, no. 2, pp. 479–522, Dec. 2019.

[35] M. J. Eppler and J. Mengis, "The concept of information overload: A review of literature from organization science, accounting, marketing, MIS, and related disciplines," *Inf. Soc.*, vol. 20, no. 5, pp. 325–344, Nov. 2004.

[36] A. Edmunds and A. Morris, "The problem of information overload in business organisations: A review of the literature," *Int. J. Inf. Manage.*, vol. 20, no. 1, pp. 17–28, Feb. 2000.

[37] E. M. Hallowell, "Overloaded circuits: Why smart people underperform," *Harvard Bus. Rev.*, vol. 83, no. 1, pp. 54–62, 2005.

[38] D. J. Levitin, *The Organized Mind: Thinking Straight in the Age of Information Overload*. Baltimore, MD, USA: Penguin, 2014.

[39] G. George, M. R. Haas, and A. Pentland, "From the Editors: Big data and management," *Acad. Manage. J.*, vol. 57, no. 2, pp. 321–326, 2014.

[40] A. McAfee, E. Brynjolfsson, T. H. Davenport, D. J. Patil, and D. Barton, "Big data: The management revolution," *Harvard Bus. Rev.*, vol. 90, no. 10, pp. 60–68, 2012.

[41] W. Gao, Z. Liu, Q. Guo, and X. Li, "The dark side of ubiquitous connectivity in smartphone-based SNS: An integrated model from information perspective," *Comput. Hum. Behav.*, vol. 84, pp. 185–193, Jul. 2018.

[42] H. Fang, "Managing data lakes in big data era: What's a data lake and why has it became popular in data management ecosystem," in *Proc. IEEE Int. Conf. Cyber Technol. Autom., Control, Intell. Syst. (CYBER)*, Jun. 2015, pp. 820–824.

[43] Y. Duan, J. S. Edwards, and Y. K. Dwivedi, "Artificial intelligence for decision making in the era of big data–evolution, challenges and research agenda," *Int. J. Inf. Manage.*, vol. 48, pp. 63–71, Oct. 2019.

[44] D. Bawden and L. Robinson, "The dark side of information: Overload, anxiety and other paradoxes and pathologies," *J. Inf. Sci.*, vol. 35, no. 2, pp. 180–191, Apr. 2009.

[45] H. A. Simon, "Designing organizations for an information-rich world," in *Designing Organizations for an Information-Rich World*, M. Greenberger, Ed. Baltimore, MD, USA: The Johns Hopkins Univ. Press, 1971, pp. 38–52.

[46] T. H. Davenport and J. C. Beck, *The Attention Economy: Understanding the New Currency of Business*. Boston, MA, USA: Harvard Business School Press, 2001.

[47] J. B. Spira and C. Burke, *Intel's War on Information Overload: A Case Study*. New York, NY, USA: Basex, 2009.

[48] V. B. Hinsz, R. S. Tindale, and D. A. Vollrath, "The emerging conceptualization of groups as information processors," *Psychol. Bull.*, vol. 121, no. 1, pp. 43–64, 1997.

[49] E. Whelan and R. Teigland, "Transactive memory systems as a collective filter for mitigating information overload in digitally enabled organizational groups," *Inf. Org.*, vol. 23, no. 3, pp. 177–197, Jul. 2013.

[50] J. R. Winquist and J. L. Larson, Jr., "Information pooling: When it impacts group decision making," *J. Personality Social Psychol.*, vol. 74, no. 2, pp. 371–377, 1998.

[51] W. A. Wickelgren, "Speed-accuracy tradeoff and information processing dynamics," *Acta Psychologica*, vol. 41, no. 1, pp. 67–85, Feb. 1977.

[52] L. Chittka, P. Skorupski, and N. E. Raine, "Speed–accuracy tradeoffs in animal decision making," *Trends Ecol. Evol.*, vol. 24, no. 7, pp. 400–407, Jul. 2009.

[53] J. A. R. Marshall, A. Dornhaus, N. R. Franks, and T. Kovacs, "Noise, cost and speed-accuracy trade-offs: Decision-making in a decentralized system," *J. Roy. Soc. Interface*, vol. 3, no. 7, pp. 243–254, Apr. 2006.

[54] G. Fioretti, "Agent-based simulation models in organization science," *Org. Res. Methods*, vol. 16, no. 2, pp. 227–242, Apr. 2013.

[55] H. A. Simon, "Invariants of human behavior," *Annu. Rev. Psychol.*, vol. 41, no. 1, pp. 1–20, Jan. 1990.

[56] S. Frederick, "Cognitive reflection and decision making," *J. Econ. Perspect.*, vol. 19, no. 4, pp. 25–42, Nov. 2005.

[57] B. Edmonds and R. Meyer, *Simulating Social Complexity*. Springer, 2015. [Online]. Available: https://www.springer.com/us/book/9783319669472

[58] P. Puranam, N. Stieglitz, M. Osman, and M. M. Pillutla, "Modelling bounded rationality in organizations: Progress and prospects," *Acad. Manage. Ann.*, vol. 9, no. 1, pp. 337–392, 2015.

[59] E. Bonabeau, "Agent-based modeling: Methods and techniques for simulating human systems," *Proc. Nat. Acad. Sci. USA*, vol. 99, no. 3, pp. 7280–7287, May 2002.

[60] R. M. Burton and B. Obel, "Computational modeling for what-is, what-might-be, and what-should-be studies–and triangulation," *Org. Sci.*, vol. 22, no. 5, pp. 1195–1202, Oct. 2011.

[61] H. A. Simon, "Information-processing theory of human problem solving," in *Handbook Learn. Cognit. processes*, vol. 5, pp. 271–295, 1978.

[62] C. Shannon and W. Weaver, *The Mathematical Theory of Communication*, 10th ed. Urbana, IL, USA: The Univ. of Illinois Press, 1949.

[63] L. Bao and J. C. Fritchman, "Information of complex systems and applications in agent based modeling," *Sci. Rep.*, vol. 8, no. 1, p. 6177, Dec. 2018.

[64] J. Holler, K. H. Kendrick, M. Casillas, and S. C. Levinson, Eds., *Turn-Taking in Human Communicative Interaction* (Frontiers Research Topics). Lausanne, Switzerland: Frontiers Media SA, 2016.

[65] J. Larrue and A. Trognon, "Organization of turn-taking and mechanisms for turn-taking repairs in a chaired meeting," *J. Pragmatics*, vol. 19, no. 2, pp. 177–196, Feb. 1993.

[66] E. Careau, C. Vincent, and B. R. Swaine, "Consensus group session of experts to describe interprofessional collaboration processes in team meetings," *J. Interprofessional Care*, vol. 25, no. 4, pp. 299–301, Jul. 2011.

[67] R. W. Proctor and K.-P. L. Vu, "Human Information Processing," in *Encyclopedia Sciences of Learning*, N. M. Seel, Ed. Boston, MA, USA: Springer, 2012, pp. 1458–1460.

[68] S. G. Green and T. D. Taber, "The effects of three social decision schemes on decision group process," *Organizational Behav. Hum. Perform.*, vol. 25, no. 1, pp. 97–106, Feb. 1980.

[69] S. G. Sohrab, M. J. Waller, and S. Kaplan, "Exploring the hidden-profile paradigm: A literature review and analysis," *Small Group Res.*, vol. 46, no. 5, pp. 489–535, Oct. 2015.

[70] G. Gigerenzer and R. Selten, *Bounded Rationality: The Adaptive Toolbox*. Cambridge, MA, USA: MIT Press, 2002.

[71] D. C. Knill and A. Pouget, "The Bayesian brain: The role of uncertainty in neural coding and computation," *Trends Neurosci.*, vol. 27, no. 12, pp. 712–719, Dec. 2004.

[72] S. Glasauer, "Sequential Bayesian updating as a model for human perception," in *Progress in Brain Research* (Mathematical Modelling in Motor Neuroscience: State of the Art and Translation to the Clinic. Gaze Orienting Mechanisms and Disease), vol. 249, S. Ramat and A. G. Shaikh, Eds. Amsterdam, The Netherlands: Elsevier, Jan. 2019, ch. 1, pp. 3–18.

[73] W. Ocasio, "Towards an attention based view of the firm," *Strategic Manage. J.*, vol. 18, no. S1, pp. 187–206, Jul. 1997.

[74] B. J. Olson, S. Parayitam, and Y. Bao, "Strategic decision making: The effects of cognitive diversity, conflict, and trust on decision outcomes," *J. Manage.*, vol. 33, no. 2, pp. 196–222, Apr. 2007.

[75] M. S. Feldman and J. G. March, "Information in organizations as signal and symbol," *Admin. Sci. Quart.*, vol. 26, no. 2, pp. 171–176, Jun. 1981.

[76] Y. Dabaghian, V. L. Brandt, and L. M. Frank, "Reconceiving the hippocampal map as a topological template," *eLife*, vol. 3, Aug. 2014.

[77] A. Babichev, S. Cheng, and Y. A. Dabaghian, "Topological schemas of cognitive maps and spatial learning," *Frontiers Comput. Neurosci.*, vol. 10, Mar. 2016, p. 18.

[78] S. T. Mueller, B. S. Perelman, and B. G. Simpkins, "Pathfinding in the cognitive map: Network models of mechanisms for search and planning," *Biologically Inspired Cognit. Archit.*, vol. 5, pp. 94–111, Jul. 2013.

[79] M. E. Turner and A. R. Pratkanis, "Twenty-five years of groupthink theory and research: Lessons from the evaluation of a theory," *Org. Behav. Hum. Decis. Processes*, vol. 73, nos. 2–3, pp. 105–115, Feb. 1998.

[80] S. Schulz-Hardt, D. Frey, C. Lüthgens, and S. Moscovici, "Biased information search in group decision making," *J. Personality Social Psychol.*, vol. 78, no. 4, pp. 655–669, 2000.

[81] A. R. Dennis, "Information exchange and use in small group decision making," *Small Group Res.*, vol. 27, no. 4, pp. 532–550, Nov. 1996.

[82] D. Gigone and R. Hastie, "The common knowledge effect: Information sharing and group judgment," *J. Personality Social Psychol.*, vol. 65, no. 5, pp. 959–974, 1993.

[83] H. Mercier and D. Sperber, "Why do humans reason? Arguments for an argumentative theory," *Behav. Brain Sci.*, vol. 34, no. 2, pp. 57–74, Apr. 2011.

[84] R. L. Priem and K. H. Price, "Process and outcome expectations for the dialectical inquiry, devil's advocacy, and consensus techniques of strategic decision making," *Group Org. Stud.*, vol. 16, no. 2, pp. 206–225, Jun. 1991.

[85] D. M. Schweiger and P. A. Finger, "The comparative effectiveness of dialectical inquiry and devil's advocacy: The impact of task biases on previous research findings," *Strategic Manage. J.*, vol. 5, no. 4, pp. 335–350, Oct. 1984.

[86] C. R. Schwenk, "Effects of devil's advocacy and dialectical inquiry on decision making: A meta-analysis," *Organizational Behav. Human Decis. Processes*, vol. 47, no. 1, pp. 161–176, Oct. 1990.

[87] D. K. Simonton, "Creativity as blind variation and selective retention: Is the creative process Darwinian?" *Psychol. Inquiry*, vol. 10, no. 3, pp. 309–328, Oct. 1999.

[88] K. E. Weick, *The Social Psychology of Organizing*, 2nd ed. Reading, MA, USA: Addison-Wesley, 1979.

[89] B. M. Staw, "An evolutionary approach to creativity and innovation," in *Innovation and Creativity at Work: Psychological and Organizational Strategies*, M. W. West and J. L. Farr, Eds. Hoboken, NJ, USA: Wiley, 1990, pp. 287–308.

[90] R. Libby and R. K. Blashfield, "Performance of a composite as a function of the number of judges," *Org. Behav. Hum. Perform.*, vol. 21, no. 2, pp. 121–129, Apr. 1978.

[91] P. Yetton and P. Bottger, "The relationships among group size, member ability, social decision schemes, and performance," *Org. Behav. Hum. Perform.*, vol. 32, no. 2, pp. 145–159, Oct. 1983.

[92] P. R. Laughlin and E. Al, "Group size, member ability, and social decision schemes on an intellective task," *J. Personality Social Psychol.*, vol. 31, no. 3, pp. 522–535, 1975.

[93] J. A. Krems and J. Wilkes, "Why are conversations limited to about four people? A theoretical exploration of the conversation size constraint," *Evol. Hum. Behav.*, vol. 40, no. 2, pp. 140–147, Mar. 2019.

[94] Z. Simsek, J. F. Veiga, M. H. Lubatkin, and R. N. Dino, "Modeling the multilevel determinants of top management team behavioral integration," *Acad. Manage. J.*, vol. 48, no. 1, pp. 69–84, Feb. 2005.

[95] D. Bang and C. D. Frith, "Making better decisions in groups," *Roy. Soc. Open Sci.*, vol. 4, no. 8, Aug. 2017, Art. no. 170193.

[96] D. Laming, "Serial position curves in free recall," *Psychol. Rev.*, vol. 117, no. 1, pp. 93–133, 2010.

[97] D. M. Oppenheimer, R. A. LeBoeuf, and N. T. Brewer, "Anchors aweigh: A demonstration of cross-modality anchoring and magnitude priming," *Cognition*, vol. 106, no. 1, pp. 13–26, Jan. 2008.

[98] M. Hilbert, "Toward a synthesis of cognitive biases: How noisy information processing can bias human decision making," *Psychol. Bull.*, vol. 138, no. 2, pp. 211–237, 2012.

[99] J. Bendor, R. M. Kramer, and S. Stout, “When in doubt...: Cooperation in a noisy Prisoner’s dilemma,” *J. Conflict Resolution*, vol. 35, no. 4, pp. 691–719, Dec. 1991.
[100] G. van Rossum and F. L. Drake, Jr., *Python Reference Manual*. Amsterdam, The Netherlands: Centrum voor Wiskunde en Informatica, 1995.
[101] D. Merkel, “Docker: lightweight linux containers for consistent development and deployment,” *Linux J.*, vol. 2014, no. 239, p. 2, 2014.
[102] P. Virtanen *et al.*, “SciPy 1.0: Fundamental algorithms for scientific computing in Python,” *Nature Methods*, vol. 17, pp. 261–272, 2020.
[103] Y. Xiao, H. Zhang, and T. M. Basadur, “Does information sharing always improve team decision making? An examination of the hidden profile condition in new product development,” *J. Bus. Res.*, vol. 69, no. 2, pp. 587–595, Feb. 2016.
[104] C. R. Schwenk and R. A. Cosier, “Effects of the expert, devil’s advocate, and dialectical inquiry methods on prediction performance,” *Org. Behav. Hum. Perform.*, vol. 26, no. 3, pp. 409–424, Dec. 1980.
[105] C. R. Schwenk and R. A. Cosier, “Effects of consensus and devil’s advocacy on strategic decision-making,” *J. Appl. Social Psychol.*, vol. 23, no. 2, pp. 126–139, Jan. 1993.
[106] T. T. Herbert and R. W. Estes, “Improving executive decisions by formalizing dissent: The corporate devil’s advocate,” *Acad. Manage. Rev.*, vol. 2, no. 4, pp. 662–667, Oct. 1977.
[107] R. Axtell, “Why agents? On the varied motivations for agent computing in the social sciences,” Brookings Inst., Center Social Econ. Dyn., Washington, DC, USA, Working Paper 17, 2000, p. 23.
[108] S. Bhattacharyya and S. Ohlsson, “Social creativity as a function of agent cognition and network properties: A computer model,” *Social Netw.*, vol. 32, no. 4, pp. 263–278, Oct. 2010.
[109] R. Hastie and T. Kameda, “The robust beauty of majority rules in group decisions,” *Psychol. Rev.*, vol. 112, no. 2, pp. 494–508, 2005.
[110] L. Ross, D. Greene, and P. House, “The ‘false consensus effect’: An egocentric bias in social perception and attribution processes,” *J. Exp. Social Psychol.*, vol. 13, no. 3, pp. 279–301, May 1977.
[111] J. M. Bowman and G. M. Wittenbaum, “Time pressure affects process and performance in hidden-profile groups,” *Small Group Res.*, vol. 43, no. 3, pp. 295–314, Jun. 2012.
[112] A. Nahavandi and E. Aranda, “Restructuring teams for the re-engineered organization,” *Acad. Manage. Perspect.*, vol. 8, no. 4, pp. 58–68, Nov. 1994.

Dirk-Jan van Veen (Graduate Student Member, IEEE) received the M.Sc. degree in applied geophysics from ETH Zürich, Zürich, Switzerland, in 2016. He is currently pursuing the Ph.D. degree with the ETH Zürich, Future Resilient Systems Lab, Singapore-ETH Centre, Singapore.

His research interests include distributed cognitive systems, heuristic search algorithms, neural networks, and social physics.

Ravi S. Kudesia received the Ph.D. degree in management from Washington University in St. Louis, St. Louis, MO, USA, in 2017.

He is currently an Assistant Professor of human resource management with the Fox School of Business, Temple University, Philadelphia, PA, USA. He studies human organizing and how people can organize more mindfully. His research concerns three processes: attention, interpretation, and energy. Using experiments, qualitative case studies, and agent-based models, he traces how these processes transfer across people as they organize into collectives and how these collectives solve problems and make sense of their environments. His research has appeared in several journals, including *Academy of Management Review*, *Organizational Behavior and Human Decision Processes*, and *Safety Science*.

Hans R. Heinimann (Member, IEEE) received the Ph.D. degree from ETH Zürich, Zürich, Switzerland, in 1986.

He held visiting professor positions at Oregon State University, Corvallis, OR, USA, The University of Tokyo, Tokyo, Japan, and The Australian National University, Canberra ACT, Australia. He has been a Professor of forest engineering with ETH Zürich, Zürich, Switzerland, since 1991. From 2004 to 2009, he was a fellow with the Collegium Helveticum, Zürich, a center of advanced studies jointly sponsored by ETH Zürich and the University of Zurich, Zürich, focusing on cross-disciplinary research. He held various leadership positions in university bodies, such as the head of the department, the director of studies, the head of the institute, and the head of the faculty recruiting committees, among others. He was the Pro-rector for Education of ETH Zürich from 2007 to 2013, where he is also a Founding Member of the ETH Risk Centre, of which he was the chairman from 2011 to 2013. From 2014 to 2019, he was the Programme Director of the Future Resilient Systems Lab, Singapore-ETH Centre, Singapore, where he is still involved as a Principal Investigator.

Dr. Heinimann is also a member of several scientific and professional societies, including the Swiss Academy of Engineering Sciences (SATW). In 2017, he was awarded the Distinguished Service Award by the IUFRO in recognition for his outstanding and valuable services to the organization. He has been actively involved in the international scientific community and coordinated the Forest Operations Engineering and Management Division of the International Union of Forest Research Organisations (IUFRO).