Trends in
Cognitive Sciences

CellPress

Feature Review

The network science of collective intelligence

Damon Centola [1,2,3,4,*]

In the last few years, breakthroughs in computational and experimental techniques have produced several key discoveries in the science of networks and human collective intelligence. This review presents the latest scientific findings from two key fields of research: collective problem-solving and the wisdom of the crowd. I demonstrate the core theoretical tensions separating these research traditions and show how recent findings offer a new synthesis for understanding how network dynamics alter collective intelligence, both positively and negatively. I conclude by highlighting current theoretical problems at the forefront of research on networked collective intelligence, as well as vital public policy challenges that require new research efforts.

Highlights

The essential puzzle of collective intelligence is whether the collective judgment from a group of people will outperform a smart individual reasoning alone. Recent computational and experimental studies have led to breakthroughs in two of the primary fields of networked collective intelligence: collective problem-solving and the wisdom of the crowd.

Collective problem-solving typically addresses the optimal design for communication networks within organizations. The key network property governing problem-solving outcomes is informational efficiency (i.e., average simple path length).

The wisdom of the crowd shows that the average response from a large group of novices can be more accurate than the opinions of individual experts. The key network property governing the wisdom of the crowd is network centralization.

The enigma of collective intelligence

Perhaps one of the most important questions in the study of human behavior is whether social interaction leads people toward greater intelligence and creativity, or whether it constrains individual intelligence, reducing productivity and undermining human potential. Dating back to antiquity, Aristotle likened human groups to insect colonies and analogized the creative patterns observed in animal societies to the virtues of human collective intelligence [1]. By contrast, the English philosopher Thomas Hobbes had a decidedly anti-Aristotelian view of human social reasoning, viewing collective behavior as a self-destructive form of mob violence, and suggested that the only way to save the human crowd from itself was through the leadership of a strong monarch [2,3].

Today, the core puzzle of collective intelligence has become synonymous with the question of whether the collective judgment of a group of people (which, in the broad scope of the literature on collective intelligence may include the statistical aggregation of independent opinions, passive information-sharing among peers, active solution-sharing on research teams, and considered and/or passionate deliberation within political bodies) will outperform a smart individual (or computer) reasoning alone [4–6]. There are many examples of collective decision-making gone horribly awry, for instance, the infamous case of groupthink at NASA leading to the Space Shuttle *Columbia* disaster [7]. However, there are also examples of collective reasoning processes significantly improving the quality of people's judgments, for instance, identifying the dangers of climate change [8] and improving the quality of scientific and medical judgments [9–11]. The study of collective intelligence impacts a remarkable range of disciplines, including finance [12], national elections [13,14], sports betting [15], medical decision-making [16–18], marketing [19], engineering and data sciences [5,20,21], and animal behavior [22–24], to name just a few.

This review offers a new perspective showing how network science may provide a unifying framework for research on collective intelligence. I focus on recent insights from two key fields of empirical study spearheading new approaches to the network dynamics of group performance, namely: 'collective problem-solving' and 'the wisdom of the crowd'. Here, I present a synthesis of this work that broaches a generalized understanding of how social networks influence collective intelligence.

[1]Annenberg School for Communication, University of Pennsylvania, Philadelphia, PA 19104, USA
[2]School of Engineering and Applied Sciences, University of Pennsylvania, Philadelphia, PA 19104, USA
[3]Department of Sociology, University of Pennsylvania, Philadelphia, PA 19104, USA
[4]Network Dynamics Group, University of Pennsylvania, Philadelphia, PA 19104, USA

*Correspondence:
dcentola@asc.upenn.edu (D. Centola).

Two fields, three principles

Collective problem-solving examines situations in which communication networks among team members are necessary for discovering innovative solutions. The primary concern for researchers in this field is determining how the structure, or topology, of communication networks among problem-solvers may improve (or even optimize) the quality of the solutions they discover. Work in this field can be applied to a broad range of topics in biology [25], military research [26], and many other areas [27]. However, recent network applications have typically focused on organizational studies [28], with an emphasis on identifying team-networking strategies that can improve the speed of scientific discovery, enhance the creativity of team members, and accelerate the overall capacity for technological and artistic innovation [29,30]. Group consensus in collective problem-solving emerges endogenously through peer information-sharing networks that reveal the performance of solutions used by competing individuals and teams. A key assumption in this field is that there is a discoverable optimal answer within the solution space and that peers typically adopt performance-increasing solutions from one another (with allowances for noise [21,30]).

The wisdom of the crowd focuses largely on estimation tasks, such as predicting stock prices or forecasting climate change, and is driven by the insight that the average 'group estimate' taken from a large number of independent individuals is typically more accurate than any of the individual group members' estimates [7]. Unlike collective problem-solving, the group estimate in the wisdom of the crowd does not emerge via endogenous consensus but is instead derived using an aggregation method to characterize the central tendency of the group (for example, calculating the mean of the individual estimates). Research on the wisdom of crowds also includes the study of discrete choice problems (such as choosing between presidential candidates), for which a variety of different aggregation methods have been used (such as majority rule, plurality rule, quorum rule, etc.) to determine the most effective way of sampling the population to improve the quality of their collective choice [10,16,31].

Both kinds of wisdom of the crowd problem (continuous estimates and discrete choices) assume that: (i) there is a correct answer, and (ii) even though individuals in the population may not know the right answer, it can nevertheless be found at the group level by using aggregation methods to combine group members' guesses. Individuals most often provide their best guess to optimize their individual performance on the task (i.e., typically unaware of how their guess will factor into the aggregation process that produces a collective judgment; however, see also [32–34]). Thus, group intelligence in the wisdom of the crowd is evaluated based on the collective response that emerges from the chosen aggregation method, irrespective of the intelligence of individual group members' responses. The theory of crowd wisdom has been broadly applied to everything from organizational decision-making [35] to medical reasoning [10,36] to democratic processes [37,38]. Importantly, unlike research on collective problem-solving, in which communication networks are typically seen as necessary for improving collective performance, social networks are not necessarily seen as a requirement for the wisdom of the crowd [4,7,31,37]. A large body of research argues that social influence can unexpectedly compromise the wisdom of the crowd [7,39,40], suggesting that collective intelligence can be increased by eliminating peer-to-peer communications.

Consequently, these two distinct fields of research on collective intelligence offer different expectations about the effects of communication networks on group performance. Historically, these differing expectations have not represented a scientific puzzle because these fields of scholarship were not in close dialogue with one another. Each field uses its own models of individual and collective behavior and makes different assumptions about the sequence of individual decision-making and the informational content that propagates through communication networks. However, the recent explosion of computational and experimental research in both fields has established a dialogue between these research areas under the shared rubric of 'collective intelligence'. From this dialogue, three key insights emerge that enable a productive synthesis of these fields:

i) The first insight is that the key network property that determines group performance in collective problem-solving tasks is informational efficiency, typically operationalized in terms of average simple path length [41,42] (i.e., the typical number of 'steps' required for information to traverse a network).
ii) The second insight is that the effect of network structure on collective problem-solving is governed by the complexity of the task [25]. Counterintuitively, collective intelligence on complex tasks is improved by reducing the network's informational efficiency.
iii) Thirdly, for wisdom of the crowd research, the key network property that governs collective intelligence is not network efficiency, but network centralization, which is operationalized in terms of degree distribution [43,44]. The findings in this field are nuanced, with varying network configurations leading to either greater or lower collective intelligence than independent (non-networked) aggregation; however, the primary insight is that collective intelligence is most reliably increased when there is reduced centralization in communication networks.

In short, collective intelligence on difficult tasks is typically improved by communication networks that have a decentralized, informationally inefficient architecture. The core insight underwriting this admittedly structural conclusion is that this communication infrastructure protects innovative and uncommon ideas from being silenced before they can take hold. A common explanation for this finding is that these networks preserve informational diversity in a population [4,21,45], but this finding also holds in settings in which collective intelligence increases as informational diversity is reduced [46].

As discussed in Box 1, the network synthesis presented here resonates with scholarly findings that span a vast array of disciplinary topics, including social coordination, network diffusion, collective behavior, and the evolution of cooperation, all of which may be reasonably included under the rubric of 'collective intelligence' (Box 1). While all of these topics deserve detailed treatment, there is sufficient nuance in each of these fields that this review focuses exclusively on the network dynamics of collective problem-solving and the wisdom of the crowd.

Collective problem-solving

The study of communication networks and group problem-solving dates back to the mid-20th century [47]. Early network researchers identified the key network measure of average 'path length' (i.e., the shortest number of 'steps' between the nodes in a graph) – as the controlling topological feature governing the efficiency of the spread of information through a population [41,48,49]. Over the next several decades, this research matured into several important insights about the informational pathways that mediate the transmission of social learning within neighborhoods [50], within organizations [51], and even across nations [52]. A key lesson from this work was that 'weak ties' [51], also known as 'crosscutting ties' [50], across communities or organizational units [53], could reduce the 'simple path length' of a network (i.e., 'degrees of separation' between people) and thereby accelerate the spread of information among them [54] (see [42] for recent improvements in measures of simple path length and complex path length).

The effects of informational efficiency on organizational performance rest upon a trade-off between the strategies of 'exploration' and 'exploitation' in team problem-solving [28]. Exploration is the activity of a researcher (or a team) working independently to discover a new way to solve

Box 1. The many fields of networked collective intelligence

An enormous variety of network-related research may be appropriately included under the rubric of 'collective intelligence'. Here, I provide a brief account of the most influential and relevant areas of study that are not included in this review.

Coordination: in coordination dilemmas, individuals select among discrete options with positive externalities, such that players' rewards are conditional on the choices of their peers (e.g., as in social norms and conventions). Rewards are based either on local coordination with one's immediate peers, or global coordination across the population [2,42,99–102,122–128].

Cooperation: in a cooperation dilemma, there is tension between individual and collective rewards. Individual rewards are greatest when individuals exploit their peers, however, collective rewards are greatest when individuals share resources with their peers [54,104,129]. Collective intelligence is determined by the success of cooperation in the face of incentives for individual exploitation [104,130].

Innovation diffusion and mobilization: diffusion studies evaluate collective intelligence in terms of the successful spread of social contagions, such as innovative technologies [131–136] and social movements [51,106]. While the spread of 'simple' contagions (e.g., news or gossip) typically benefits from informationally efficient networks, the spread of 'complex' contagions (e.g., complex knowledge and novel behaviors) typically benefits from clustered, informationally inefficient networks [60].

Information cascades and herding: information cascades are beneficial for the emergence of evolutionarily adaptive collective behaviors in many species, such as ungulate herding and fish schooling [1,22,137,138]. However, studies of herding dynamics among humans emphasize the negative impact of sequential choice [40], in which each person's decision is directly influenced by the people choosing prior to them, which can result in the 'madness of crowds', such as market bubbles and crashes [112,113,139].

Democratic deliberation: is public deliberation (e.g., a town hall forum) a superior form of democratic participation compared with public voting (e.g., a plebiscite or a referendum)? Proponents of deliberative democracy [38,105,140–144] suggest that the open exchange of ideas in a democratic 'marketplace' elicits 'the force of the better argument' [140], while opponents argue that deliberation leads to political polarization and reduced collective intelligence [7,31,38,71,109,140,141,145–148].

Opinion models: cross-disciplinary work in statistical physics, applied mathematics, and computer science has generated a large array of models of opinion dynamics (e.g., the voter model, the Ising model, cultural dissemination models, cultural evolution models, etc.) [54,149,150], which offer promising directions for future empirical research within ecologically valid settings [60].

Because each research area relies on its own formal models of individual and collective behavior, these varied traditions have resisted synthesis. The present article offers a preliminary step toward such a synthesis by identifying common network mechanisms that underwrite the fields of collective problem-solving and the wisdom of the crowd.

a difficult problem [55]. By contrast, exploitation occurs when researchers can see other researchers' (or other teams') solutions and copy them [30,56]. Exploitation can lead to rapid gains in a team's or a firm's competitive capabilities since a good solution found by one researcher can be quickly incorporated into other researchers' approaches to problem-solving [30]. However, a competitive strategy based on exploiting other researchers' good solutions may drain valuable resources away from efforts to pioneer the discovery of more innovative, superior solutions [28]. Exploration is riskier than exploitation because the outcome is less certain. However, while exploitation offers the short-term promise of certain gain, it may also yield long-term disadvantages by reducing the capacity for innovation [57].

Network studies typically formalize the trade-off between exploitation and exploration in terms of 'simple path length' in the social networks among team members [21,30,58]. Organizational networks that are 'informationally efficient' have minimal average simple path length (see [48]), which accelerates information flow among researchers and therefore increases the rate of exploitation among them [56,59]. Conversely, organizational networks that are informationally inefficient are typically highly clustered with a larger average simple path length [60]. These informationally inefficient networks typically reduce exploitation (and increase exploration) by reducing researchers' exposure to new solutions discovered by their peers.

Simple and complex problems

The ability to evaluate the performance of new solutions (or 'innovations') relies on an evolutionary conception of 'fitness' within a problem landscape [25]. Theoretical predictions for the effects of an organization's network structure on its collective intelligence are conditional on the 'complexity' of the problem-solving task at hand [25]. As shown in Figure 1, task complexity is defined in terms of the 'ruggedness' of the solution landscape.

If the solution landscape is 'simple' (i.e., single peaked), then a single strategy can typically be applied along a single dimension until the optimal solution is found (e.g., 'hill climbing' algorithms [61]). Theoretical studies found that efficient networks optimize the rate at which teams solve simple problems [30]. By contrast, complex problem-solving tasks have a 'rugged' (i.e., multipeaked) solution landscape. Theoretical research found that while efficient networks typically lead teams to rapid convergence on a 'local' optimum (i.e., a small hilltop), they rarely reach the 'global' optimum (i.e., the highest peak). By contrast, simulations showed that while inefficient networks take slightly longer to converge, they enabled teams to discover higher quality solutions, often the best possible solution [30] (Figure 2).

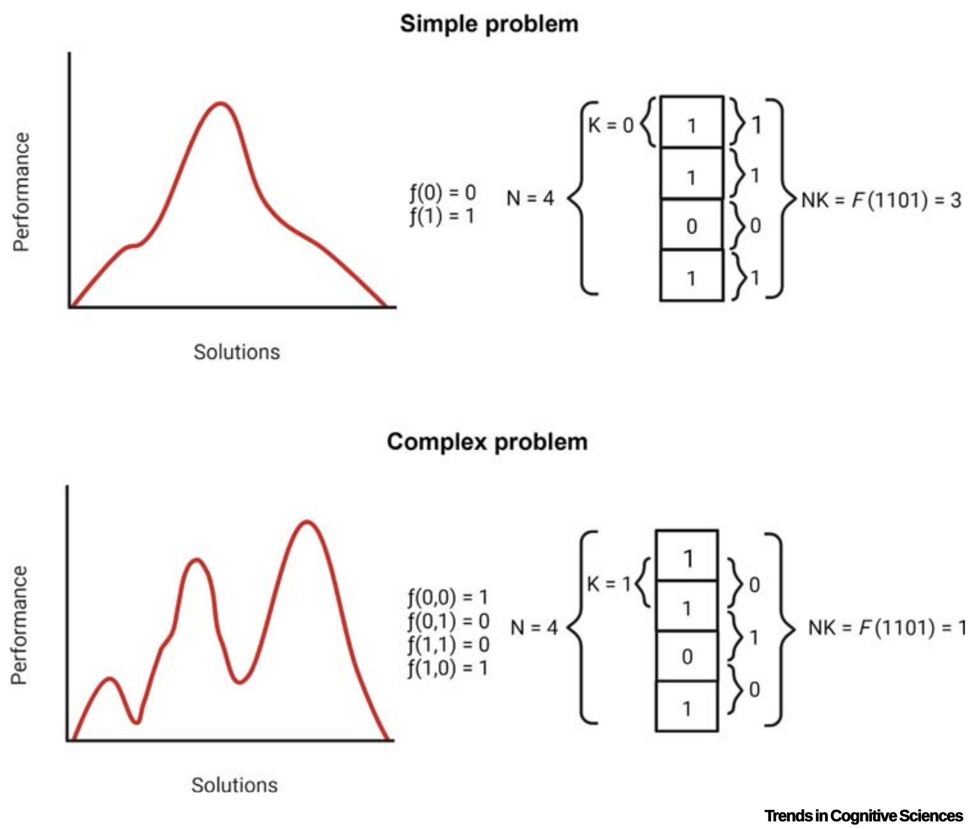

Figure 1. Smooth and rugged solution landscapes. In the *NK* model, each solution is a unique bitstring. *N* is the number of loci in the bitstring and *K* is the interdependence among loci, which determines the 'complexity' of the problem space. In the shown simple problem, $K = 0$, such that the landscape is 'single peaked', in which overall performance of the solutions increases monotonically as each locus is transformed from 0 to 1. Thus, changing the third locus from 0 to 1 increases the overall performance score from 3 to 4. In the shown complex problem, $K = 1$, such that overall performance is a complex interaction across neighboring loci, resulting in a solution landscape that is 'rugged'. In the complex example, changing the third locus from 0 to 1 reduces the overall performance score from 1 to 0. This figure was created using BioRender (https://biorender.com/).

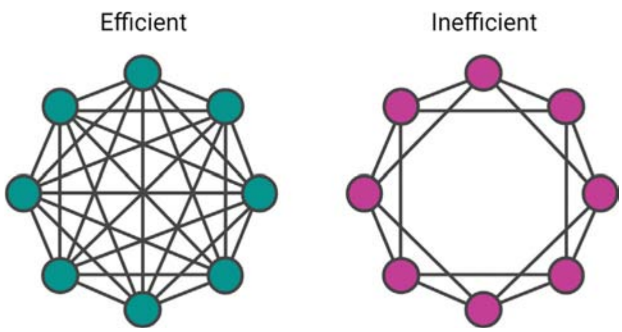

Figure 2. Informationally efficient and inefficient networks. Informational efficiency is measured by simple path length. This informationally efficient network is a 'fully connected network' (i.e., a complete graph), such that the average simple path length PLs = 1 (i.e., everyone in the network is directly connected to everyone else). In the informationally inefficient network, several steps are required to traverse the network. This figure was created using BioRender (https://biorender.com/).

Four scope conditions circumscribe these theoretical results on networks and collective problem-solving: (i) problem-solving team members continually (i.e., round after round) evaluate whether to explore the problem space (i.e., look for a new solution using their existing approach) or whether to exploit one of their contacts' solutions (i.e., copy a solution already found by one of their contacts); (ii) team members can only copy the strategies of the peers to whom they are directly connected in the network; (iii) team members can observe the performance quality of the solutions used by their network contacts; and (iv) team members cannot observe any global information about the overall popularity or quality of different solutions beyond observing their immediate network contacts.

Experimental findings on networks and collective problem-solving

An early experimental test of these computational predictions did not find clear support for or against them. This study used Amazon Mechanical Turk (or 'MTurk') participants to study how a large variety of constructed networks among participants would affect players' performance in an experimentally designed video game [58]. Researchers conducted hundreds of experimental trials with networks of varying path lengths but found inconsistent results. Subsequent studies uncovered two explanations for these findings: first, MTurk subjects did not offer an ecologically valid comparison to professional researchers working under competitive time pressures, and second, the design of the experimental video game allowed subjects to observe global information about the popularity of the solutions that other players were using, which violated the scope conditions of the model by engaging a learning model based on conformity rather than the theoretically proposed learning model based on performance [62,63].

A subsequent test of the theory of networked collective problem-solving recruited graduate and professional data scientists to participate in an online data science competition [21]. Researchers randomly assigned competitors to teams with either maximally efficient team communication networks (i.e., a complete graph) or highly inefficient team communication networks (i.e., a one-dimensional lattice) (Figure 2). All competitors were asked to solve complex data science problems and were provided a set of statistical tools for exploring the solution space. After eight competitions, teams with inefficient communication networks exhibited significantly greater collective intelligence than teams with efficient networks. While teams with efficient networks never found the optimal solution, this solution was found by half of the teams with inefficient networks. These teams also outperformed several common machine learning strategies ([64] *cf.* [6]).

The expanding literature on networks and collective problem-solving (Box 2) has explored a wide variety of theoretical and empirical extensions, including: (i) studying agents dynamically rewiring their peer networks during the problem-solving process [65]; (ii) examining the differences

Box 2. Team composition in collective problem-solving

In contrast to the considerations of network structure discussed in the main text, there are several non-network approaches to improving teams' collective problem-solving abilities, which focus instead on ideal team composition [4,45,151–154]. These studies typically identify optimal strategies for selecting team members based on: (i) their cognitive (i.e., problem-representation) diversity as compared to their technical ability (e.g., the 'diversity trumps ability' theorem) [4,153]; (ii) their demographic characteristics as compared with their social network characteristics [151,152]; (iii) the number of experienced incumbents versus innovative newcomers on a team [155]; or (iv) essentialist characteristics of team members, such as gender or sociability, that are hypothesized to create task- and context-independent group intelligence, sometimes referred to as 'c-factor' [151]. While these approaches to team construction are promising, they have yet to be integrated with structural studies of team problem-solving dynamics in varying network topologies [153]. This opens several exciting directions for future work investigating the interaction of team assembly strategies and team network structures. For instance, promising work at this frontier has sought to identify the effects of clustering team members within organizational networks based on individuals' attributes, such as diversity of ability versus diversity of knowledge [156].

between individuals engaged in lone exploration, versus small, informationally efficient groups engaged in collaborative discovery [66]; (iii) exploring variations in the frequency of interactions in a population (e.g., continuous, intermittent, or no interaction) while holding networks constant to determine whether manipulating meeting frequency impacted teams' problem-solving abilities [56]; (iv) studying how the changing structure of collaboration networks affects team productivity and artistic creativity on Broadway [29]; (v) examining the novel capacity for distributed cognition through complex interactions between humans and computers within modern networking technologies [5,6,67,68]; and (vi) identifying differences between 'information transfer' versus 'knowledge transfer' within organizational networks, showing that informationally efficient 'weak tie' networks accelerate rapid information transfer across teams, while informationally inefficient 'wide bridge' networks improve knowledge transfer across teams [60,69,70].

Conclusions for collective problem-solving

Each of the aforementioned explorations reports findings that are consistent with the conclusion that increased clustering and moderate information flow among teams improves collective intelligence in complex problem-solving tasks. The mechanism governing these findings is that informational inefficiency maintains informational diversity among team members [4], thereby enabling the exploration of unlikely or uncommon solutions. This structural capacity to 'protect' innovative ideas is essential because, although novel approaches to problem-solving can yield high rewards given sufficient time, they typically exhibit lower performance early on in the discovery process, when they can be easily outcompeted by more familiar solutions [30,64]. While there are many nuances to this literature, the overall findings suggest that organizations, teams, and industries faced with complex problem-solving tasks can optimize group performance by strategically designing institutional structures (such as network bridge-width across teams, meeting frequency among colleagues, meeting size among workgroups, and the duration and structure of conference opportunities within industries) to enable a modest exchange of information on an intermittent basis, thereby preserving informational diversity and expanding intellectual creativity within research communities.

The wisdom of the crowd

Unlike collective problem-solving research, which originates from questions about how to optimize communication networks, research on the wisdom of the crowd began from a suspicion about the intelligence of the 'common person'. The father of eugenics, Sir Francis Galton, designed an experiment to demonstrate the folly of democracy, in which the attendees of a county fair were asked to guess the weight of an 'undressed ox'. To Galton's surprise, the average estimate of the 800 fair-goers was stunningly accurate, within 1% of the true answer [37]. This phenomenon, later dubbed 'the wisdom of the crowd' [7], is a result of the fact that, in large groups, diverse individuals' overestimates and underestimates cancel each other out, resulting in a group estimate that can be closer to the truth than any of the individuals' estimates.

Continuous and discrete tasks

Before discussing the network dynamics of the wisdom of the crowd, I will take a few moments to position this work amid the non-network refinements to Galton's original approach, which have produced many fascinating subgenres of research on independent aggregation techniques (i.e., without social influence) for evaluating collective intelligence. This literature benefits from a useful distinction between continuous variable estimates versus discrete choice problems. For estimation tasks, such as Galton's challenge of estimating an ox's weight, the wisdom of the crowd is typically evaluated in terms of a comparison between the overall accuracy of the group's central tendency (e.g., the accuracy of the group mean) versus the individual accuracy of each of the group members. There are two slightly different characterizations of this comparison. The first emphasizes that the accuracy of the group is typically greater than the accuracy of any individual group member (or of any individual 'expert') [7,37,46]. The second characterization notes that the error of the group's central tendency (e.g., the error of the group mean) is lower than the average error among the individual members of the group (i.e., the 'crowds beat averages' law [4]; Box 3). Both characterizations of the wisdom of the crowd have in common the idea that the error of the group's central tendency will consistently be lower than (and will never be greater than) the average individual error of the group members.

Box 3. Aggregation for estimation tasks

For independent (non-network) aggregation approaches to estimation tasks, the wisdom of the crowd can be easily represented in terms of the 'crowds beats averages' law [4], which states that the average individual error of group members is never less than (and is typically greater than) the collective error of the group. Put formally, if N = group size, c = collective estimate, θ = true value, and B_i = individual estimates, then:

$$(c-\theta)^2+\frac{1}{N}\sum_{i=1}^{N}(B_i-c)^2=\frac{1}{N}\sum_{i=1}^{N}(B_i-\theta)^2 \qquad \text{[I]}$$

In other words, group error + group diversity = average individual error. A corollary referred to as the 'diversity prediction theorem', shows that group diversity and average individual ability contribute equally to collective intelligence. More formally:

$$(c-\theta)^2=\frac{1}{N}\sum_{i=1}^{N}(B_i-\theta)^2-\frac{1}{N}\sum_{i=1}^{N}(B_i-c)^2 \qquad \text{[II]}$$

In other words, group error = average individual error – group diversity. Assuming the independence of group diversity and individual error, reducing group diversity increases group error while increasing group diversity reduces group error.

Empirical findings on network dynamics [77,81,86] show that social influences that reduce group diversity typically interact with average individual error. In decentralized networks, social influences that reduce group diversity by magnitude α have been found to simultaneously reduce average individual error by magnitude $\delta > \alpha$. In such cases, decreasing group diversity can lead to significant reductions in group error [46]. More generally, the key network factor controlling the effects of changing group diversity on group error is network centralization. In decentralized networks, decreasing group diversity typically leads to rapidly decreasing group error [86]. As discussed in the main text, in centralized networks these effects are governed by whether or not the central actor is 'in the same direction' as the correct answer [46].

Related studies of collective estimation tasks, such as the Delphi method and prediction markets, allow social influence among participants, but do not consider network structure. These approaches use carefully specified rules, such as message moderation procedures and market pricing rules, to manage the effects of social influence on group decisions. In the Delphi method, the content and frequency of information exchange among group participants are carefully monitored (e.g., via written notes) by a group moderator. In prediction markets, people respond to one another indirectly through a market mechanism that prices forecasts based on people's willingness to invest in, or divest from, those predictions. Related approaches use sequential interactions among small groups of experts to explore how iterative estimations among knowledgeable people [157–159], sometimes with voluntary opt-in and opt-out strategies [160–163], or even among lay persons engaged in direct conversation [164], can improve upon the wisdom of the crowd that emerges from groups of independent actors.

In contrast to Galton's interest in estimation tasks, another tradition in the wisdom of the crowd focuses on discrete choice problems, such as determining the guilt or innocence of a prisoner, or selecting between two candidates for elected office. Several important results in this tradition, such as Condorcet's jury theorem (CJT) [31] and versions of the miracle of aggregation (TMA) [4,71], have been used to provide a wisdom-of-the-crowd justification for modern democratic voting theory (i.e., an 'epistemic defense' of majority rule) [71]. Importantly, these arguments typically assume statistical independence among voters and require strong assumptions about the accuracy of each individual voter, sometimes referred to as the 'hardness' of the voting problem (Box 4).

The best approach to increasing the wisdom of the crowd can vary considerably depending on whether a group is facing a Galton-type estimation problem or a Condorcet-type discrete choice problem (Boxes 3 and 4). For instance, collective intelligence on an estimation task typically improves with increasing population size [4,7,46]. While this positive effect of group size can also hold for 'easy' discrete choice problems (for which people have >50% chance of choosing correctly), the opposite is typically true for 'hard' problems (for which people have <50% chance of choosing correctly), in which case reducing the inclusiveness of the aggregation method (e.g., using a plurality rule or a quorum rule) can outperform a more inclusive majority rule approach [16,31,72–74] (Box 4).

The difference between Condorcet-type discrete choice problems versus Galton-type estimation tasks can sometimes be reduced to a matter of operationalization (i.e., 'will it rain' versus 'what is the likelihood of rain'), yet this distinction in problem structure can nevertheless have important implications for a group's collective intelligence. For instance, collective intelligence on discrete choice problems relies on the fact that some fraction of the population will make the correct

Box 4. Aggregation for discrete choice problems

There are two classic 'majority rule' results for discrete choice problems, assuming independence among voters [31]: (i) Condorcet's jury theorem (CJT); and (ii) the miracle of aggregation (TMA).

(i) CJT states that if individuals choosing between A and B have a predisposition to make the correct choice, then increasing population size (i.e., a more inclusive vote) increases the likelihood of a correct collective choice. For instance, for population size $N = 10$ and an average likelihood of choosing correctly, $p_i = 0.51$, a minimal majority of six voters has a probability $P_{\text{majority}} = 0.52$ of choosing correctly, and a supermajority of eight has a $P_{\text{supermajority}} = 0.56$ of choosing correctly. However, for $N = 1000$ and $p_i = 0.51$, a minimal majority of 501 voters has a $P_{\text{majority}} = 0.73$ of choosing correctly.

(ii) The simplest version of TMA is to suppose only 10% of voters are sufficiently knowledgeable to make the correct choice between A and B, while everyone else chooses randomly. In a small population ($N = 10$), stochastic variation among uniformed individuals may lead the ignorant 90% to outweigh the informed 10%. However, if $N = 100\,000$, random choices among the 90% will yield 45% of the vote for each option, empowering the 10% of knowledgeable individuals to 'influence' the population's majority vote, selecting the correct choice by a margin of 55% to 45%.

An important difficulty arises for both the CJT and TMA accounts of majority rule when ignorant voters are not random in their choices (i.e., flipping an unweighted coin), but rather are systematically biased away from the correct choice [165]. In this case, CJT and TMA results switch from favoring larger groups to favoring smaller groups, since a smaller group is less likely to reflect the majority bias [31,165].

Recent work exploring sampling strategies that select minority votes rather than using the entire population has found useful results from 'quorum' methods, which use population information to preselect an effective minority for improving the population's likelihood of making a good choice [10,16,17,73,115,166]. In medicine, quorum rules can outperform the choices of both the most accurate individual clinician and the majority of clinicians [16].

Notably, CJT and TMA are based on the assumption that the goal of voting is to select the correct answer [1,31,71]. However, this assumption may be inappropriate if the goal of democratic voting is not to select the best-performing choice, but to effectively, fairly, and peacefully balance competing interests, even if the resulting choice reduces collective performance (*cf.* constitutional [144] versus biological [24,167] accounts of social choice).

categorical choice ([31], see [4]) (Box 4). By contrast, groups can achieve high levels of collective accuracy on estimation tasks even when no individual provides the correct answer (e.g., if 50% overestimate and 50% underestimate, the central tendency can be perfectly accurate although everyone's guess is wrong) [37,75].

Following Galton, a growing literature on estimation tasks revolves around refinements for measuring the central tendency of populations [76]. Galton favored the group median as a measure of central tendency, while Rousseau (who offered a similar, if less scientific proposal a century earlier) nominated the group mean as the best measure of central tendency (see also [4,7]). Recently, scholars have explored a much wider variety of methods (e.g., arithmetic mean, geometric mean, average of the mean and median, maximum-likelihood estimation, etc.), finding that the success of each measure can vary depending on the level of bias within the underlying distribution of estimates in the population [75].

As I return next to the main discussion of networks and the wisdom of crowds, I use the group mean as the primary method for identifying a group's central tendency because it has proven to be a reliable metric in dynamical network contexts [46]. For instance, for estimation problems with initially skewed distributions (e.g., whole number responses to count questions) [22,46], social influence can improve the group median (as individuals converge toward the mean) while failing to improve the group mean [46,77]. By contrast, for estimation tasks with symmetrically distributed responses (e.g., probability estimates between 0 and 1 with a normal distribution), improvements in the group median are generally coextensive with improvements in the group mean. Thus, within the scope of this discussion, the measure of the group mean provides a conservative approach for identifying changes in group performance across a range of estimation tasks. There are many productive ways in which a population's central tendency might be measured and the exploration of alternative measures of central tendency within networked contexts offers an important area of ongoing research [75].

Network centralization and estimation tasks

Returning now to the main theme of this section on the effects of network structure on the wisdom of crowds (see Box 5), I begin with the observation that the network dynamics of collective intelligence vary considerably between Galton-type estimation tasks and Condorcet-type discrete choice problems. Recent findings show that network structures that consistently increase group intelligence on estimation tasks can nevertheless reduce group intelligence for discrete choice problems [78,79]. For clarity, the immediate focus of this discussion will be on the network dynamics of estimation problems, followed by a concluding discussion on the challenges, and outstanding questions for the network dynamics of discrete choice problems (see in particular the Concluding remarks and further directions).

Building on Galton's original results for estimation tasks, a mushrooming literature explores the potential for the wisdom of the crowd to be used to improve everything from the science of climate change [39] to procedures for democratic elections [7]. An initial challenge to the theory of crowd wisdom came from scholars who appreciated that peer influences could undermine collective intelligence by correlating the error terms in people's estimates, in essence, turning the wisdom of the crowd into 'groupthink' [31,74,75,80]. Experimental studies of the negative effects of peer influence on collective intelligence suggested that social influence decreases group intelligence by driving people toward converging opinions (i.e., reduced informational diversity) while increasing individual confidence in the group answer [39]. However, a reanalysis of these experimental data showed that while informational diversity decreased as a result of social influence, the group estimate did not change significantly and the accuracy of individuals' estimates improved significantly [77]. These findings were supported by a wide range of theoretical and empirical studies

Box 5. Network-based aggregation in the wisdom of crowds

A significant difference between independent aggregation versus network approaches to the wisdom of crowds is that the former do not enable any form of individual learning. For instance, for discrete choice tasks, while the miracle of aggregation (TMA) and Condorcet's jury theorem (CJT) offer methods for arriving at the correct collective choice, the independence assumption entails that individual voters do not learn anything from one another during the voting process. Similarly, a quorum rule for eliciting collective decisions from groups of physicians may select the best choice, but it does not enable participating clinicians to learn from their peers in the process . This is also true for estimation tasks. Calculating the central tendency of a group does not provide a way for individual group members to improve their guesses. Relatedly, while the diversity prediction theorem illustrates how group error can decrease as a result of increasing group diversity, it does not provide any mechanisms through which average individual error may be reduced and thus does not shed light on how individual population members may improve the quality of their responses as a result of increasing group diversity.

As discussed in the main text, for Galton-type estimation tasks, learning dynamics in social networks can reduce individual errors among group members while also reducing overall group error [36,86]. As discussed in the Concluding remarks and future directions, network structures that consistently improve collective intelligence for estimation tasks may nevertheless compromise collective intelligence for discrete choice problems [79]. For discrete choice problems in which social influence may lead to polarization, an effective solution may be to reframe the problem in a continuous variable estimation format, which may effectively engage the network dynamics of social learning to improve individual accuracy and collective intelligence while reducing group polarization [8,78,79].

([81,82] see also [83–85]), which suggested that decreasing informational diversity as a result of social influence does not necessarily undermine collective intelligence (Box 3).

Building on these ideas, a series of experiments investigated the effects of network structure on the dynamics of peer influence in the wisdom of crowds [86]. Many key network properties were investigated (e.g., simple path length, density, clustering, etc.); however, network centralization was found to be the governing network parameter affecting the wisdom of the crowd [43,86]. As shown in Figure 3A, network centralization is controlled by the degree distribution in the social network, which can impact both individual and collective intelligence [8,46]. In decentralized networks, social influence was found to significantly reduce informational diversity while improving individual performance. Consistent with earlier findings, the mechanism governing this result was individual convergence toward the group mean. Unexpectedly, however, the findings also showed that decentralized networks reliably improved the wisdom of the crowd itself (i.e., the group mean became more accurate), leading decentralized networks to exhibit significantly lower group error than found among independent control groups without social influence; these control groups were equivalent to Galton's original wisdom of the crowd. Researchers identified network size as an important scope condition for this finding, indicating that networks above a minimal size ($n > 30$) were needed to reliably detect the effects of peer influence on improving the accuracy of the group mean.

Changes in network structure can significantly affect these results. While decentralized networks consistently improved collective intelligence, centralized networks were found to either significantly improve or significantly reduce the wisdom of the crowd, depending on the relationship between the central person, the correct answer, and the population mean (Figure 3B). Results showed that collective intelligence (i.e., average error of the group mean) improved if the central person pulled the group mean toward the correct answer, but decreased if the central person pulled the group mean away from the correct answer ([86,87] see also [43] on 'persuasion bias').

Mechanisms and extensions

The network mechanisms governing improvements in the wisdom of the crowd differ according to network structure. In centralized networks, the mechanism for any positive effects of social influence on the wisdom of crowd is the disproportionate influence of a central node 'in the same direction' as the correct answer. By contrast, in decentralized networks the mechanism for

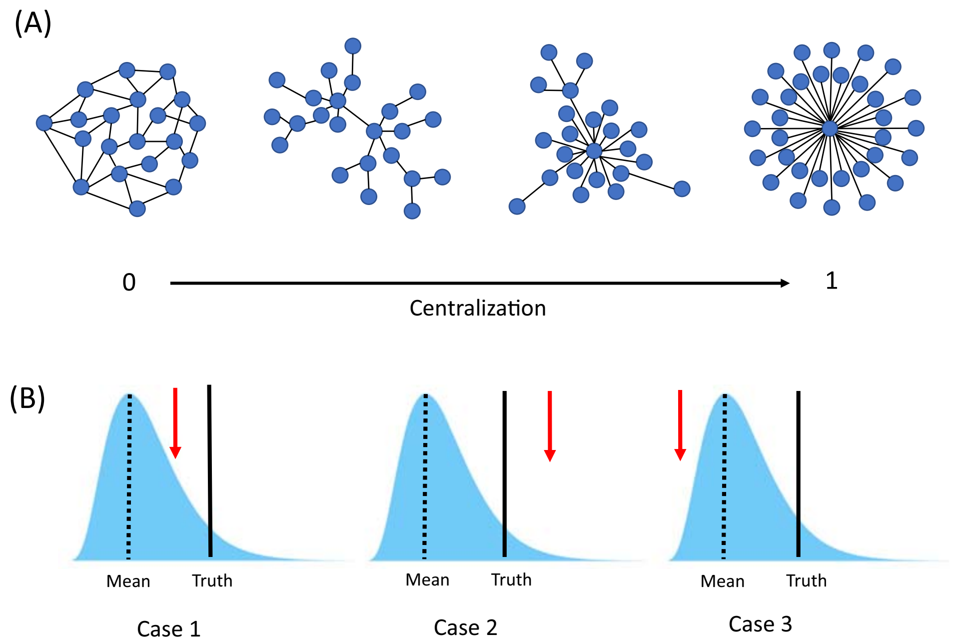

Figure 3. Networks and the wisdom of crowds. (A) A continuum of network centralization. In networks with highly skewed degree distributions, a small number of people (or even a single person) have a disproportionate number of contacts and therefore control the flow of information throughout the network. By contrast, as the degree distribution becomes more uniform (i.e., decentralized), everyone in the population has greater equality of connectivity. In the shown graph, the decentralized network is less informationally efficient than the centralized network because several additional steps are required for information to traverse the decentralized graph. (B) Three different possible cases for the relationship between the central node, the group mean, and the correct answer. Because the vast majority of people have opinions that are in the same direction as the group bias, the vast majority of people (if located in the center of the network) are likely to undermine the wisdom of the crowd rather than improve it.

improvements in crowd wisdom is a population-wide correlation between individual accuracy and the magnitude of revision, referred to as the 'revision coefficient' [86]. When there is a positive revision coefficient in the network, more accurate individuals revise their responses less (i.e., a high 'self-weight'), while less accurate individuals revise their response more (i.e., a low 'self-weight') [75,86,88]. In centralized networks, a positive revision coefficient does not affect the collective outcome. However, in decentralized networks, the effect of a positive revision coefficient is that accurate individuals become 'centers of gravity' within the network, pulling the rest of the population toward them. The result is a collective convergence toward the most accurate responses, which reduces the diversity of information in the population while simultaneously increasing both individual performance (i.e., reducing individual error) and the wisdom of the crowd (i.e., reducing group error).

Recent extensions of these findings on 'networked collective intelligence' have explored the implications of improved collective judgment within networks for reducing well-known group biases. For instance, studies have examined the effects of decentralized communication networks for reducing political polarization and partisan bias on both climate change [8] and immigration [87] using, respectively, (politically heterogeneous) cross-party networks and (politically homogeneous, or 'homophilous') echo chamber networks. In both cases, decentralized networks reduced polarization among initially polarized groups while increasing collective accuracy among both parties [78], with a notable caveat in settings in which polarizing images triggered partisan priming effects [8,89]. Subsequent studies extended the findings on bias reduction to healthcare and medicine, showing that peer influence in decentralized networks reduced smokers' biases regarding smoking risks ([90] see also [91,92]) and reduced implicit race and gender bias in clinicians' treatment recommendations for patients of varying race and gender ([36] *cf.* [93]).

Further extensions of this network paradigm relax traditional assumptions about informational content and network structure by allowing people to see the accuracy of one another's guesses while also allowing them to dynamically select their network ties [63]. In the presence of informational feedback about peer accuracy, group intelligence increases as networks become more centralized around accurate individuals. Other studies have noted that information about individual accuracy is often not available in real-time decision-making [94], suggesting that the potential for centralized networks to reliably improve the wisdom of crowds relies upon further research parameterizing the relative empirical timescales of: (i) network evolution, (ii) access to information about peer accuracy, and (iii) the process of individual decision-making, within various ecological settings. Finally, the newest frontier in this network tradition is the exploration of the evolutionary space of 'adaptive collective intelligence', which engages important new questions about the robustness of network strategies in situations in which the kinds of collective intelligence problems that groups may face are either unknown or changing (Box 6).

A preliminary synthesis of networked collective intelligence

The summary in Table 1 (Key table) offers three conclusions, which together provide an initial synthesis of the network dynamics governing both collective problem-solving and the wisdom of the crowd.

First, a collective problem-solving task can be transformed into a wisdom of the crowd estimation task by selecting a 'simple' (i.e., single peaked) problem and concealing information about the quality of individuals' solutions. In that case, the network dynamics of the wisdom of the crowd will govern the outcome and decentralized networks will be the most reliable way to minimize group bias and improve collective intelligence.

Box 6. New frontiers in collective adaptation

A new frontier in collective intelligence research has begun to explore the unintended consequences of collective adaptation [168,169], in which a network solution that increases group performance for one kind of task (e.g., increased network centralization to improve group performance on a simple collective problem-solving task), inadvertently leads groups toward a social configuration that reduces collective intelligence for other kinds of problems (e.g., a wisdom of the crowd task). Additional nuance in this burgeoning research area comes from the fact that the evolutionary adaption of social networks in response to a particular challenge of collective intelligence (e.g., increasing connectivity to increase the speed of a group's response to external threats) may unwittingly alter a population's trajectory through a complex, evolving problem landscape. For instance, when increased connectivity in the network inadvertently increases the frequency of emergent cooperation dilemmas in the population) [57,170,171]. Thus, one consequence of network adaptations that resolve particular collective intelligence problems may be that they alter the kinds of future problems that a group will face.

This evolutionary space of 'adaptive collective intelligence' opens a fascinating new territory of research. For instance, one question of interest is whether there are structures that may be more robust across multiple kinds of challenges. Are there network structures that maximize performance quality for one class of challenges, such as complex collective problem-solving, while also minimizing performance impairment for other classes of problems, such as wisdom of the crowd tasks [172]? This question may arise in several different ways. For example, in situations of 'problem ambiguity', it may be unclear at the outset what kind of collective intelligence problem a group is facing. Alternatively, when collectives face multiple problems, a group may need to find a social configuration that can sustain high performance while navigating several different kinds of collective intelligence challenges simultaneously [169]. As discussed in the main text, initial findings indicate that decentralized, informationally inefficient networks may offer a surprisingly robust architecture for navigating a variety of sequential and simultaneous collective intelligence tasks under conditions of problem uncertainty.

Key table

Table 1. Implications of problem type and informational content on key network properties for collective intelligence[a]

Field of collective intelligence	Collective problem-solving	Collective problem-solving	Wisdom of the crowd	Wisdom of the crowd
Problem type	Simple	Complex	Estimation tasks	Discrete choice tasks
Informational content	Peer solution + solution quality	Peer solution + solution quality	Average of peer estimates	Distribution of peer choices
Key network property	Informational efficiency (simple path length)	Informational efficiency (simple path length)	Centralization (degree distribution)	TBD
Desirable network	Efficient: complete graph or highly centralized graph	Inefficient: regular lattice network	Decentralized: complete graph or regular lattice network	TBD Expected to depend on 'hardness' of voting problem

[a]This table briefly summarizes the key factors that distinguish research on collective problem-solving from research on the wisdom of the crowd, namely: (i) problem type, and (ii) informational content and their consequences for network structure. In the wisdom of the crowd estimation tasks, findings show that providing the average of peer estimates exhibits similar collective dynamics as providing individual peer estimates [8]. In the wisdom of the crowd discrete choice tasks, while current results suggest that peer influence typically results in the majority choice, additional research is needed to explore variations that may arise under different 'hardness' conditions, as well as variations in the voting rules.

Second, if a collective problem-solving task is simple and the participants do have information about the performance of one another's solutions, then reducing the simple path length in the network will optimize collective intelligence. There are many ways to reduce simple path length. Because a highly centralized network has a very low simple path length, it can reliably improve collective intelligence for this kind of task. This network would not reliably increase collective intelligence for a wisdom of the crowd estimation task unless the central individual was 'in the same direction' as the correct answer. However, there is a general network solution that reliably improves collective intelligence for both simple collective problem-solving tasks and wisdom of the crowd estimation tasks: a fully connected network. For a simple collective problem-solving task, this network offers slightly better performance than a highly centralized network because it has a slightly lower simple path length than a centralized network. Moreover, a fully connected network is also decentralized and thus would benefit the wisdom of the crowd.

A minor caveat that invites future research in this area concerns the extent to which increasing population size can create *de facto* centralization in networks that are fully connected. A very large, fully connected network discussion (e.g., on social media), may result in many voices competing to be heard simultaneously, leading the loudest or most charismatic voice to dominate the conversation (i.e., *de facto* centralization). This concern suggests that, at large scales, informationally inefficient decentralized networks (e.g., moderate-degree lattices) may be more reliable than informationally efficient decentralized networks (i.e., fully connected graphs) for improving the wisdom of the crowd. One suggestive implication of the tendency toward centralization in large groups may be that it helps to explain why large scientific teams are more prone to producing more conservative 'normal science', further developing and expanding existing models, while small teams with more egalitarian structures are more likely to produce innovative work that disrupts existing scientific models [95,96].

Third, for a complex problem-solving task, in which participants have information about the performance of one another's solutions, a fully connected network performs suboptimally because it is too informationally efficient. However, a decentralized network with limited informational efficiency, such as a regular lattice network with modest average connectivity, can reliably improve a group's collective intelligence. Because this network is decentralized, it can also reliably improve the wisdom of the crowd for estimation tasks.

Concluding remarks and future directions

The ability of decentralized, informationally inefficient communication networks to improve collective intelligence is governed by different mechanisms for different tasks. For complex collective problem-solving, informational inefficiency improves collective intelligence because it preserves informational diversity. However, for the wisdom of the crowd estimation tasks, a decentralized structure improves collective intelligence because it leads to convergence (i.e., reduced informational diversity) on the most accurate responses. These different mechanisms notwithstanding, the common network principle that facilitates collective intelligence for both complex collective problem-solving and the wisdom of the crowd is the ability of the network structure to protect unpopular ideas (i.e., ideas that challenge group biases and beliefs) from being overpowered early on in a decision process by too many countervailing influences [97,98]. In particular, networks that enable innovative ideas to emerge and take hold without being silenced by a highly influential 'opinion leader' are reliable structures for nurturing the growth of opinions that can challenge the status quo and improve collective reasoning [64].

Generalization of network insights

This network principle effectively generalizes to another area of collective intelligence research, often referred to as 'social coordination' (Box 1). Research in this area is grounded in the game-theoretic challenge of enabling a population to coordinate on an optimal behavior [99]. Unlike collective problem-solving and the wisdom of the crowd, in which each person is typically rewarded based on the quality of their individual response (for alternative reward strategies see [32–34]), social coordination research assumes the existence of positive externalities, such that an individual's payoff increases with the number of others who adopt the same option, for instance, people coordinating on the use of a social media technology or a linguistic convention [100,101]. Consistent with findings on complex collective problem-solving and the wisdom of the crowd, increasing informational efficiency and network centralization accelerate the rate of convergence, but can lead to suboptimal coordination equilibria [102]. By contrast, decentralized, informationally inefficient (e.g., clustered) social networks can accelerate social coordination on an optimal choice [103]. Consistent with the earlier discussion, this positive effect of network structure on collective intelligence for coordination tasks comes from the fact that a decentralized, informally inefficient architecture protects innovative and unfamiliar options, such as a superior challenger technology, from being overwhelmed early on by countervailing influences from more familiar options, like an inferior but widely used incumbent technology [60].

Similar insights may also extend to the classic problem of cooperation (Box 1). Clustered, informationally efficient networks are effective for increasing both the rapid growth and the evolutionary stability of cooperation, particularly when cooperators must continually face competition from defectors [60,104–106]. As earlier, the network mechanism for the success of cooperation is the ability of decentralized, informationally inefficient network structures to protect clusters of cooperators from being exploited by excessive contact with selfish actors throughout the population [60].

These scientific conclusions highlight a fairly immediate practical implication. The endogenous evolutionary dynamics of social media networks tend to evolve toward highly centralized, informationally efficient topologies. These network structures are optimized for widespread coordination on popular solutions, as well as the rapid spread of familiar ideas that conform to group biases, even when those ideas are incorrect [60]. These networks can thus undermine the wisdom of the crowd and derail efforts to solve complex problems by amplifying inaccurate but

Outstanding questions

What is the psychological model that underwrites the striking consistency with which decentralized social networks improve the wisdom of the crowd for estimation tasks? The consistent finding that people who are more accurate tend to revise their responses less (and therefore tend to be more influential) is not well understood. Several studies have shown that confidence is not always highly correlated with accuracy, indicating that individual confidence levels may not explain the pattern of social influence observed in decentralized networks.

What is the empirically relevant timescale on which people learn about the quality of peer solutions? Assuming people can reassign their network ties based on the quality of peer solutions, how can we account for inevitable lags in information and how might these lags affect the quality of the collective outcomes that result from network evolution concomitant with real-time decision-making?

Are continuous variable estimates (such as risk estimates) a reliable proxy for people's categorical choices (such as medical treatment decisions and voting decisions)? The answer to this question informs a great deal of research, since the dynamics of network aggregation for continuous variables does not easily generalize to categorical variables.

How do characteristics such as race, gender, income, health status, political affiliation, etc. affect collective intelligence differently in different contexts? For instance, it is not well understood why political graphics that stimulate 'partisan priming' undermine collective intelligence in cross-party interactions, but not in homogeneous political groups. Similarly, revealing political identifiers in cross-party interactions can increase partisan bias and reduce collective intelligence; however, revealing health-status identifiers in cross-group interactions with smokers and nonsmokers can reduce health bias and increase collective intelligence. New work is needed to explore the interactions between network composition, peer identity, and network structure in different substantive domains.

popular beliefs [107–109]. An important direction for future work is to explore how to design social media channels that improve the intelligence of networked populations by both: (i) reducing the general spread of misinformation [78,110] and (ii) limiting the susceptibility of targeted communities to weaponized disinformation campaigns (e.g., regarding vaccine safety) [111].

New frontiers for the wisdom of the crowd

The dawn of network science has reinvigorated the centuries-old problem of collective intelligence, revealing hundreds of novel theoretical and applied problems for researchers working at this interdisciplinary frontier. To start, most real-world group decisions involve discrete choices that are either binary (i.e., vote for candidate A or candidate B) or categorical (i.e., send a patient home, prescribe medication, or refer to surgery). However, techniques for network aggregation that reliably improve the wisdom of the crowd have only been shown to be robust for estimation tasks ([79]; see also [31,40,112,113]).

Strategies to address this 'discrete variable problem' for peer influence in the wisdom of the crowd have suggested solutions in which actors participate in the network exchange of continuous estimates before providing a discrete choice [64]. The effectiveness of this solution is nuanced, however. Recent theoretical and experimental explorations have identified specific cases in which improving estimate accuracy can reduce the number of individuals making the correct discrete choice [79]. Nevertheless, recent experimental work has also shown that using decentralized peer influence networks to improve estimate accuracy can reliably increase the quality of physicians' categorical decisions regarding patient treatment [36], as well as improve social media users' ability to correctly classify true and false news stories [36]. One important direction for future work is to explore the extent to which discrete decisions can, in practice, be reduced to an underlying estimation task [114–116] (see Outstanding questions). This open area of research is particularly relevant for 'crowdsourcing' approaches to detecting specific diagnoses in medicine and other fields [117,118].

Another fascinating new research area concerns work trying to understand the psychological model that underwrites the striking consistency with which decentralized social networks improve the wisdom of the crowd for estimation tasks [119]. One suggestion is that more accurate individuals are also more confident [120], however, the relationship between accuracy and confidence is not without complications [63,121] and has even been found to be correlated in the opposite direction when group biases come into play [88,90]. While the relationship between individual confidence and accuracy has been found to vary across contexts, the dynamics of collective intelligence in decentralized networks are strikingly consistent across domains, suggesting that the effects of individual accuracy on the network dynamics of collective intelligence, rooted in a positive revision coefficient, are not reducible simply to the influence of the most confident individuals. The interplay between interpersonal network dynamics and individual psychology offers an exciting frontier for the exploration of how individual learning mechanisms interact with features of network structure (in particular when confronting longstanding group biases) to improve individual and collective intelligence.

Acknowledgments

I wish to thank Devon Brackbill, Joshua Becker, Doug Guilbeault, Arnout van de Rijt, and several other readers for valuable suggestions and comments, as well as four anonymous peer reviewers and the TICS editor for helpful advice on this article. I would also like to thank the Robert Wood Johnson Foundation, Perbak Capital Partners, the Templeton Foundation, and Penn's Medical Communication Research Institute for research funding related to this article.

Declaration of interests

No interests are declared.

References

1. Aristotle (1999) *Politics*, Batoche Books, Kitchener
2. Hobbes, T. (2010) *Leviathan*, Yale University Press, Shapiro
3. Le Bon, G. (2001) *The Crowd: A Study of the Popular Mind*, Batoche Books, Kitchener
4. Page, S.E. (2007) *The Difference: How the Power of Diversity Creates Better Groups, Firms, Schools, and Societies*, Princeton University Press
5. Malone, T. (2018) *Superminds*, Brown Little and Company
6. Patel, B.N. *et al.* (2019) Human–machine partnership with artificial intelligence for chest radiograph diagnosis. *NPJ Digit. Med.* 2, 111
7. Surowiecki, J. (2004) *The Wisdom of Crowds*, Anchor
8. Guilbeault, D. *et al.* (2018) Social learning and partisan bias in the interpretation of climate trends. *Proc. Natl. Acad. Sci. U. S. A.* 115, 9714–9719
9. Jayles, B. *et al.* (2017) How social information can improve estimation accuracy in human groups. *Proc. Natl. Acad. Sci. U. S. A.* 114, 12620–12625
10. Kurvers, R. *et al.* (2016) Boosting medical diagnostics by pooling independent judgments. *Proc. Natl. Acad. Sci. U. S. A.* 113, 8777–8782
11. Fihn, S.D. (2019) Collective intelligence for clinical diagnosis—are 2 (or 3) heads better than 1? *JAMA Netw. Open* 2, e191071
12. Nofer, M. and Hinz, O. (2014) Are crowds on the internet wiser than experts? The case of a stock prediction community. *J. Bus. Econ.* 84, 303–338
13. Sjoberg, L. (2009) Are all crowds equally wise? A comparison of political election forecasts by experts and the public. *J. Forecast.* 28, 1–18
14. Sunstein, C.R. (2006) *Infotopia: How Many Minds Produce Knowledge*, Oxford University Press
15. Herzog, S.M. and Hertwig, R. (2011) The wisdom of ignorant crowds: predicting sport outcomes by mere recognition. *Judgm. Decis. Mak.* 6, 58–72
16. Wolf, M. *et al.* (2015) Collective intelligence meets medical decision-making: the collective outperforms the best radiologist. *PLoS One* 10, e0134269
17. Kämmer, J.E. *et al.* (2017) The potential of collective intelligence in emergency medicine: pooling medical students' independent decisions improves diagnostic performance. *Med. Decis. Mak.* 37, 715–724
18. Barnett, M.L. *et al.* (2019) Comparative accuracy of diagnosis by collective intelligence of multiple physicians vs individual physicians. *JAMA Netw. Open* 2, e190096
19. Kelley, E.K. and Tetlock, P.C. (2013) How wise are crowds? Insights from retail orders and stock returns. *J. Finance* 68, 1229–1265
20. Shore, J. *et al.* (2015) Facts and figuring: an experimental investigation of network structure and performance in information and solution spaces. *Organ. Sci.* 26, 1432–1446
21. Brackbill, D. and Centola, D. (2020) Impact of network structure on collective learning: an experimental study in a data science competition. *PLoS One* 15, e0237978
22. Kao, A.B. and Couzin, I.D. (2014) Decision accuracy in complex environments is often maximized by small group sizes. *Proc. R. Soc. B Biol. Sci.* 281, 20133305
23. Sosna, M. *et al.* (2019) Individual and collective encoding of risk in animal groups. *Proc. Natl. Acad. Sci. U. S. A.* 116, 20556–20561
24. Couzin, I. (2007) (2007) Collective minds. *Nature* 445, 715
25. Kauffman, S. and Weinberger, E. (1989) The NK model of rugged fitness landscapes and its application to the maturation of the immune response. *J. Theor. Biol.* 141, 211–245
26. Baum, D. (2005) *New Yorker 17 January*. pp. 42–48
27. Simon, H.A. (1962) The architecture of complexity. *Proc. Am. Philos. Soc.* 106, 467–482
28. March, J.G. (1991) Exploration and exploitation in organizational learning. *Organ. Sci.* 2, 71–87
29. Uzzi, B. and Spiro, J. (2005) Collaboration and creativity: the small world problem. *Am. J. Sociol.* 111, 447–504
30. Lazer, D. and Friedman, A. (2007) The network structure of exploration and exploitation. *Adm. Sci. Q.* 52, 667–694
31. Condorcet, N. (1785) *Essai sur l'Application de l'Analyse à la Probabilité des Décisions Rendues à la Pluralité de Voix*, Imprimerire Royale
32. Bazazi, S. *et al.* (2019) Self-serving incentives impair collective decisions by increasing conformity. *PLoS One* 14, e0224725
33. Hahn, U. (2022) Collectives and epistemic rationality. *Top. Cogn. Sci.* 14, 602–620
34. Pfeifer, P.E. *et al.* (2014) The promise of prediction contests. *Am. Stat.* 68, 264–270
35. Bonabeau, E. (2009) Decisions 2.0: the power of collective intelligence. MIT Sloan Manag. Rev. 50, pp. 45-52.
36. Centola, D. *et al.* (2021) The reduction of race and gender bias in clinical treatment recommendations using clinician peer networks in an experimental setting. *Nat. Commun.* 12, 6585
37. Galton, F. (1907) Vox populi. *Nature* 75, 450–451
38. Fishkin, J.S. and Luskin, R.C. (2005) Experimenting with a democratic ideal: deliberative polling and public opinion. *Acta Politica* 40, 284–298
39. Lorenz, J. *et al.* (2011) How social influence can undermine the wisdom of crowd effect. *Proc. Natl. Acad. Sci. U. S. A.* 108, 9020–9025
40. Frey, V. and van de Rijt, A. (2020) Social influence undermines the wisdom of the crowd in sequential decision making. *Manag. Sci.* 67, 3985–4642
41. Solomonoff, R. and Rapoport, A. (1951) Connectivity of random nets. *Bull. Math. Biophys.* 13, 107–117
42. Guilbeault, D. and Centola, D. (2021) Topological measures for identifying and predicting the spread of complex contagions. *Nat. Commun.* 12, 4430
43. DeMarzo, P.M. *et al.* (2003) Persuasion bias, social influence, and unidimensional opinions. *Q. J. Econ.* 118, 909–968
44. Newman, M. (2010) *Networks: An Introduction*, Oxford University Press
45. Malone, T.W. and Bernstein, M.S. (2015) *Handbook of Collective Intelligence*, MIT Press
46. Becker, J. *et al.* (2017) Learning is robust to noise in decentralized networks. *Proc. Natl. Acad. Sci. U. S. A.* 114, E9184
47. Bavelas, A. (1950) Communication patterns in task-oriented groups. *J. Acoust. Soc. Am.* 22, 725–730
48. Rapoport, A. (1953) Spread of information through a population with socio-structural bias. I. Assumption of transitivity. *Bull. Math. Biophys.* 15, 523–533
49. Rapoport, A. and Horvath, W. (1961) A study of a large sociogram. *Behav. Sci.* 6, 279–291
50. Schwartz, P. *et al.* (1984) *Crosscutting Social Circles*, Academic Press
51. Granovetter, Mark (1973) The strength of weak ties. *Am. J. Sociol.* 78, 1360–1380
52. Milgram, S. (1967) The small world. *Psychol. Today* 2, 60–67
53. Burt, R. (1992) *Structural Holes: The Social Structure of Competition*, Harvard University Press
54. Centola, D. and Macy, M. (2007) Complex contagions and the weakness of long ties. *Am. J. Sociol.* 113, 702–734
55. Hills, T.T. *et al.* (2015) Exploration versus exploitation in space, mind, and society. *Trends Cogn. Sci.* 19, 46–54
56. Bernstein, E. *et al.* (2018) How intermittent breaks in interaction improve collective intelligence. *Proc. Natl. Acad. Sci.* 115, 8734–8739
57. Levinthal, D. and March, J.G. (1981) A model of adaptive organizational search. *J. Econ. Behav. Organ.* 2, 307–333
58. Mason, W. and Watts, D.J. (2012) Collaborative learning in networks. *Proc. Natl. Acad. Sci. U. S. A.* 109, 764–769
59. Bonaccio, S. and Dalal, R.S. (2006) Advice taking and decision-making: an integrative literature review, and implications for the organizational sciences. *Organ. Behav. Hum. Decis. Process.* 101, 127–151
60. Centola, D. (2018) *How Behavior Spreads*, Princeton University Press
61. Russell, S.J. and Norvig, P. (2009) *Artificial Intelligence: A Modern Approach* (3rd edn), Prentice Hall
62. Barkoczi, D. and Galesic, M. (2016) Social learning strategies modify the effect of network structure on group performance. *Nat. Commun.* 7, 13109
63. Almaatouq, A. *et al.* (2020) Adaptive social networks promote the wisdom of crowds. *Proc. Natl. Acad. Sci. U. S. A.* 117, 11379–11386

64. Centola, D. (2021) *Change: How to Make Big Things Happen*, Little, Brown & Co
65. Songhori, M.J. and García-Díaz, C. (2018) Collective Problem-Solving in Evolving Networks: An Agent-Based Model. In *2018 Winter Simulation Conference (WSC)*, pp. 965–976, IEEE
66. Almaatouq, A. *et al.* (2021) Task complexity moderates group synergy. *Proc. Natl. Acad. Sci. U. S. A.* 118, 36
67. Baran, P. (1962) On Distributed Communications Networks. In *RAND Corporation papers*, pp. 2626, IEEE
68. de Sola Pool, I. (1983) *Technologies of Freedom*, Harvard University Press
69. Hansen, M.T. (1999) The search-transfer problem: the role of weak ties in sharing knowledge across organization subunits. *Adm. Sci. Q.* 44, 82–111
70. Ancona, D.G. and Caldwell, D. (1990) Beyond boundary spanning: managing external dependence in product development teams. *J. High Technol. Manag. Res.* 1, 119–135
71. Elster, J., Landemore, H., eds (2012) *Collective Wisdom: Principles and Mechanisms*, Cambridge University Press
72. Kuncheva, L. *et al.* (2003) Limits on the majority vote accuracy in classifier fusion. *Pattern Anal. Appl.* 6, 22–31
73. Galesic, M. *et al.* (2018) Asking about social circles improves election predictions. *Nat. Hum. Behav.* 2, 187–193
74. Laan, A. *et al.* (2017) Rescuing collective wisdom when the average group opinion is wrong. *Front. Robot. AI* 4, 56
75. Kao, A.B. *et al.* (2018) Counteracting estimation bias and social influence to improve the wisdom of crowds. *J. R. Soc. Interface* 15, 4977653
76. McAndrew, T. *et al.* (2021) Aggregating predictions from experts: a review of statistical methods, experiments, and applications. *WIREs: Comput. Stat.* 13, e1514
77. Farrell, S. (2011) Social influence benefits the wisdom of individuals in the crowd. *Proc. Natl. Acad. Sci. U. S. A.* 108, E625
78. Guilbeault, D. *et al.* (2021) Probabilistic social learning improves the public's judgments of news veracity. *PLoS One* 16, 3
79. Becker, J.A. *et al.* (2021) The crowd classification problem: social dynamics of binary-choice accuracy. *Manag. Sci.* 68, 3949–3965
80. Janis, I.L. (1982) *Groupthink: Psychological Studies of Policy Decisions and Fiascoes*, Houghton Mifflin Boston
81. Gürçay, B. *et al.* (2015) The power of social influence on estimation accuracy. *J. Behav. Decis. Mak.* 28, 250–261
82. Golub, B. and Jackson, M.O. (2010) Naive learning in social networks and the wisdom of crowds. *Am. Econ. J. Microecon.* 2, 112–149
83. DeGroot, M.H. (1974) Reaching a consensus. *J. Am. Stat. Assoc.* 69, 118–121
84. Bala, V. and Goyal, S. (1998) Learning from neighbours. *Rev. Econ. Stud.* 65, 595–621
85. Mossel, E. *et al.* (2015) Strategic learning and the topology of social networks. *Econometrica* 83, 1755–1794
86. Becker, J. *et al.* (2017) Network dynamics of social influence in the wisdom of crowds. *Proc. Natl. Acad. Sci. U. S. A.* 114, 5070–5076
87. Becker, J. *et al.* (2019) The wisdom of partisan crowds. *Proc. Natl. Acad. Sci. U. S. A.* 116, 10717–10722
88. Madirolas, G. and de Polavieja, G.G. (2015) Improving collective estimations using resistance to social influence. *PLoS Comput. Biol.* 11, e1004594
89. Bail, C.A. *et al.* (2018) Exposure to opposition views on social media can increase political polarization. *Proc. Natl. Acad. Sci. U. S. A.* 115, 9216–9221
90. Guilbeault, D. and Centola, D. (2020) Networked collective intelligence improves dissemination of scientific information regarding smoking risks. *PLoS One* 15, e0227813
91. Poland, B. *et al.* (2006) The social context of smoking: the next frontier in tobacco control? *Tob. Control.* 15, 59–63
92. Mercken, L. *et al.* (2010) Dynamics of adolescent friendship networks and smoking behavior. *Soc. Net.* 32, 72–82
93. Sukhera, J. *et al.* (2020) Implicit bias in health professions: from recognition to transformation. *Acad. Med.* 95, 717–723
94. Healey, M.P. *et al.* (2021) Costs of collective wisdom: how resources influence information aggregation in organizational decision making. *Strateg. Organ.* Published online April 13, 2021. https://doi.org/10.1177/14761270211003849
95. Wu, L. *et al.* (2019) Large teams develop and small teams disrupt science and technology. *Nature* 566, 378–382
96. Xu, F. *et al.* (2022) Flat teams drive scientific innovation. *Proc. Natl. Acad. Sci.* 119, e2200927119
97. Centola, D. (2019) Influential networks. *Nat. Hum. Behav.* 3, 664–665
98. Centola, D. (2021) Influencers, backfire effects and the power of the periphery. In *Personal Networks: Classic Readings and New Directions in Ego-centric Analysis* (Brea, L. *et al.*, eds), Cambridge University Press
99. Young, H.P. (1993) The evolution of conventions. *Econometrica* 61, 57–84
100. Centola, D. and Baronchelli, A. (2015) The spontaneous emergence of conventions: an experimental study of cultural evolution. *Proc. Natl. Acad. Sci. U. S. A.* 112, 1989–1994
101. Guilbeault, D. *et al.* (2021) Experimental evidence for scale-induced category convergence across populations. *Nat. Commun.* 12, 327
102. Ellison, G. (1993) Learning, local interaction, and coordination. *Econometrica* 61, 1047–1071
103. Montanari, A. and Saberi, A. (2010) The spread of innovations in social networks. *Proc. Natl. Acad. Sci. U. S. A.* 107, 20196–20201
104. Axelrod, R. (1984) *The Evolution of Cooperation*, Basic Books
105. Cohen, M.D. *et al.* (2001) The role of social structure in the maintenance of cooperative regimes. *Ration. Soc.* 13, 5–32
106. Centola, D. (2013) Homophily, networks, and critical mass: solving the start-up problem in large group collective action. *Ration. Soc.* 25, 3–40
107. Pennycook, G. and Rand, D.G. (2019) Fighting misinformation on social media using crowdsourced judgments of news source quality. *Proc. Natl. Acad. Sci. U. S. A.* 116, 2521–2526
108. Moscovici, S. and Zavalloni, M. (1969) The group as a polarizer of attitudes. *J. Pers. Soc. Psychol.* 12, 125–135
109. Sunstein, C.R. (2002) The law of group polarization. *J. Polit. Philos.* 10, 175–195
110. Vosoughi, S. *et al.* (2018) The spread of true and false news online. *Science* 359, 1146–1151
111. Broniatowski, D. *et al.* (2018) Weaponized health communication: Twitter bots and Russian trolls amplify the vaccine debate. *Am. J. Public Health* 108, 1378–1384
112. Banerjee, A.V. (1992) A simple model of herd behavior. *Q. J. Econ.* 107, 797–817
113. Bikhchandani, S. *et al.* (1992) A theory of fads, fashion, custom, and cultural change as informational cascades. *J. Polit. Econ.* 100, 992–1026
114. Pauker, S.G. and Kassirer, J.P. (1980) The threshold approach to clinical decision making. *N. Engl. J. Med.* 302, 1109–1117
115. Poses, R.M. *et al.* (1990) Are two (inexperienced) heads better than one (experienced) head? Averaging house officers' prognostic judgments for critically ill patients. *Arch. Intern. Med.* 150, 1874–1878
116. Kattan, M.W. *et al.* (2016) The wisdom of crowds of doctors: their average predictions outperform their individual ones. *Med. Decis. Mak.* 36, 536–540
117. King, A.J. *et al.* (2013) Self-examinations and visual identification of atypical nevi: comparing individual and crowdsourcing approaches. *Cancer Epidemiol.* 37, 979–984
118. Sims, M.H. *et al.* (2014) Crowdsourcing medical expertise in near real time. *J. Hosp. Med.* 9, 451–456
119. Schultze, T. *et al.* (2012) Why groups perform better than individuals at quantitative judgment tasks: group-to-individual transfer as an alternative to differential weighting. *Organ. Behav. Hum. Decis. Process.* 118, 24–36
120. Bahrami, B. *et al.* (2010) Optimally interacting minds. *Science* 329, 1081–1085
121. Koriat, A. (2012) The self-consistency model of subjective confidence. *Psychol. Rev.* 119, 80–113
122. Hume, D. (1740/1976) *A Treatise of Human Nature* (3rd edn), Clarendon Press
123. Dewey, J. (1966) *Democracy and Education: An Introduction to the Philosophy of Education*, The Free Press
124. Granovetter, M. (1978) Threshold models of collective behavior. *Am. J. Sociol.* 83, 1420–1443
125. Schelling, T. (1978) *Micromotives and Macrobehavior*, Norton

126. Axelrod, R. (1986) An evolutionary approach to norms. *Am. Polit. Sci. Rev.* 80, 1095–1111
127. Elster, J. (1989) *The Cement of Society: A Survey of Social Order*, Cambridge University Press
128. Bicchieri, C. (2006) *The Grammar of Society: The Nature and Dynamics of Social Norms*, Cambridge University Press
129. Watts, D.J. (1999) Networks, dynamics, and the small-world phenomenon. *Am. J. Sociol.* 105, 493–527
130. Sober, E. and Wilson, D.S. (1998) *Unto Others: The Evolution and Psychology of Unselfish Behavior*, Harvard University Press
131. Ryan, B. and Gross, N. (1943) The diffusion of hybrid seed corn in two Iowa communities. *Rural. Sociol.* 8, 15
132. Coleman, J. *et al.* (1957) The diffusion of an innovation among physicians. *Sociometry* 20, 253–270
133. Davis, G.F. and Greve, H.R. (1997) Corporate elite networks and governance changes in the 1980s. *Am. J. Sociol.* 103, 1–37
134. Young, H.P. (2011) The dynamics of social innovation. *Proc. Natl. Acad. Sci. U. S. A.* 108, 21285–21291
135. Centola, D. (2010) The spread of behavior in an online social network experiment. *Science* 329, 1194–1197
136. Centola, D. (2011) An experimental study of homophily in the adoption of health behavior. *Science* 334, 1269–1272
137. Rosenthal, S.R. *et al.* (2015) Behavioral contagion in mobile animal groups. *Proc. Natl. Acad. Sci. U. S. A.* 112, 4690–4695
138. Rosenberg, L. (2015) Human swarms, a real time method for collective intelligence. *Proc. Eur. Conf. Artif. Life* 658–665
139. Mackay, C. (1932) *Extraordinary Popular Delusions and the Madness of Crowds*, Richard Bentley
140. Habermas, J. (1984) *Theory of Communicative Action, Volume One: Reason and the Rationalization of Society*, Beacon Press
141. Thompson, D.F. and Gutmann, A. (2009) *Why Deliberative Democracy?*, Princeton University Press
142. Delli Carpini, M.X. and Keeter, S. (1996) *What Americans Know about Politics and Why It Matters*, Yale University Press
143. Mill, J.S. (1859) *On Liberty*, John W. Parker and Son
144. Madison, J. (1787) *New York Daily Advertiser 22 November*.
145. Arrow, K.J. (1950) A difficulty in the concept of social welfare. *J. Polit. Econ.* 58, 328–346
146. Cohen, J. (1989) Delibration and democratic legitimacy. In *Debates in Contemporary Political Philosophy: An Anthology* (Matravers, D. and Pike, J.E., eds), Routledge, in Association with the Open University
147. Grofman, B. *et al.* (1983) Thirteen theorems in search of the truth. *Theor. Decis.* 15, 261–278
148. Girard, P. *et al.* (2015) Belief diffusion in social networks. *Inform. Syst. Behav. Soc. Methods* Published online June 13, 2015. https://doi.org/10.2139/ssrn.2617554
149. Axelrod, R. (1997) The dissemination of culture: a model with local convergence and global polarization. *J. Confl. Resolut.* 41, 203–226
150. Durrett, R. and Levin, S.A. (2005) Can stable social groups be maintained by homophilous imitation alone? *J. Econ. Behav. Organ.* 57, 267–286
151. Wooley, A.W. *et al.* (2010) Evidence for a collective intelligence factor in the performance of human groups. *Science* 330, 686–688
152. Reagans, R. and Zuckerman, E.W. (2001) Networks, diversity, and productivity: the social capital of corporate R&D teams organization. *Science* 12, 502–517
153. Hong, Lu and Page, S.E. (2004) Groups of diverse problem solvers can outperform groups of high-ability problem solvers. Proc. Natl. Acad. Sci. U. S. A. 101, 16385-16389.
154. Feng, S. *et al.* (2019) The wisdom of polarized crowds. *Nat. Hum. Behav.* 3, 329–336
155. Guimerà, R. *et al.* (2005) Team assembly mechanisms determine collaboration network structure and team performance. *Science* 308, 697–702
156. Gomez, C.J. and Lazer, D.M.J. (2019) Clustering knowledge and dispersing abilities enhances collective problem solving in a network. *Nat. Commun.* 10, 5146
157. Green, K. *et al.* (2007) Methods to elicit forecasts from groups: Delphi and prediction markets compared. *Foresight Int. J. Appl. Forecast.* 8, 17–20
158. Wolfers, J. and Zitzewitz, E. (2004) Prediction markets. *J. Econ. Perspect.* 18, 107–126
159. Hogarth, R.M. (1978) A note on aggregating opinions. *Organ. Behav. Hum. Perform.* 21, 40–46
160. Mellers, B. *et al.* (2014) Psychological strategies for winning a geopolitical forecasting tournament. *Psychol. Sci.* 25, 1106–1115
161. Bennett, S.T. *et al.* (2018) Making a wiser crowd: benefits of individual metacognitive control on crowd performance. *Comp. Brain Behav.* 1, 90–99
162. Merkle, E.C. *et al.* (2017) A neglected dimension of good forecasting judgment: the questions we choose also matter. *Int. J. Forecast.* 33, 817–832
163. Kameda, T. and Nakanishi, D. (2003) Does social/cultural learning increase human adaptability? Rogers's question revisited. *Evol. Hum. Behav.* 24, 242–260
164. Navajas, J. *et al.* (2018) Aggregated knowledge from a small number of debates outperforms the wisdom of large crowds. *Nat. Hum. Behav.* 2, 126–132
165. Caplan, B. (2007) *The Myth of the Rational Voter: Why Democracies Choose Bad Policies*, Princeton University Press
166. Winkler, R.L. and Poses, R.M. (1993) Evaluating and combining physicians' probabilities of survival in an intensive care unit. *Manag. Sci.* 39, 1526–1543
167. Conradt, L. and Roper, T.J. (2003) Group decision-making in animals. *Nature* 421, 155–158
168. Galesic, M. *et al.* (2022) Beyond collective intelligence: collective adaptation. *SocArXiv* Published online January 15, 2022. https://doi.org/10.31235/osf.io/5f2ad
169. Herzog, S.M. *et al.* (2019) The ecological rationality of the wisdom of crowds. In *Taming Uncertainty* (Hertwig, R. *et al*., eds), pp. 245–262, MIT Press
170. Turner, M.E. and Horvitz, T. (2001) The dilemma of threat: group effectiveness and ineffectiveness under adversity. In *Groups at Work* (Turner, M.E., ed.), Psychology Press
171. Akçay, E. (2018) Collapse and rescue of cooperation in evolving dynamic networks. *Nat. Commun.* 9, 2692
172. Burton, J.W. *et al.* (2021) Rewiring the wisdom of the crowd. In *Proceedings of the Annual Meeting of the Cognitive Science Society* (Vol. 43)