The Benefits of Network Centralization for Collective Intelligence in Exploration/Exploitation Problems

Jesse Shore,[a] Ethan Bernstein,[b] and Alice Jang[a]

[a]Boston University; [b]Harvard Business School

September 15, 2019

Abstract

This paper examines how centralization of organizational communication networks impacts problem solving performance on tasks requiring a shift from one inferior solution to a dissimilar, novel one—or, in the language of March, tasks that require exploration rather than just exploitation. Drawing on a 1,620-subject experiment using a new platform designed to run randomized experiments with network structure as the independent variable, we tested the effect of seven network structures—representing different levels of centralization—on problem solving success. To simulate a dynamic environment with shifting information, we designed a murder mystery task and manipulated when each piece of information could be found: early information encourages an incorrect consensus, requiring a collective shift of solution when more information emerged later. We find that when the communication network within an organization is more centralized, it achieves the benefits of social influence (learning) without the costs (herding). In centralized networks, peripheral nodes are relatively independent and less likely to get stuck at an earlier consensus solution (reduced herding). Central nodes then learn from members of the periphery and quickly spread the correct answer to others in the network (increased learning). We also find, however, that these benefits of centralization come with a major caveat: they only materialize in undirected networks, not in directed core-periphery structures that typify self-organized networks like digitally-mediated ones. We draw on these findings to reconceptualize existing theory on the impact of centralization on collective intelligence in problem solving that demands exploration and adaptation by an organizations members.

1 Introduction

How does the centralization of communication network structures within an organization affect the collective intelligence—or capacity to collectively solve problems—of that organization's members? As solving problems is a key driver of organizational success (Grant 1996, Nickerson and Zenger 2004) and network structure is known to be a determinant of collective problem-solving capability (Mason and Watts 2012, Becker et al. 2017, Argote et al. 2018), interest has grown among organizational theorists and managers in understanding the link between collective intelligence and network "centralization," defined as the degree to which communication tends to flow disproportionately through one or more members of the organization rather than being equally shared among all members.

On the one hand, centralized communication has been shown to enhance coordination efficiency (Bavelas 1950, Tushman 1979, Butts et al. 2007); because communication must pass through the network "core" (i.e. central nodes in the network), centralization all but guarantees that someone in the core will have access to all critical information. On the other hand, that has been shown to come at a cost to quality of collective judgment and collaboration success (Tushman 1979, Huang and Cummings 2011, Becker et al. 2017, Argote et al. 2018), because the core can easily be imperfect, acting more as a troublesome bottleneck than a helpful filter. As a result of such established shortcomings of the network core, managers today increasingly use enterprise social network technologies (Majchrzak et al. 2009, Leonardi 2014) to facilitate the emergence of informal decentralized information flows that bypass the centralized formal structures of traditional bureaucratic organizations (Gulati and Puranam 2009). In doing so, organizations have attempted to halt two mechanisms that prior research has suggested links centralization with impaired collective judgment. First, people in the network periphery cannot share knowledge or judgements directly with each other, which limits the extent they can learn from each other or build on each other's ideas. Second, the central node or nodes influence many people in the periphery, which can quickly spread any judgements, including bad ones.

We present experimental evidence that these same two negative effects of network centralization can be *beneficial* for collective intelligence for a particular, yet fast-growing, class of problem-solving settings: those in which organizations need to adapt (Christensen 2013), shifting from one seemingly correct, established solution (which turns out to be no longer accurate) to a dissimilar, novel solution (drawn from a diversity of new ideas rather than a consensus of past experience). Using the language of organizational theory, we present evidence that centralization can be beneficial when problem solving requires what March (1991, pp. 71, 85) called exploration ("captured by terms such as search, variation, risk taking, experimentation, play, flexibility, discovery, innovation" from which "returns are uncertain, distant, and often negative") versus exploitation ("such things as refinement, choice, production, efficiency, selection, implementation, execution") because of the core's unique capacity to operate as a bottleneck to the periphery sharing information directly with each other, preventing detrimental social influence.

We ran a 1620-person experiment on a new experimental platform that we developed for running randomized experiments with network structure as the independent variable (i.e., investigating how network structure causally impacts collective performance). To collectively solve problems, participants could search for information, share information they found with each other, and communicate with free-text, much as they would in organizational problem-solving settings enabled by contemporary enterprise social network technologies, as well as enter their solutions and access their neighbors' solutions to the problems posed in the experiment. We designed the experimental platform to permit the creation of a shifting environment (analogous to dynamic environments faced by product managers and organization leaders) by controlling when certain pieces of information can be discovered during the course of a single run of the experiment. We used this feature to create a collaborative murder-mystery task with a pivotal phase in which the problem-solving requires the participants to collectively shift away from a premature consensus on an incorrect answer, adapt to new information, and adopt a new solution contrary to what was previously believed to be true. To study the main effect of network centralization on problem-solving performance and describe the mechanisms through which it works, we tested seven realistic network structures (four variations on centralized network structures and three controls) with participants randomly assigned to nodes, so that our results would be robust to the attributes of the individuals in the core.

We find that, in problem solving settings requiring both exploration and exploitation, the inability of peripheral nodes to directly influence each other preserves a degree of independence, and increases the likelihood of having a good idea that contradicts the majority opinion. And central nodes not only influence many people, but they can also play the role of a filter—identifying, adopting and spreading promising ideas in the periphery.

Our findings suggest that while it remains true that central nodes influence many people, and that they can represent a bottleneck that prevents the recombination of ideas in the periphery (Huang and Cummings 2011), they can also use their special position to learn from many people who have been kept relatively independent from one another. The same features that make centralized communication structures hinder the completion of a project or interfere with the Wisdom of the Crowds phenomenon make them resilient and adaptable to new information.

By testing both directed and undirected network ties, we also provide a major caveat to that core finding. The benefits to centralization rely on central nodes learning from the more independent peripheral nodes. However, many centralized network structures within organizations are asymmetrical (i.e. have directed ties), in which central nodes influence but are not influenced by the periphery, like traditional bureaucratic/administrative pyramidal structures (Puranam 2018) and thus do not allow the benefits of centralization to occur.

Our work addresses a major current question in organizational theory on the role of centralized structures in collective intelligence. The findings on centralization directly contribute to both network theory and organization design, while also adding to the growing literature on the fit between organization structure and the problems being solved, and specifically to the literature on how to organize to promote social learning without succumbing to lock-in on existing ideas.

2 The Benefits and Costs of Centralization for Different Task Types

Centralized networks contain one or more "core" nodes—nodes through which communication disproportionately flows—and prior research on the value of centralization has been focused on the mechanisms through which the core generates benefits (or costs) to problem-solving in various types of tasks. We review prior findings about the impact of centralization on the three major categories of tasks that have been studied–coordination tasks, wisdom of crowd tasks, and collaboration tasks–all three of which, as executed in prior work, have generally fallen into March's (1991) definition of exploitation tasks. From each of the three task categories, we identify the mechanism through which centralization benefits or detriments problem-solving performance. We then use those mechanisms to develop hypotheses around our research question: how centralization might affect problem solving in tasks marked by the kind of exploration required for adaptation characterized by a major shift from one consensus to a new solution.

2.1 The Benefits of Centralization in Coordination Tasks

Coordination tasks, such as ensuring that everybody shows up at the same time for a meeting, or allocating ambulances to different accidents such that they do not converge on a single accident, are based on division or standardization of labor and facilitated by at least one individual having full information (Puranam 2018). In executing such tasks collectively, the existence of a central node through which all information passes ensures that at least one person knows everything and can "put the pieces together" and communicate essential information to the rest of the network (Bavelas 1950, Mulder 1960, Hinds and McGrath 2006). Centralized networks are therefore associated with collective intelligence on coordination tasks (Tushman 1979), a result reinforced by the finding that under time pressure (when coordination becomes more important), people naturally form more centralized network structures (Butts et al. 2007).

These findings from prior literature on coordination tasks appear on the left side of Figure 1: because peripheral nodes are not directly interconnected in centralized networks, communication between them must pass through the core, increasing access to the periphery's information and, in the case of these tasks, ensuring that at least one node access to all information and can process it accordingly. The core can then use its influence over the periphery to usher it towards coordinated behavior.

2.2 The Costs of Centralization in Wisdom of the Crowd Tasks

Much research has investigated the phenomenon of the "wisdom of the crowd" – namely that the average of a group of independent estimates tends to be more accurate than the average individual estimate (Galton 1907). Recent work in this broad literature studies interference from social influence in establishing unbiased collective judgements (e.g. Salganik et al. 2006, Muchnik et al. 2013, Levine et al. 2014, Becker et al. 2017, van de Rijt 2019). The benefits that accrue from the wisdom of the crowd phenomenon rely on people's individual estimation errors being uncorrelated; herding (when the range of people's judgements about something coalesce, through

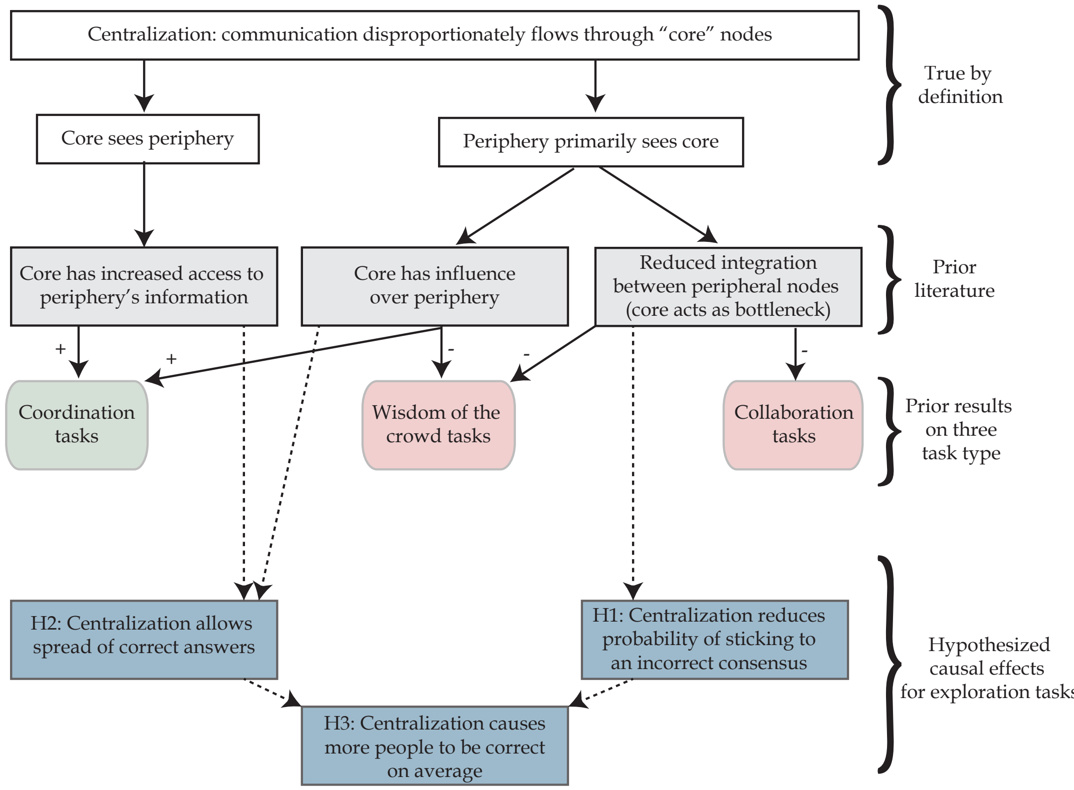

Figure 1: Theoretical diagram

communication and social influence, on an answer regardless of whether it is correct) results in correlated errors and can sometimes lead to a reduction of the wisdom of the crowd effect (Lorenz et al. 2011, Pan et al. 2012).

In practice, however, uncentralized social influence can be beneficial to collective accuracy in wisdom of crowd tasks, as less confident or knowledgeable individuals adopt better judgments from their peers (Bahrami et al. 2010, see also Tchernichovski et al. (2019), Madirolas and De Polavieja (2014), Becker et al. (2017)). Centralized social influence, on the other hand, interferes with crowd wisdom by subjecting many peripheral nodes to the estimate of the central node(s), which are not necessarily good (Becker et al. 2017). This finding echoes formal models of social influence (DeGroot 1974, Golub and Jackson 2010) that conclude that in the presence of social influence, crowd wisdom occurs only if no single individual is too influential. Although the network may contain diverse judgments, the out-sized influence of the core prevents judgments in the periphery from influencing the rest of the network.

These findings from prior literature on wisdom of the crowd tasks appear in Figure 1: because the periphery sees primarily the core, the core has an outsized influence over the periphery, which can harm performance on wisdom of the crowd tasks. Additionally, centralization reduces the capacity of the periphery to self-correct (van de Rijt 2019), as there is reduced integration, further harming wisdom of the crowd tasks.

2.3 The Costs of a Central Core in Group Collaboration Tasks

If coordination tasks benefit from division of labor (i.e., differentiation), collaboration-based tasks benefit from integration. Rather than simply staying out of each-other's way (as in Steiner's 1972 "additive" tasks), in collaboration tasks, people do different but complementary and interdependent subtasks. By developing shared tacit knowledge (Hansen 1999) and a "transactive memory system" (Liang et al. 1995, Lewis et al. 2005, Argote et al. 2018), collaborators more effectively perform complementary work and effectively use knowledge that is distributed among team members. As a result, groups composed of people who have a strong theory of mind (i.e. who reason about others' mental states) tend to be collectively intelligent (Woolley et al. 2010, Engel et al. 2014). Indeed the degree to which a group's behavior is "integrated" – as opposed to just a sum of independent individuals – has been proposed as a structural metric or indicator of its collective intelligence (Engel and Malone 2018).

Network centralization has been found to interfere with integration of knowledge held in the periphery, so centralization harms performance on collaborative tasks (Tushman 1979, Cummings and Cross 2003, Huang and Cummings 2011). Peripheral nodes cannot learn directly from one another, and because central nodes have a finite capacity for communication, they can act as a bottleneck. Thus, critical knowledge held by peripheral individuals may not be shared widely and integrated with others' knowledge, reducing collaborative performance tasks (Huang and Cummings 2011).

These findings from prior literature on collaboration tasks appear on the right side of Figure 1: because the periphery sees primarily the core, communication among peripheral nodes must pass through the core, making the core a bottleneck to integration and use of knowledge in the periphery, inhibiting success on collaboration tasks.

2.4 Hypotheses of the Benefits and Costs of Centralization for Tasks Requiring Exploration for Adaptation

Prior research investigating the impact of centralization on problem-solving outcomes has produced the three mechanisms above. However, all three categories of tasks share a key characteristic: they all fall under March's description of exploitation-related tasks—how to make best use of the information or knowledge already present within the network rather than the ability of network members to shift from one potential solution to another (i.e., an exploration task).

Recently, organization theory has placed particular emphasis on how to use collective intelligence to improve such exploration-related tasks (March 1991, Lavie et al. 2010, Levine et al. 2018). Whether in a problem-solving context involving creating (or defending against) a disruptive innovation (Christensen 2013) or simply trying to solve an escape room or adventure racing challenge (Barton and Sutcliffe 2018, Englmaier et al. 2018), benefits to greater exploration occur in creative or strategic work where improving a solution is not as simple as "hill-climbing" by refining an existing idea. Rather, good solutions may be dissimilar to one-another, and if one explores by making incremental changes, one is likely to go through a "valley" of bad ideas before finding another "peak" of better ideas. Additionally, dynamic environments introduce the same tension between making the most of current knowledge and adapting to new information (Lavie et al. 2010).

How is the presence of a central core likely to affect such exploration tasks? To shift from one consensus to a new solution, an organization must first detect the presence of information that undermines the pre-existing consensus. In all networks, social influence results in conformity pressure and makes it less likely for an organization's members to explore alternative solutions. Rather than strike out on their own in search of something new and potentially better, they exhibit a tendency to remain at a current solution, at which they can remain in agreement with their peers. But network structures that limit integration of peripheral members' ideas—like more centralized networks—can support greater exploration (Lazer and Friedman 2007): in centralized networks, peripheral individuals are less able to integrate their ideas with others in the network because the core is a bottleneck (Tushman 1979, Huang and Cummings 2011) and thus the collective may be less subject to conformity pressure, especially if disconfirming information surfaces. In essence, the core node or nodes in a centralized network are influenced by a relatively disconnected set of nodes and thus face relatively less conformity pressure than nodes in a less centralized network. Relatively less conformity pressure on peripheral nodes makes it more likely for them to adapt to new information by changing their solution, rather than getting stuck at the majority opinion. We thus hypothesize:

Hypothesis 1 *Centralization reduces the probability of sticking to an incorrect consensus (i.e., reduces the herding effect of social influence).*

Adaptation in problem solving through a shift from prior consensus to a new solution requires not just the surfacing of new ideas (variation) but also a collective recognition of that idea (selection) (March 1991). Put differently, there is the need for both independent exploration, to generate diversity of ideas, and the opportunity to influence and learn from others, to facilitate the spread and combination of those ideas (as in the intermittent interaction of Bernstein et al. 2018). If centralized networks are hypothesized to support the first need (Hypothesis 1), they can be hypothesized to support the second need through the influence of the core. The core, because it has access to all solutions adopted by the periphery, can more effectively select and spread anew solution when it is shown to be better. In essence, while social influence leads to herding behavior, it is also the mechanism by which *good* ideas spread, especially in the case of truly disconfirming information. In other words, social influence can be the vector for social "learning," through which a collective ultimately selects which good ideas to use and refine (March 1991). We thus hypothesize:

Hypothesis 2 *Centralization allows the spread of correct answers (maintains the learning effect of social influence).*

Taken together, the lack of herding on poor solutions, and the spread of good ones, should result in a higher average performance:

Hypothesis 3 *Centralized networks cause more people to be correct on average.*

3 Experimental design

Using a new online lab platform, we conducted a randomized experiment of the effect of network centralization in a problem-solving setting requiring shifting from one solution to another as new information emerges (see Figure 4 for network treatments). Our task was a fictional murder mystery, characterized by a prominent red herring and timed release of critical clues, designed such that information leading to the correct answer would only become available after participants had already come to a different conclusion or consensus. Thus, our task design focuses on the ability to shift from one idea to a dissimilar one in the context of different communication structures, and our main dependent variable is how many people per trial were correct at the end of the trial.

Our experimental interface (Figure 2) allowed participants for search for clues about the mystery by entering keywords into a search bar, to share any clues they find with network neighbors (with or without free text messages), send messages to network neighbors without clues, reply to or forward messages received from network neighbors, register one's own solution (along with a confidence that the answer is correct), and see the solutions (and confidence levels) of network neighbors.

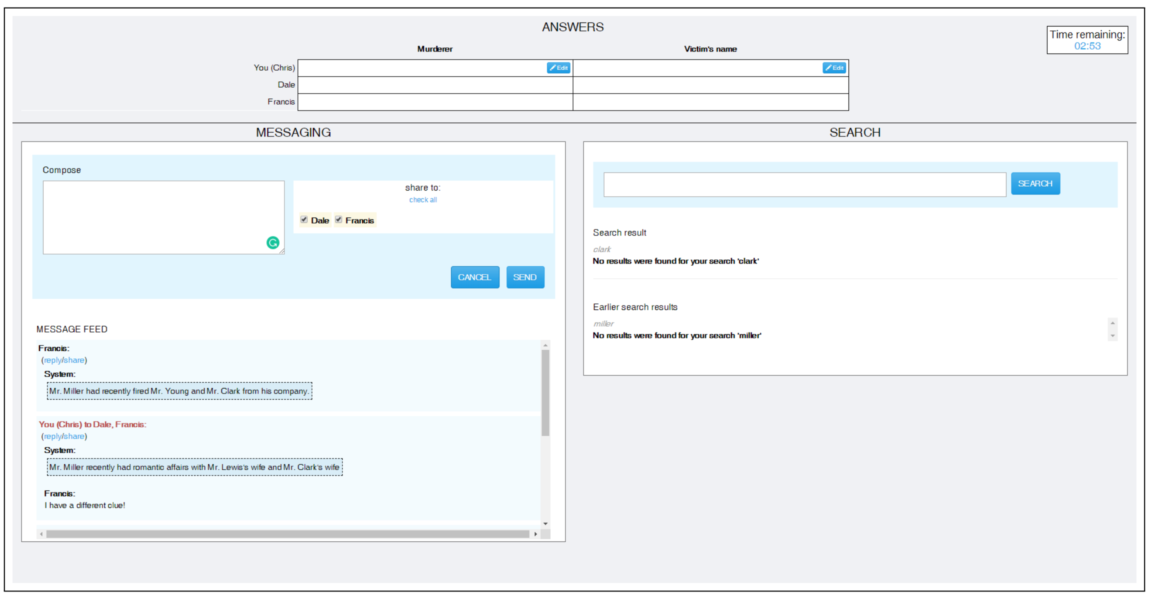

Figure 2: Screenshot of the experimental interface

3.1 Structure of task

At the beginning of the trial every participant received an initial clue in the form of a "message from headquarters." Subsequent clues could be found by conducting a search using keywords from the initial clue. For example, the initial clue indicates that the fictional murder victim passed away unexpectedly after dining at "Cafe Achilles." Searching for more information about Cafe Achilles produces clues about other diners who were there that night (i.e. possible suspects in the murder case); searching further with another diner's name as a search term produces more information about that suspect, including possible motives or connections to the murder victim. The first set of clues establishes a red herring by repeatedly bringing up a specific suspect and establishing a motive for that suspect to commit the crime.

Each trial lasted for fifteen minutes and was divided into three effective phases. As each new phase began, new information became accessible. After five minutes had elapsed, a new message from headquarters was broadcast to all participants that revealed new information (and therefore new search terms) that led to information that was not accessible during the first five minutes of the experiment. With sufficient search and deduction, the information available during this second phase of the experiment was sufficient to rule out the main red herring and discover the real culprit. At the ten minute mark, a new "message from headquarters" was broadcast with a critical clue that simply and logically ruled out the most prominent red herring from earlier in the experiment.

Figure 3 shows the overall percentage of participants with each solution registered over time and illustrates the predominant switch away from the red herring and toward the correct answer in the last five minutes of the experiment. During the second phase, the majority of participants had registered the red herring as the solution and fewer than 10% of people had identified the correct answer, even though it was logically possible to do so. During the third phase, when it was logically straightforward to deduce the correct answer, a sizeable minority retained incorrect answers.

3.2 Network Treatments

We examine six different network treatments with nine nodes (individual participants) each. The two centralized networks that are the focus of our study are the hub-and-spoke network and the core-periphery network. The hub-and-spoke network is the maximally-centralized network, with a single central node to which all remaining nodes are connected. The hub-and-spoke structure is also a representation of the communication pattern induced by hierarchical organization.

The core-periphery network structure is a less centralized structure that is nevertheless an empirically common form of centralization in self-organized social networks (see e.g. Colizza et al. 2006, Cattani and Ferriani 2008, Dahlander and Frederiksen 2012, Connelly et al. 2011, Kane and Alavi 2007) consisting of a central group (or "core") rather than a single central individual and peripheral individuals connected to the core but not each other (Borgatti and Everett 2000). The

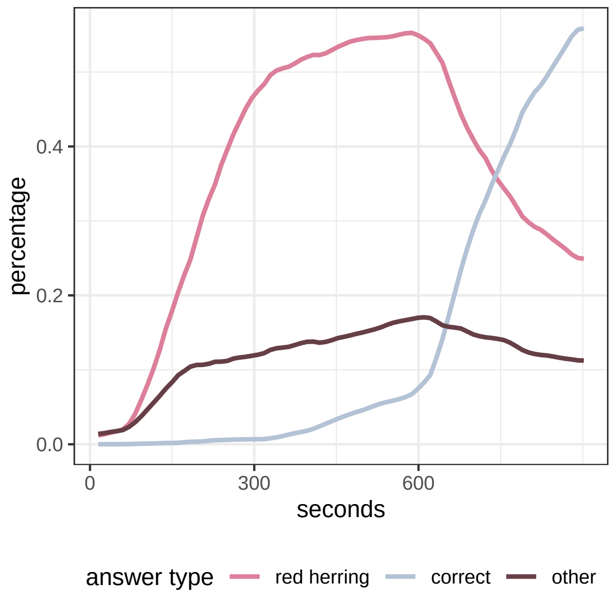

Figure 3: Percentage of participants with each solution registered over time. "Other" includes four different suspects.

existence of a central group is practically important because of the greater conformity pressure felt in network positions with a clustered pattern of ties (Centola 2010).

In addition to these two focal centralized networks, we also tested two versions of the core-periphery network with directed ties for the purposes of assessing mechanisms underlying the effects of the undirected centralized networks. One, the out-only core-periphery network, differs from the undirected version of the core-periphery network in that the core has fewer incoming ties from the periphery. In addition to serving as a mechanism check, this network structure also occurs in self-organized networks in which influence is not bilateral; digital social media in which one user can unilaterally 'follow' another can take this form (Frahm and Shepelyansky 2012). The second directed core-periphery network, the in-only core-periphery network, reverses the pattern such that the core has many incoming ties from the periphery but few outgoing ties.

Finally, we tested three uncentralized network structures as controls. The locally-clustered network consists of three interconnected cliques of three individuals, akin to interconnected teams or groups of friends. The complete clique is a group of nine people, all of whom are connected to each other, like a large workgroup or people using a shared digital communication channel. Finally, we constructed a synthetic treatment of "isolates," which are groups of nine individuals resampled with replacement from the set of all participants in the in-only core-periphery network who had no incoming ties (positions 5, 7, and 9). Isolates represent the disconnected nature of "crowds" (Afuah and Tucci 2012) or participants in a competition (Boudreau et al. 2011).

These three controls represent what we believe are empirically plausible alternatives to centralization, with different levels of social influence. Although random (e.g. Becker et al. 2017) or extremal (e.g. Mason and Watts 2012) network structures have sometimes been used as controls for network experiments, we attempted to create controls that could also function as plausible null hypotheses—that is, plausible alternative modes of organizing for problem solving.

We collected 30 independent trials for each experimental condition, and 90 bootstrap-sampled synthetic trials of isolates. Network treatments were collected in randomly-permuted order to avoid systematic biases due to time effects. Fewer than 2% of participants dropped out of trials mid-way, but this resulted in approximately 16% of trials being affected by participant drop out at some point. In estimating causal effects of treatments, we analyze all trials, regardless of participant drop-out to avoid inadvertently conditioning on post-treatment latent variables (Montgomery et al. 2018).[1]

3.3 Participant pipeline

We recruited participants through Amazon Mechanical Turk (www.mturk.com), located in the United States. Building on prior work by Mason and Suri (2012), we constructed a pipeline consisting of (1) a three-minute individual task in which participants learn to search for clues and enter their solutions, as well as the basic format of a murder-mystery task, (2) a five-minute task for three people in which participants learn to use the social features of the interface and finally (3) the main fifteen-minute trial for nine people.

Requiring two tasks to be complete before entering the main trial played the dual role of providing sufficient instruction (to ensure participants were familiar with the platform before the main experiment) and filtering out participants who were inattentive, disinterested, unable to complete the task, or otherwise likely to drop out in the middle of a synchronous trial. Low levels of attention among Amazon Mechanical Turk workers is a rational response to the preponderance of low-paying

[1] If drop-out is affected by experimental treatment in any way (e.g. if one treatment was more difficult and this caused more drop-out), then omitting trials risks biasing estimates of the causal effects of the treatment (e.g. if we omitted observations where people dropped out due to difficulty, our estimates would be positively biased because we ignored observations where the treatment caused people to have trouble).

Figure 4: Visualizations of network treatments

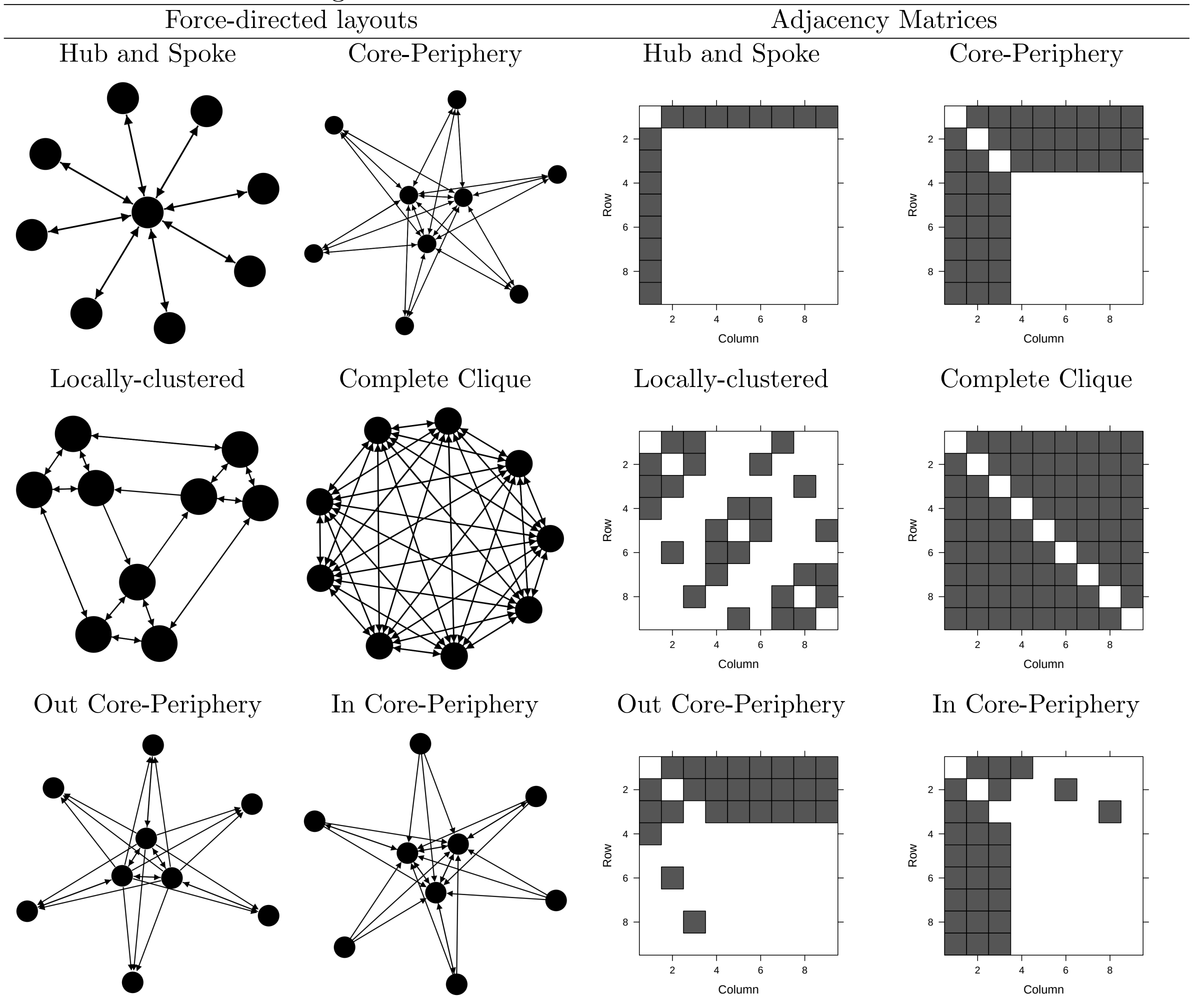

and poorly-designed tasks (Bigham 2014). However, drop-outs and inattentive participants affect the performance of everybody in the trial—not only the person who does not complete the task—so minimizing these was a priority. In addition, random drop-out (i.e. drop-out that is not caused by something specific to network treatment) is not a problem from a statistical inference perspective, but it introduces problems of validity of interpretation: if the central node in the hub-and-spoke network drops out, this creates a large de-facto change in the network structure.

Participants were eligible to move on to stage 2 if they were able to successfully complete stage 1 by following instructions to operate the search and answer features and performing a simple logical deduction on the six total clues that could be found (details on the logical task are in Appendix A). If participants were able to deduce the correct answer from the available information and enter their answer into the answer field, they were eligible to move on to the next stage. Similarly, participants were eligible to take part in the stage 3 (main) trial if they could successfully pass the second stage trial, which was a slightly longer task played synchronously with two other people. Additionally, this stage introduced an attention check question that was necessary to fill in to pass the trial. If participants did not successfully pass a stage 1 or stage 2 trial, they could try again up to four times before becoming disqualified from further attempts.

The three-stage trial pipeline had a steep drop-off in numbers of participants at each stage. Approximately 43% of instructional trials were successfully passed; of these 43%, 74% opted in to a stage 2 trial (i.e. the net number of people opting into the stage 2 trial was $43\% \times 74\% \approx 32\%$ of the number of people who opted into stage 1 trials). Approximately 80% of second stage trials were successfully passed; of these, 78% opted in to the main stage 3 trial. Overall, only 20% of subjects progressed all the way through a stage 3 trial, which we believe corresponds to effective filtering of attentive, interested, and capable participants.

Subjects were paid according to the length of time that they had the correct answers entered into the user interface (including one trivial attention-check question as well as the main question corresponding to the solution to the mystery). Median wages were $8.47 per hour.

Online labor market workers actively review "requesters" (people that post tasks) to help identify reputable employers (Gray et al. 2016). On Turkerview.com, our experiments had an average reported overall wage of $11.11 per hour, which includes pay for time in the waiting room, as well as stage 1, stage 2, and the main stage 3 trials. Our experiments received the overall rating of "Workers feel this requester pays fairly" as well as the "good communication," "approves quickly," "no rejections," and "no blocks" tags. We used MTurk's qualifications functionality (in which a requester can assign a worker a custom qualification upon successfully completing a task to allow them to see further tasks in a series) to qualify and disqualify workers for specific tasks. We never blocked workers or rejected their work, as these actions can have serious negative effects on their ability to qualify for other work on the MTurk platform. Full description of subject payments is in Appendix A.

4 Results

4.1 Main result

Centralized networks performed best in this experiment, in which it was important to adapt to new information and shift from one solution to a better one. Social influence in problems requiring exploration and exploitation introduces the possibility of both negative effects ("herding" on the wrong answer), and positive effects ("learning" the right answer from others). As we show below centralized networks minimized herding, while maximizing learning, thus getting the benefits without the costs of collaboration.

Table 1: Descriptive statistics by network treatment

	Hub-and-spoke	Core-Periphery	Isolates	Locally-Clust.	Comp. Clique	Out C-P	In C-P
N	30	30	90	30	30	30	30
mean(# correct)	4.7	4.367	3.189	3.2	3.533	3.133	4
sd(# correct)	2.706	3.079	1.381	2.497	3.391	2.417	1.965
median(# correct)	4.5	4	3	3	1	2.5	4

Basic descriptive statistics are in Table 1. Figure 5 provides an overview of the distribution of how many people were correct at the end of the trial, by experimental condition, in the form of empirical cumulative distribution functions (eCDFs). The further the eCDF curves are to the right-hand side of the panel, the better the performance is in terms of the percent of people correct per trial. Panel A shows the eCDFs for the isolates (minimal social influence) and complete clique (maximal social influence) conditions. Compared to nominal groups of isolates, the complete clique tended to have more trials with very few people correct (marked "herding" in the figure) and more trials with many people correct (marked "learning" in the trial). Subsequent panels overlay the eCDF from another network treatment on top of those of the isolates and complete clique for easier visual interpretation. Panel B shows results for the locally-clustered network, which has less herding and less learning than the complete clique, relative to isolated solvers.

Panels C and D show results for the centralized network treatments and panels E and F show results for the mechanism check treatments that are discussed further in section 4.3. Panel C shows that the hub-and-spoke network had few trials with few people correct (it did not experience the herding effect), but did have many trials with many people correct (it did benefit from social learning). Panel D shows that the core-periphery was subject to some herding effect, though less than the complete clique treatment, while enjoying similar levels of learning.

Models of the number of people correct at the end of a trial put these comparisons in a regression framework and are reported in Table 2. Because we are interested in distinguishing between learning and herding effects, we model the quantiles of the response variable (number of people correct at the end of the trial) as well as the mean. Quantile regressions for count variables are fit through the method of Machado and Silva (2005), in the R package `lqmm` (Geraci and Bottai 2014). Regression models for the mean number of people correct were fitted as quasipoisson generalized linear model to account for overdispersion from social influence. Appendix C discusses model choice at greater length.

Table 2: Number of people correct at end of trial

	Dependent variable:							
	25^{th} $\%^{ile}$	median	mean	75^{th} $\%^{ile}$	25^{th} $\%^{ile}$	median	mean	75^{th} $\%^{ile}$
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Hub-and-spoke					0.099	0.083	0.074	0.068
					(0.566)	(0.354)	(0.149)	(0.247)
Core-periphery	−0.341	−0.046	−0.074	0.058				
	(0.536)	(0.268)	(0.149)	(0.196)				
Isolates	0.036	−0.327	-0.388^{***}	-0.502^{***}	0.133	−0.261	-0.314^{**}	-0.466^{**}
	(0.165)	(0.222)	(0.126)	(0.187)	(0.555)	(0.163)	(0.129)	(0.190)
Locally-clustered	-0.602^{**}	−0.399	-0.384^{**}	-0.323^{*}	−0.513	-0.329^{*}	-0.311^{*}	−0.291
	(0.295)	(0.248)	(0.162)	(0.193)	(0.607)	(0.197)	(0.165)	(0.197)
Complete clique	-0.947^{***}	−0.863	-0.285^{*}	0.024	−0.854	-0.898^{*}	−0.212	0.070
	(0.266)	(0.658)	(0.157)	(0.186)	(0.593)	(0.520)	(0.160)	(0.189)
Out core-periphery	-0.524^{*}	−0.468	-0.405^{**}	-0.467^{**}	−0.414	−0.388	-0.332^{**}	-0.429^{**}
	(0.277)	(0.519)	(0.163)	(0.194)	(0.598)	(0.500)	(0.166)	(0.197)
In core-periphery	0.110	−0.095	−0.161	−0.257	0.177	−0.020	−0.088	−0.214
	(0.231)	(0.242)	(0.152)	(0.192)	(0.581)	(0.189)	(0.155)	(0.195)
Constant	0.891^{***}	1.468^{***}	1.548^{***}	1.867^{***}	0.795	1.403^{***}	1.474^{***}	1.831^{***}
	(0.140)	(0.218)	(0.103)	(0.178)	(0.548)	(0.157)	(0.107)	(0.182)
Observations	270	270	270	270	270	270	270	270

Note: *p<0.1; **p<0.05; ***p<0.01
All coefficients are on the log scale

The hub-and-spoke network dominated the three controls: it had fewer very poor trials (those with the majority choosing an incorrect answer) than the complete clique and locally-clustered networks, and had more very good trials (those with the majority choosing the correct answer) than the isolates. Additionally, the average number of people correct in the hub-and-spoke network

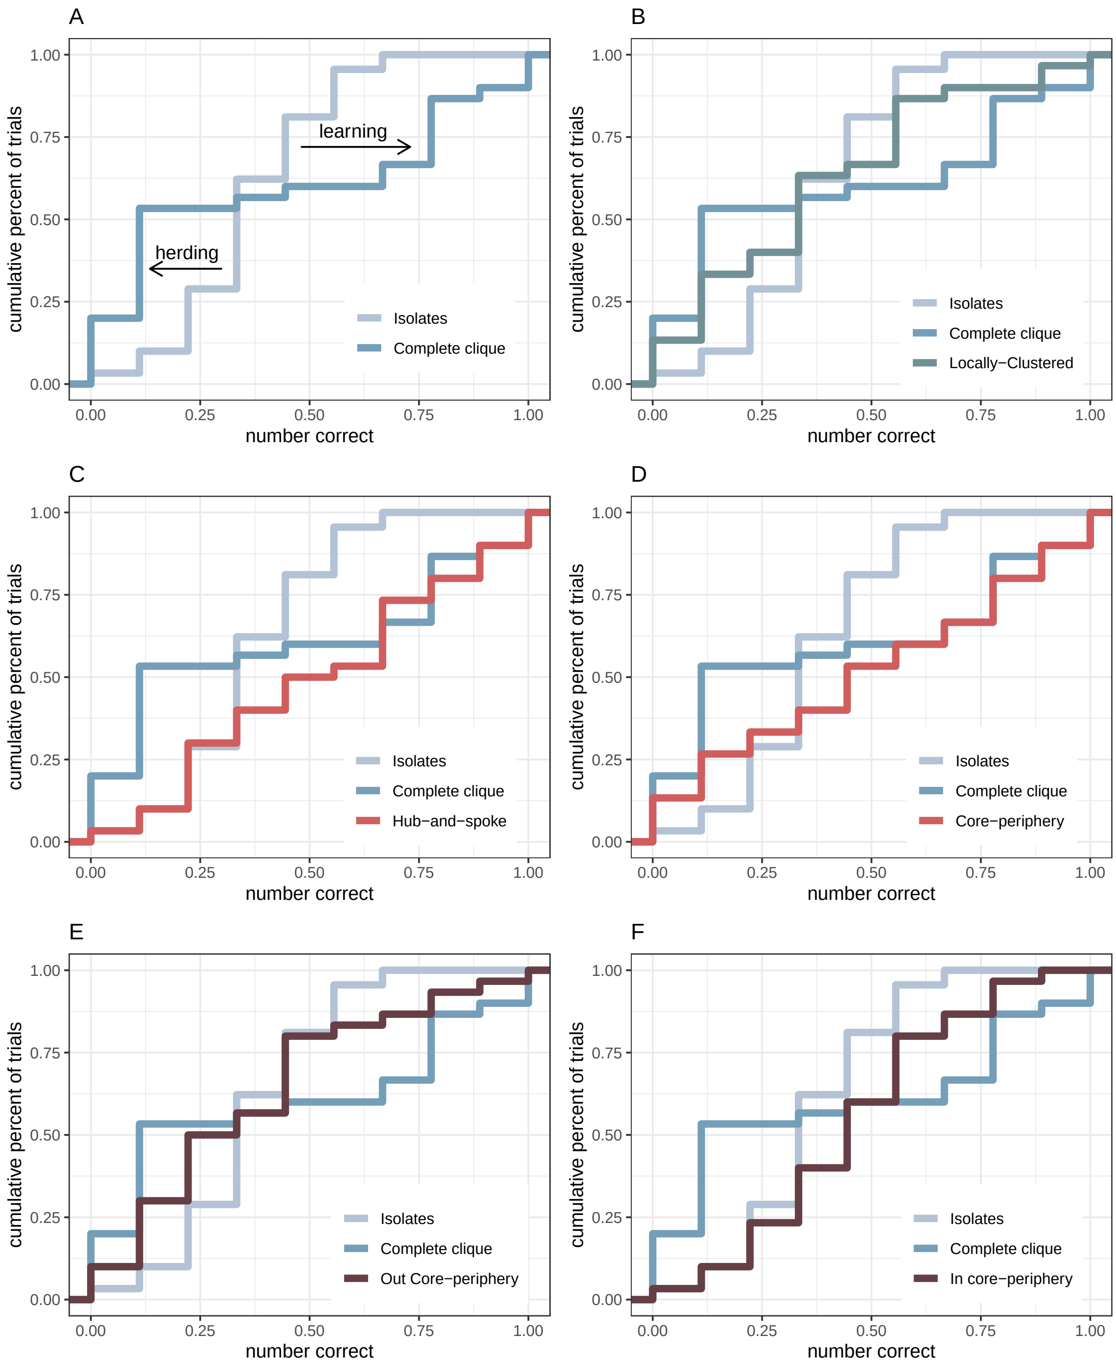

Figure 5: Empirical cumulative distribution functions of the number of individuals correct per trial. Panel A shows only the isolates (no social influence) and complete clique (maximal social influence) conditions. The complete clique condition has more trials with few people correct ("herding" on the wrong answer) and more people with many people correct ("learning" the right answer from peers). Panel C shows that the hub-and-spoke network caused the learning effect without the herding effect.

was higher at 4.7 as compared to 3.19 for the isolates, 3.53 for the complete clique, and 3.2 for the locally-clustered network.

Results for the core-periphery network were more equivocal. Panel D shows a similar effect to that of the hub-and-spoke network, but with some herding occurring due to the presence of a central cluster and fewer peripheral individuals than the hub-and-spoke. The core-periphery network thus outperformed the isolates and locally-clustered network with respect to the learning effect (i.e. 75^{th} percentile) and the mean, but it was only directionally consistent with the hub-and-spoke network and not statistically significantly better with respect to the herding effect.

With respect to the hub-and-spoke network, our experimental data are sufficient to reject the null for Hypotheses 1 (reduced herding, as measured by 25th quantile regression), 2 (maintained learning, as measured by the 75th percentile regression) and 3 (overall positive average effect). For the less centralized core-periphery network, we reject the null for Hypotheses 2 and 3. Overall, given that the effect sizes are smaller but directionally consistent for the less centralized network, we interpret the full body of data as supporting all three main hypotheses.

4.2 Mechanisms

The prior section reported causal effects at a whole-network level and showed that centralized networks are able to adapt to new information and spread the correct solution when it is found. By inspection of the eCDFs of the number of people correct at the end of the trial, we argue that centralized networks allow learning with less herding (in the case of core-periphery networks) or no herding (in the case of hub-and-spoke networks).

In Section 2.4, we argued that the the features of centralized networks that were found to be detrimental to studies of collaboration or the wisdom of the crowds phenomenon might be beneficial when exploration and adaptability are important. Specifically, we posited that (1) the inability for peripheral nodes to communicate directly with each other could reduce conformity pressure and allow those individuals to adapt to new information and (2) that the privileged position of the central node or nodes could allow them to spread good ideas that arose in the periphery.

In this subsection, we present step-by-step evidence of those mechanisms. Figure 6 identifies the sub-mechanisms that are implicit in that argument in the boxes labeled A, B, and C, that is, a more detailed account of what we should expect to observe in our data if the hypothesized

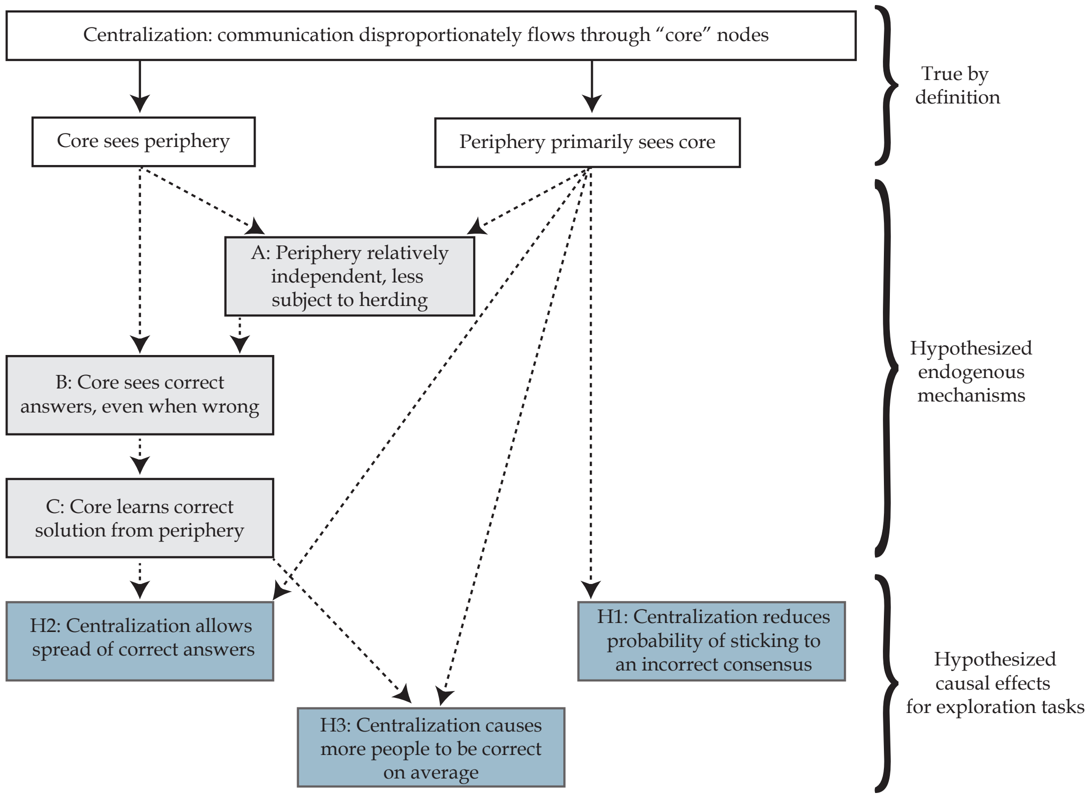

Figure 6: Diagram of mechanisms underlying performance on exploration/exploitation tasks

pathways are correct. We begin by providing evidence for mechanism A: peripheral individuals in centralized networks are less subject to herding than other network positions.

If this is true, then it follows directly that central individuals would have diverse solutions among their network neighbors in the periphery. If the environment changes to reveal information that could lead to the correct solution, then we should expect that the more diverse solutions adopted by peripheral nodes would include correct answers, even when their neighbor(s) in the network core were not correct. Thus we will provide evidence for mechanism B: central/core individuals in centralized networks see more correct answers among their neighbors than other network positions, even when they are wrong themselves.

It further follows that if mechanism B is correct—and the hypothesized effect is sufficiently large to make a practical difference—then this would reduce the conformity pressure for central nodes to adhere to an incorrect solution. Indeed, it would provide them with opportunities to learn from peripheral individuals who had already adapted to the new information and found the correct answer to the problem. As the correct answer is uncertain, complex contagions theory predicts that central individuals would adopt the correct answer after more than one peripheral individual did. Thus, we will provide evidence for mechanism C: central/core individuals in centralized networks are more likely to be correct after more than one person in the periphery is correct.

The following subsections provide evidence for each mechanism. In contrast to the previous section, which provided causal results, this section provides descriptive results of mechanisms unfolding within trials. Results from sections 4.2.1 and 4.2.2 are calculated from data collected every thirty seconds, consisting of which solution each person had registered (if any) and the solutions registered by any visible network neighbors; results from section 4.2.3 are calculated from timestamped records of when each individual registered each answer.

4.2.1 A: Peripheral nodes are relatively independent

Figure 7 plots the probability of having the same answer as one's neighbors for both correct answers (top) and the red herring (bottom). Peripheral positions covary less with their neighbors than other positions. Specifically the left side of each panel shows that they are more likely than other positions to adopt an answer (whether correct or otherwise), even if they do not have a neighbor that has adopted that answer. Similarly, they are less likely to adopt an answer that all of their neighbors have adopted than other positions.

Inspection of Figure 7 reveals that both the number and share of neighbors with a given solution matters in deriving the probability of adopting a matching solution. For example, the periphery of the hub-and-spoke network has only a single neighbor (the central node); although it is a single person, it is 100% of what each peripheral person can see. Thus, that single solution has a higher weight than one of several solutions visible by a position with multiple neighbors, because it is a larger share of the total. Conversely, when the same share of neighbors adopts a solution, the conformity pressure increases with the number of neighbors: when all of one's eight neighbors agree, the probability of matching them is 85-100%, whereas the probability of matching one's single neighbor in positions with only one connection is only approximately 55%.

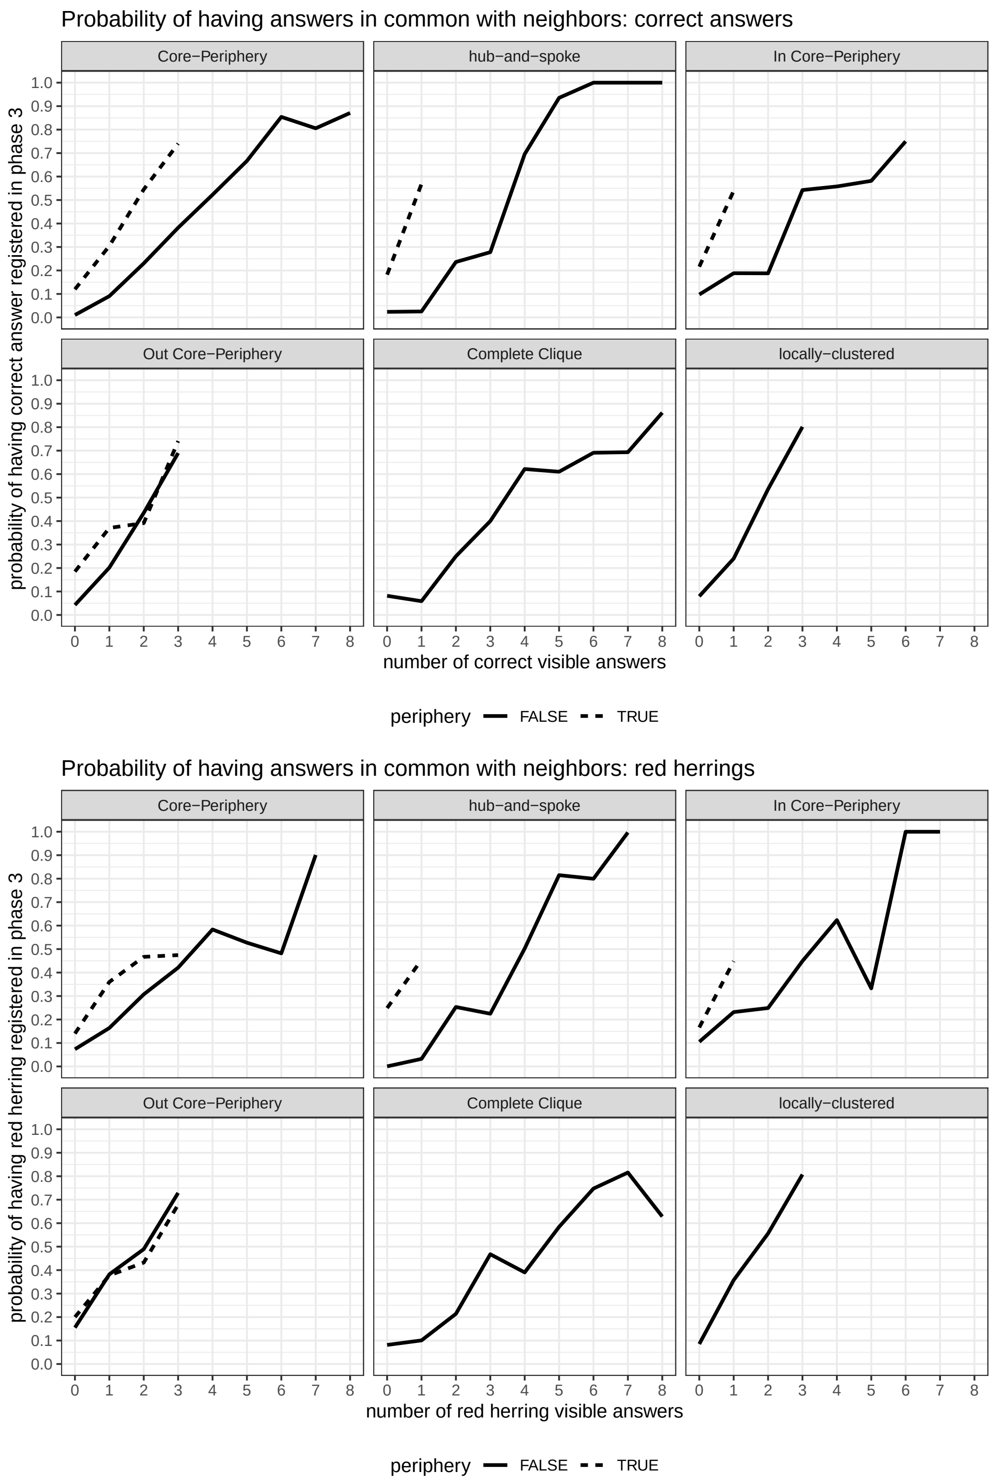

Figure 7: Probability of having the same answer as neighbors by position during phase 3 of the experiment.

4.2.2 B: Core sees more correct answers, even when wrong

The relative independence of peripheral nodes documented in Section 4.2.1 implies that peripheral nodes are less subject to deleterious effects of social influence and are more likely to consider different points of view. In the context of a changing information environment, this also means that they are more likely to consider new information, rather than disregard it in favor of a current (incorrect) consensus.

The top panel in Figure 8 visualizes the moving average of the number of unique solutions visible by individuals in each network position over time. While the general trend is similar across positions, the core nodes in the hub-and-spoke and core-periphery saw more unique solutions on average (3.12 and 2.65, respectively) than did individuals in the complete clique, locally-clustered, or isolate nodes positions (2.44, 1.58, and 0, respectively). Isolates are not shown in the figure because they saw no peer solutions by construction. Section 4.3, below, discusses results for the directed core-periphery networks.

Most importantly for the opportunity to learn from peers, the proportion of alters with correct answers when one has the wrong answer oneself varies significantly by position. In cases where people had the wrong answer registered in the final phase, individuals in the core position of the hub-and-spoke network saw 0.267, and people in the core of the core-periphery network saw 0.203 more correct answers among their neighbors than people in the complete clique ($p = 0.008$ and 0.011, respectively).

Comparing core nodes of centralized networks to nodes in the locally-clustered network, the core node in the hub-and-spoke network saw 1.074 more correct alters, and core nodes in the core-periphery network saw 1.009 more correct alters.

4.2.3 C: core learns from periphery

Section 4.2.2 demonstrates that core nodes in centralized networks saw more correct answers among their more independent alters in their peripheries. Figure 7 demonstrates that those who saw more correct answers were more likely to be correct themselves.

What remains to be shown to connect mechanism C to the main result is that the effects happened in the sequence described in the mechanism diagram: identifying and then spreading the

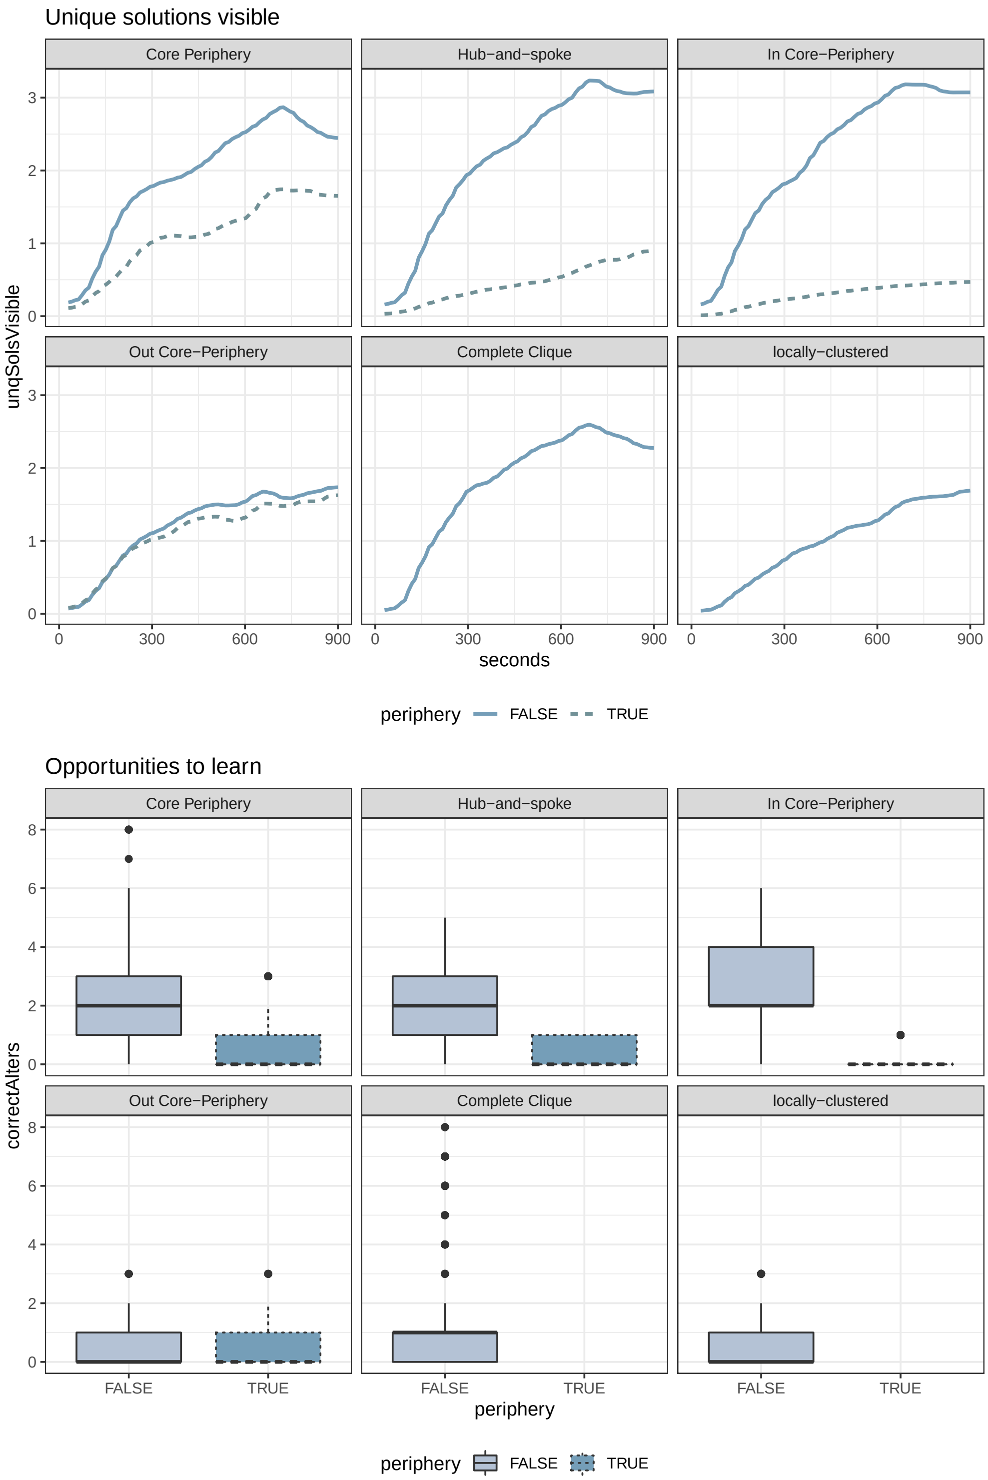

Figure 8: Core nodes see diverse answers, including correct answers, even when they are wrong themselves. Top: moving average of the number of unique solutions visible to each position over the course of the experiment. Bottom: boxplot of the number of correct answers visible to individuals who did not have the correct answer registered during the last minute of the experiment.

correct answer (i.e. that individuals in the core adopted after some in the periphery and before others in the periphery), and that more people were correct in those trials in which the core adopted the correct answer.

To assess the claim that core nodes are likely to adopt the right answer after some in the periphery and before others, we conducted a random permutation test of the order in which individuals in a given trial adopted a correct answer that they kept until the end of the trial. We randomly permuted the sequence of adopting the correct answer—coding anybody who did not end the trial with the correct answer as a 0 and everyone else with the whole number indicating when in the sequence they adopted the correct answer—in groups of thirty trials and repeated this process 10,000 times to create a null distribution against which to compare the real data. We found that the central node in the hub-and-spoke network was statistically significantly ($p = 2 \times 10^{-4}$) more likely to be the third person in the trial to adopt the correct answer, and members of the core in the core-periphery network were significantly ($p = 0.01$) more likely to be the fourth person in the trial to adopt the correct answer.

In the hub-and-spoke network, in cases in which the central node has already adopted the correct answer that they maintain through the end of the trial, 64% of peripheral nodes that have not yet adopted the correct answer go on to do so, as compared to 29% for cases in which the central node has not yet adopted the correct answer (the difference is statistically significant at $p = 10^{-6}$). In Core-periphery networks, the probability that a peripheral node is correct at the end of the trial goes up sharply if two or three core nodes have already adopted the correct answer. The probability of completing the trial with the correct answer is 13% if no core nodes have already adopted the correct answer, but 36% if two core nodes have, and 57% if three core nodes have already adopted the correct answer (p = 0.022 and 2.92e-04, respectively). Figure 9 illustrates both of these findings for the hub-and-spoke network.

4.3 Mechanism Check

The individual-level data discussed in reporting sub-mechanisms A, B, and C in Figure 6 are endogenously correlated within each trial. Thus, although we have detailed descriptive data, they do not support making causal inferences. However, data from the directed core-periphery network treatments can be studied to increase the plausibility of those causal mechanisms, because they

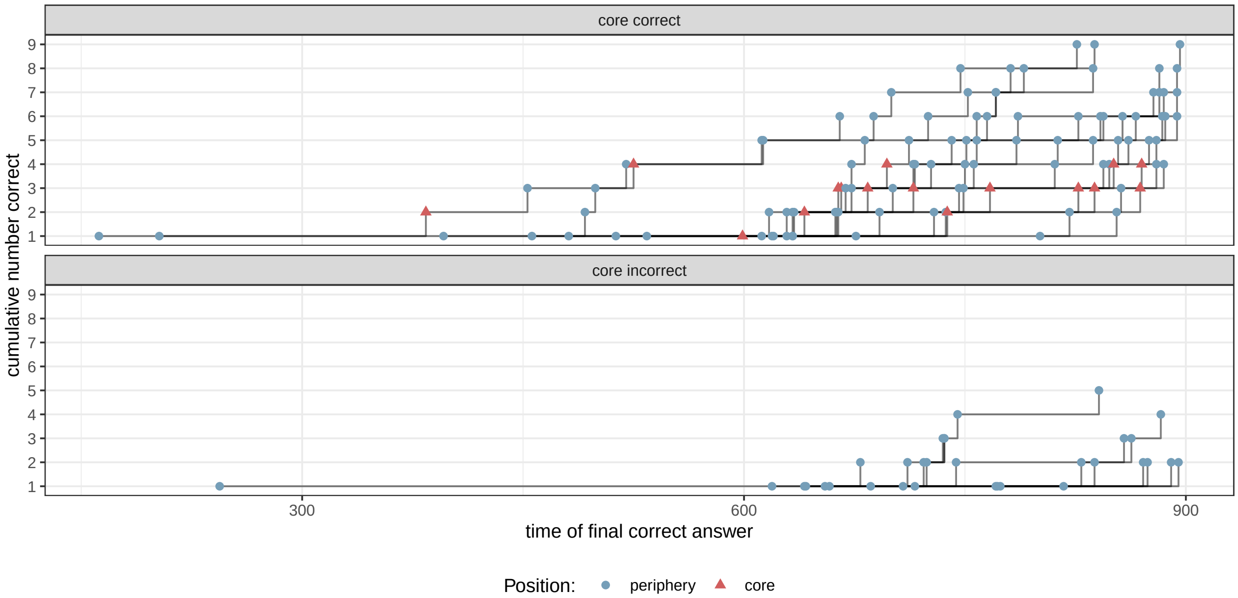

Figure 9: Illustration of time course of final correct answers for hub-and-spoke network trials. If the central node is correct, it typically adopts the correct answer after it has already been found in the periphery, and spreads the correct answer widely. Top: trials in which the central node adopted the correct answer. Bottom: trials in which the central node did not finish with the correct answer.

exogenously manipulate key variables in the causal diagram.

The out-only core-periphery does not have the feature that the core can see the periphery; thus all statements logically downstream from "core sees periphery" should not be true. The core cannot learn from the periphery because it cannot see it. Thus, individuals in the core see less diverse answers, including fewer correct answers than either the hub-and-spoke network or the (undirected) core-periphery network (see Figure 8). Individuals in the core of the out-only core-periphery are thus less likely to be correct than individuals in the core positions of uncentralized networks (estimated on those individuals who completed the trial): they are correct 32% of the time, compared to 57% for individuals in the core of the hub-and-spoke network and 51% for individuals in the core of the undirected core-periphery network ($p-$values from χ^2 tests are 0.029 and 0.014, respectively). With fewer trials with correct members of the core, there are fewer trials in which the core spreads the correct answer to the periphery (see Figure 5 and results from Table 2, indicating little social learning).

The in-only directed core-periphery network had a different pattern of results. In this network, the core could see the periphery, but the periphery could not see the core. Thus, the mechanism-diagram predicts that the core would see many correct answers among the independent peripheral nodes (indeed, due to the fact that peripheral individuals in the in-only core-periphery had fewer or no network neighbors, they were even more independent than peripheral individuals in other networks; see Figure 8, bottom), and thus that individuals in the core would be correct frequently, like those in the undirected centralized networks (which they were, 49% of the time). Despite these structural benefits, the correct answers in the core could not spread widely, since three peripheral individuals had no incoming information from the core and three had only single links from a core member. Results in Figure 5 and Table 2 indicate that while results were not significantly different from the uncentralized network structures, there were fewer trials with many correct individuals, and the 75^{th} percentile of the number of people correct per trial was lower, at 5, compared to 6.47, for the hub-and-spoke network and 6.24, for the undirected core-periphery network.

5 Discussion and Conclusion

Prior research on centralization had found that centralization is harmful for collective intelligence and collaboration because it interferes with the integration of information and knowledge held by peripheral individuals. We show experimentally that the same effects of centralization that limit effective integration of current ideas have benefits for problems that require shifting from one idea to another in response to new information. Thus, we argue that centralization aids in complex problems that benefit from both exploration and exploitation.

In centralized networks, peripheral nodes are relatively independent, which makes them more adaptable to new information and less likely to retain a consensus (but wrong) answer. Core individuals can get information from many of these independent peripheral individuals, and this gives them the opportunity to learn from the correct answers they can see among the periphery. Once they adopt the correct answer, it spreads quickly to many of the remaining members of the periphery. In this way, centralization promotes social learning without deleterious herding.

For problems requiring exploration and adaptability, Lazer and Friedman (2007) theorized that network structure implies "tradeoff between maintaining the diversity necessary for obtaining high performance in a system and the rapid dissemination of high-performing strategies." Our results show that centralization breaks that tradeoff. Like intermittent interaction (Bernstein et al. 2018), centralization creates the benefits of social influence without the associated costs. Indeed, even with network ties that are constant, centralization might even create a de facto intermittent pattern of interaction among peripheral individuals, because they are only exposed to each other's ideas when members of the core pass them on.

Our findings share some aspects in common with prior research on coordination. Prior research shows that centralization helps ensure that everybody shares the same information in (simpler) coordination tasks: having a central individual ensures that one person has access to all the information and can 'put the pieces together' and then ensure that everyone else has the same, integrated information (Bavelas 1950). Our results on complex problems that are harmed by groupthink or herding rely on centralization keeping the beliefs of peripheral indivuals more separate and hindering quick integration. For sufficiently simple tasks, centralization aids in exploitation of existing information, but as tasks get more complex, more bandwidth is necessary to effectively integrate the beliefs and information held by peripheral individuals. It is in this regime that we show that the features of centralization that harm exploitation aid in exploration.

Our experimental findings run exactly opposite to the popular narrative of centralization suppressing diversity. It is possible that the everyday experience of little diversity of thought in centralized structures is misattributed to centralization, but actually a result of other forces. Another possibility is that individual experience of diversity is lessened in centralized networks precisely because centralization isolates peripheral individuals from each other, thus maintaining beneficial aggregate diversity that is the source of adaptability. And of course, it is also possible that this popular negative view of centralization is not based in fact.

5.1 Hierarchy versus Centralization

Empirically, centralization is confounded by hierarchy (Lin 1999), which is the differentiation of individuals by power or status. There is a general tendency toward hierarchical differentiation in organizations (Tiedens and Fragale 2003, Gruenfeld and Tiedens 2010), despite the fact that hierarchy has been linked with the same generally negative effects as centralization (both are generally seen to be good for coordination (Galinsky et al. 2012) but bad for collaborative performance (Van der Vegt et al. 2010, Greer et al. 2018)). Because of their generally confounded nature, hierarchy and social network centralization are unlikely to be distinguishable via observational data. In the present experiment, we manipulate centralization without hierarchical differences among participants to focus on influence of communication network structure per se, rather than hierarchical structure.

5.2 Boundary Conditions and limitations

5.2.1 The central cluster in the core-periphery network

The benefits to centralization that we found were greatest for the maximally centralized hub-and-spoke network, and somewhat less for the core-periphery network. Clustering in social networks increases the spread of uncertain behaviors and thus the conformity pressure for those who can see the behaviors of those in clustered positions (Centola 2010). The central cluster of the core-periphery network thus increases the conformity pressure on peripheral individuals and increases the amount of herding on the wrong answer (see Figure 5, panel D, left side).

This is a major caveat for our findings. While the hub-and-spoke network is a typical structure for formal hierarchies, the core-periphery structure is also empirically common, especially for self-organized or informal networks. Thus, centralization 'in the wild' may not enjoy all of the benefits that we see in the maximally-centralized hub-and-spoke network.

5.2.2 Directed ties in centralized networks

We tested two network structures intended to isolate the mechanisms behind the benefits to centralization: the Out core-periphery and the In core-periphery. In the Out core-periphery, peripheral individuals are aware of what the core thinks, and can receive information from the core, but the core does not receive (as much) information from those peripheral individuals in return. This network structure also occurs empirically, and a striking example can be seen in the pattern of 'following' on social media platforms that allow directed ties, such as Twitter (Wu et al. 2011).

Thus, our research would not necessarily suggest that the form of centralization seen in online communities would engender the same benefits as we see in the undirected centralized network treatments in this experiment. The ability to follow someone unilaterally means that well-known individuals don't have the cognitive bandwidth to pay attention to all of the people who pay attention to them. The key question implied by our research is whether they follow enough independent people to put them in the position to act as an effective filter and spreader of good ideas, such as we see here. At least in the political sphere, prior research is mixed, showing that influential accounts on Twitter do follow an especially diverse group of other accounts, but that they tend to have the most polarized behavior (Shore et al. 2018). Similar to our results, prior research shows that social movements (analogous to a shift from a current consensus to a hopefully better one) tend to begin in the network periphery and spread from there (Barberá et al. 2015).

5.2.3 Random central individuals

As a randomized experiment, it was an intentional feature of the design that the people who were located in central positions were not "special" but instead randomly drawn from the participants in each trial. Thus, our design separates the phenomenon of interest—network centralization—from confounds like hierarchy and power. However, this departure from realism for the sake of valid causal inference also suggests a caveat in interpreting our findings. Centralized networks performed best with random individuals in the key central positions in the network. In practice central individuals are not randomly assigned to, but rather gained, that position, perhaps through greater experience or influence. Thus it is not clear whether central individuals would play the role of identifying and spreading good ideas by others in the periphery, or if the necessities of their social position would sometimes induce them to value their own existing ideas more heavily (or indeed if peripheral nodes would not volunteer their ideas out of deference to the more powerful central node).

Put another way, our results point to the necessity of humility among central individuals: collective performance can be greater if central individuals can identify and spread good ideas found by others in the periphery.

5.2.4 Organization scale and independence of peripheral individuals

In real organizations, centralization can be a multi-level phenomenon. Our experiment tested centralization at the inter-individual level, but centralization can also be an inter-team or inter-unit phenomenon. It is not clear from this experiment how centralization would operate at those higher levels, especially if peripheral units were not separate individuals but rather teams of individuals, for example. It may be that clustering among team-members is sufficient to suppress the phenomenon that peripheral individuals are relatively independent from conformity pressure (submechanism A in this paper). We leave this as an open question that would be very challenging to study with a randomized controlled experiment.

5.3 Conclusion

Centralization has previously been found to interfere with the integration of knowledge and ideas held by individuals in the network periphery. By designing an experimental protocol that manipulates when each piece of information can and cannot be found, we test the effects of centralization in a different class of problems. We show experimentally that centralization helps the collective adaptability to new information and thus improves performance in dynamic problem-solving settings. The caveat is that the central people in the network must be able to learn from and influence the periphery, identifying and spreading the good ideas that are discovered there.

References

Afuah A, Tucci CL (2012) Crowdsourcing as a solution to distant search. *Academy of Management Review* 37(3):355–375.

Argote L, Aven BL, Kush J (2018) The effects of communication networks and turnover on transactive memory and group performance. *Organization Science* 29(2):191–206.

Bahrami B, Olsen K, Latham PE, Roepstorff A, Rees G, Frith CD (2010) Optimally interacting minds. *Science* 329(5995):1081–1085.

Barberá P, Wang N, Bonneau R, Jost JT, Nagler J, Tucker J, González-Bailón S (2015) The critical periphery in the growth of social protests. *PloS one* 10(11):e0143611.

Barton M, Sutcliffe K (2018) Contextualized engagement as resilience-in-action: A study of expedition racing. *Presented at the Academy of Management Annual Meeting.*

Bavelas A (1950) Communication patterns in task-oriented groups. *The Journal of the Acoustical Society of America* 22(6):725–730.

Becker J, Brackbill D, Centola D (2017) Network dynamics of social influence in the wisdom of crowds. *Proceedings of the national academy of sciences* 201615978.

Bernstein E, Shore J, Lazer D (2018) How intermittent breaks in interaction improve collective intelligence. *Proceedings of the National Academy of Sciences* 115(35):8734–8739.

Bigham J (2014) My mturk (half) workday. URL `http://www.cs.cmu.edu/~jbigham/posts/2014/half-workday-as-turker.html`.

Borgatti SP, Everett MG (2000) Models of core/periphery structures. *Social networks* 21(4):375–395.

Boudreau KJ, Lacetera N, Lakhani KR (2011) Incentives and problem uncertainty in innovation contests: An empirical analysis. *Management science* 57(5):843–863.

Butts CT, Petrescu-Prahova M, Remy Cross B (2007) Responder communication networks in the world trade center disaster: Implications for modeling of communication within emergency settings. *Mathematical Sociology* 31(2):121–147.

Cattani G, Ferriani S (2008) A core/periphery perspective on individual creative performance: Social networks and cinematic achievements in the hollywood film industry. *Organization science* 19(6):824–844.

Centola D (2010) The spread of behavior in an online social network experiment. *science* 329(5996):1194–1197.

Christensen CM (2013) *The innovator's dilemma: when new technologies cause great firms to fail* (Harvard Business Review Press).

Colizza V, Flammini A, Serrano MA, Vespignani A (2006) Detecting rich-club ordering in complex networks. *Nature physics* 2(2):110.

Connelly BL, Johnson JL, Tihanyi L, Ellstrand AE (2011) More than adopters: Competing influences in the interlocking directorate. *Organization Science* 22(3):688–703.

Cummings JN, Cross R (2003) Structural properties of work groups and their consequences for performance. *Social networks* 25(3):197–210.

Dahlander L, Frederiksen L (2012) The core and cosmopolitans: A relational view of innovation in user communities. *Organization science* 23(4):988–1007.

DeGroot MH (1974) Reaching a consensus. *Journal of the American Statistical Association* 69(345):118–121.

Engel D, Malone TW (2018) Integrated information as a metric for group interaction. *PloS one* 13(10):e0205335.

Engel D, Woolley AW, Jing LX, Chabris CF, Malone TW (2014) Reading the mind in the eyes or reading between the lines? theory of mind predicts collective intelligence equally well online and face-to-face. *PloS one* 9(12):e115212.

Englmaier F, Grimm S, Schindler D, Schudy S (2018) The effect of incentives in non-routine analytical teams tasks-evidence from a field experiment .

Frahm KM, Shepelyansky DL (2012) Google matrix of twitter. *The European Physical Journal B* 85(10):355.

Galinsky AD, Chou EY, Halevy N, Van Kleef GA (2012) The far-reaching effects of power: At the individual, dyadic, and group levels. *Looking back, moving forward: A review of group and team-based research*, 81–113 (Emerald Group Publishing Limited).

Galton F (1907) Vox populi. *Nature* 75:450–451.

Geraci M, Bottai M (2014) Linear quantile mixed models. *Statistics and Computing* 24(3):461–479, URL `http://dx.doi.org/10.1007/s11222-013-9381-9`.

Golub B, Jackson MO (2010) Naive learning in social networks and the wisdom of crowds. *American Economic Journal: Microeconomics* 2(1):112–49.

Grant RM (1996) Toward a knowledge-based theory of the firm. *Strategic management journal* 17(S2):109–122.

Gray ML, Suri S, Ali SS, Kulkarni D (2016) The crowd is a collaborative network. *Proceedings of the 19th ACM conference on computer-supported cooperative work & social computing*, 134–147 (ACM).

Greer LL, de Jong BA, Schouten ME, Dannals JE (2018) Why and when hierarchy impacts team effectiveness: A meta-analytic integration. *Journal of Applied Psychology* 103(6):591.

Gruenfeld DH, Tiedens LZ (2010) Organizational preferences and their consequences. *Handbook of Social Psychology* .

Gulati R, Puranam P (2009) Renewal through reorganization: The value of inconsistencies between formal and informal organization. *Organization science* 20(2):422–440.

Hansen MT (1999) The search-transfer problem: The role of weak ties in sharing knowledge across organization subunits. *Administrative science quarterly* 44(1):82–111.

Hinds P, McGrath C (2006) Structures that work: social structure, work structure and coordination ease in geographically distributed teams. *Proceedings of the 2006 20th anniversary conference on Computer supported cooperative work*, 343–352 (ACM).

Huang S, Cummings JN (2011) When critical knowledge is most critical: Centralization in knowledge-intensive teams. *Small Group Research* 42(6):669–699.

Kane GC, Alavi M (2007) Information technology and organizational learning: An investigation of exploration and exploitation processes. *Organization Science* 18(5):796–812.

Lavie D, Stettner U, Tushman ML (2010) Exploration and exploitation within and across organizations. *The Academy of Management Annals* 4(1):109–155.

Lazer D, Friedman A (2007) The network structure of exploration and exploitation. *Administrative Science Quarterly* 52(4):667–694.

Leonardi PM (2014) Social media, knowledge sharing, and innovation: Toward a theory of communication visibility. *Information systems research* 25(4):796–816.

Levine SS, Apfelbaum EP, Bernard M, Bartelt VL, Zajac EJ, Stark D (2014) Ethnic diversity deflates price bubbles. *Proceedings of the National Academy of Sciences* 111(52):18524–18529.

Levine SS, Reypens C, Riedl C (2018) Should i stay or should i go? the cognition of exploration and exploitation. *Academy of Management Proceedings*, volume 2018, 10248 (Academy of Management Briarcliff Manor, NY 10510).

Lewis K, Lange D, Gillis L (2005) Transactive memory systems, learning, and learning transfer. *Organization Science* 16(6):581–598.

Liang DW, Moreland R, Argote L (1995) Group versus individual training and group performance: The mediating role of transactive memory. *Personality and social psychology bulletin* 21(4):384–393.

Lin N (1999) Social networks and status attainment. *Annual review of sociology* 25(1):467–487.

Lorenz J, Rauhut H, Schweitzer F, Helbing D (2011) How social influence can undermine the wisdom of crowd effect. *Proceedings of the national academy of sciences* 108(22):9020–9025.

Machado JAF, Silva JS (2005) Quantiles for counts. *Journal of the American Statistical Association* 100(472):1226–1237.

Madirolas G, De Polavieja G (2014) Wisdom of the confident: Using social interactions to eliminate the bias in wisdom of the crowds. *Proceedings of the Collective Intelligence Conference*, 10–12.

Majchrzak A, Cherbakov L, Ives B (2009) Harnessing the power of the crowds with corporate social networking tools: How ibm does it. *MIS Quarterly Executive* 8(2):7.

March JG (1991) Exploration and exploitation in organizational learning. *Organization science* 2(1):71–87.

Mason W, Suri S (2012) Conducting behavioral research on amazons mechanical turk. *Behavior research methods* 44(1):1–23.

Mason W, Watts DJ (2012) Collaborative learning in networks. *Proceedings of the National Academy of Sciences* 109(3):764–769.

Montgomery JM, Nyhan B, Torres M (2018) How conditioning on posttreatment variables can ruin your experiment and what to do about it. *American Journal of Political Science* 62(3):760–775.

Muchnik L, Aral S, Taylor SJ (2013) Social influence bias: A randomized experiment. *Science* 341(6146):647–651.

Mulder M (1960) Communication structure, decision structure and group performance. *Sociometry* 23(1):1–14.

Nickerson JA, Zenger TR (2004) A knowledge-based theory of the firmthe problem-solving perspective. *Organization science* 15(6):617–632.

Pan W, Altshuler Y, Pentland A (2012) Decoding social influence and the wisdom of the crowd in financial trading network. *2012 International Conference on Privacy, Security, Risk and Trust and 2012 International Confernece on Social Computing*, 203–209 (IEEE).

Puranam P (2018) *The microstructure of organizations* (Oxford University Press).

Salganik MJ, Dodds PS, Watts DJ (2006) Experimental study of inequality and unpredictability in an artificial cultural market. *science* 311(5762):854–856.

Shore J, Baek J, Dellarocas C (2018) Network structure and patterns of information diversity on twitter. *MIS Quarterly* 42(3):849–872.

Steiner ID (1972) *Group processes and productivity* (Academic press).

Tchernichovski O, Parra LC, Fimiarz D, Lotem A, Conley D (2019) Crowd wisdom enhanced by costly signaling in a virtual rating system. *Proceedings of the National Academy of Sciences* 116(15):7256–7265.

Tiedens LZ, Fragale AR (2003) Power moves: complementarity in dominant and submissive nonverbal behavior. *Journal of personality and social psychology* 84(3):558.

Tushman ML (1979) Work characteristics and subunit communication structure: A contingency analysis. *Administrative Science Quarterly* 24(1).

van de Rijt A (2019) Self-correcting dynamics in social influence processes. *American Journal of Sociology* 124(5):1468–1495.

Van der Vegt GS, De Jong SB, Bunderson JS, Molleman E (2010) Power asymmetry and learning in teams: The moderating role of performance feedback. *Organization Science* 21(2):347–361.

Woolley AW, Chabris CF, Pentland A, Hashmi N, Malone TW (2010) Evidence for a collective intelligence factor in the performance of human groups. *science* 330(6004):686–688.

Wu S, Hofman JM, Mason WA, Watts DJ (2011) Who says what to whom on twitter. *Proceedings of the 20th international conference on World wide web*, 705–714 (ACM).

Appendices

A Additional details of subject pipeline

A.1 Details of phase 1 trial

Phase 1 served to teach participants how to use the experimental interface and also to ensure they could both perform basic logical deduction and were paying sufficient attention to do so. We provide details on the logical task here.

Participants received the following clue as part of the instructions: "The murderer is either Mr. Smith, Mr. Lee, or Ms. White. The murder weapon was poison." The instructions continue:

> "Other clues must be found by using the search bar to search for keywords. DO IT NOW: Use the search bar to find more information about each suspect. Tip 1: searches should be single words, so type "Smith" instead of "Mr. Smith". Tip 2: you can search the same keyword more than once to find more clues. Tip 3: some clues contain useful information and some don't."

The clues that could be found with the search tool were as follows:

- Mr. Smith was out of town at the time of the murder.
- Mr. Smith was very sad to learn about the murder.
- Ms. White was with Mr. Smith at the time of the murder.
- Ms. White grew up in California.
- Mr. Lee's fingerprints were found at the scene of the crime.
- Mr. Lee had poison at his house that matched the poison used in the murder.

To pass the stage 1 trial, subjects had to correctly identify "Mr. Lee" as the culprit.

A.2 Payment

In designing our payment scheme, we worked to balance the need to incentivize attentive participation to gather valid data against the desire to pay participants a fair wage for their time.

In addition to using the pipeline to filter out inattentive participants, our solution was have a three-part payment structure. We provided a guaranteed but modest ($0.40) base payment; we paid

$0.14 per minute that they had the correct answer to a simple attention check question registered; finally, we paid $0.14 per minute that they had the correct answer to the murder mystery entered. The attention check (which simply asked the name of the murder victim) could be answered based on the first "message from headquarters" clue provided to everyone at the start of the trial, and thus accounted for the majority of subject payment, regardless of whether they eventually solved the mystery. If participants took thirty seconds to enter the answer to the attention check and did not solve the mystery, they would earn approximately $8.58 per hour, assuming that they spent 15 minutes on the trial itself, plus 2 minutes on reading the consent form and waiting for others to be ready to join the trial. In practice, our median hourly wage for participants that completed a trial and provided an answer to the attention check problem was $8.47, assuming a total of 17 minutes spent on the task. That said, participants varied, and the 10th and 90th percentile wages were $6.52, and $10.41.

Following Mason and Suri (2012), we used a waiting room to facilitate synchronous start to multi-person trials. Prior to each trial beginning, we assigned more participants than necessary to a trial, with the understanding that some workers either do not consent to a trial, or have other reasons for not continuing, such as disinterest or inattention. If more than the required number of participants consent to a trial and click the button indicating they are ready to begin, we randomly assigned the required number to begin the actual trial. All participants were paid $0.40 regardless of whether they consented and indicated they wre ready to join (an estimated $8 to $12 per hour equivalent). We paid an additional bonus of $0.40 to those that indicated they were ready to join but were not assigned to an active trial ($16 to $24 per hour equivalent), to encourage them to re-enter the queue and eventually join a full trial.

B Results on full trials only

In the main results, we presented results fitted on the full data, without regard for whether subjects dropped out of any of the trials. Here we present regression results fitted to only those data in which no subjects dropped out of trials. Table 3 shows that results on this subset closely match the full data.

Table 3: Number of people correct at end of trial: trials with no dropout only

	Dependent variable:							
	25^{th} $\%^{ile}$	median	mean	75^{th} $\%^{ile}$	25^{th} $\%^{ile}$	median	mean	75^{th} $\%^{ile}$
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
hub and spoke					−0.019	0.038	−0.000	−0.032
					(0.487)	(0.340)	(0.167)	(0.181)
core-periphery	−0.173	0.105	0.000	0.116				
	(0.558)	(0.344)	(0.167)	(0.161)				
isolates	0.077	−0.163	−0.299**	−0.403***	0.096	−0.201	−0.299**	−0.472***
	(0.187)	(0.318)	(0.146)	(0.142)	(0.471)	(0.156)	(0.146)	(0.125)
locally clustered	−0.602**	−0.295	−0.333*	−0.271*	−0.589	−0.333	−0.333*	−0.340**
	(0.285)	(0.353)	(0.181)	(0.150)	(0.518)	(0.217)	(0.181)	(0.136)
complete clique	−1.040***	−1.120***	−0.379**	−0.018	−1.040*	−1.150***	−0.379**	−0.040
	(0.330)	(0.367)	(0.185)	(0.311)	(0.542)	(0.240)	(0.185)	(0.197)
out core-periphery	−0.495	−0.425	−0.395**	−0.448***	−0.485	−0.433	−0.395**	−0.520***
	(0.318)	(0.506)	(0.180)	(0.145)	(0.537)	(0.459)	(0.180)	(0.130)
in core-periphery	0.311	0.088	−0.041	−0.204	0.323	0.056	−0.041	−0.268*
	(0.192)	(0.330)	(0.171)	(0.145)	(0.473)	(0.179)	(0.171)	(0.144)
Constant	0.879***	1.350***	1.520***	1.830***	0.860*	1.390***	1.520***	1.900***
	(0.158)	(0.312)	(0.118)	(0.126)	(0.460)	(0.144)	(0.118)	(0.108)
Observations	216	216	216	216	216	216	216	216
Note:								*p<0.1; **p<0.05; ***p<0.01

C Evaluation of statistical models

Our experimental data on the number of people correct at the end of the trial do not conform neatly to the assumptions of standard statistical models. Individual successes are dichotomous, but they are nested within trials, within which social influence occurs. Thus, at a trial level, the data could be considered the results of a binomial process (measured as the number of successes and failures out of nine possible people per trial), but with overdispersion such that the success or failure of different individuals in the same trial is correlated. Alternatively, the trial-level data could be considered a poisson-like process, but with the caveat that the data are bounded above with a maximum number of possible successes equal to nine. A further complication is that the amount and pattern of social influence varies systematically (but in a way unknown to the experimenters *a priori*) by treatment condition.

Thus there is no single statistical model for our data that is obviously best, and we turn to a simulation-based approach to evaluate the merits of different approaches.

C.1 False positives

We simulate thirty trials of nine people for each of three synthetic conditions intended to mimic different types of social influence in networks, such as we observe in our experimental data. Each condition has the same global mean number of individuals correct, but the number of correct individuals per trial is drawn from different distributions. The first is akin to an "isolates" condition, in which each individual is independently and identically bernoulli random variable (each trial is a binomial draw). The second is akin to the "complete clique" condition, in which all nine people in a single trial have the same answer, such that the trial as a whole is determined by a single draw of a bernoulli random variable. The third condition is akin to a more moderate level of social influence, and the number of successes per trial is drawn from the beta binomial distribution (i.e. an overdispersed binomial). For each condition, the grand mean was set such that on average three out of nine people were correct (chosen to approximate the average for the isolates condition in the experimental data).

We create 1000 simulated datasets as above and fit several statistical models in order to assess the false positive rate. We considered variations on standard models for dichotomous or count data. Specifically, we fit

1. a standard binomial regression (unit of observation is the trial, data is number of successes and failures per trial)
2. a quasibinomial regression
3. a beta binomial regression
4. a logistic regression with standard errors clustered by trial (unit of observation is the individual)
5. a random-effects logistic regression with separate random variances for the two SI conditions, and
6. a quasipoisson regression (unit of observation is the trial).

Fixed effects models could not be fit on this data due to the perfect separation of the dependent variable in the individual trials in the maximum social-influence condition.

C.1.1 Results

Our primary concerns in choosing a regression framework for analyzing this data were obtaining unbiased estimates of the mean number people correct at the end of the trial and avoiding false positive findings[2]. Therefore, we assessed coefficient bias and the false positive rate on data for which we know there is no true difference in means by construction.

Table 4 gives the coefficients estimated by each model. Most models give an unbiased estimate of the ground truth quantity of number of people correct, but the estimates from random effects logistic regression and the beta binomial regression were extremely biased. Although these two models are intended to model endogenous correlation, they perform poorly and we rejected them for use in analyzing this experiment.

The proportion of datasets for which a false positive finding of statistical significance occured (for different decision thresholds, α) is in Table 5. The standard binomial model does not account for overdispersion from social influence within trials, and unsurprisingly has an unacceptably high false positive rate (almost 20% at $\alpha = 0.05$). The remaining models—logistic regression with clustered standard errors, quasibinomial regression and quasipoisson regression—had more reasonable false positive rates. Among these three, logistic regression with clustered standard errors produced more false positives (approximately 60% more at $\alpha = 0.05$), so we reject it in favor of one of the remaining two.

C.2 False negatives

For completeness, we also created simulated data in which the moderate social influence condition had a higher true mean, similar to the mean of the hub-and-spoke network in the real results. Tables 6 and 7 give mean coefficients and False negative rates when there is a true difference in means. Results on this second simulated data set are similar to the first one, above. Of particular note is that the false negative rate for the quasibinomial and quasipoisson regressions is substantial: around 81% for $\alpha = 0.05$. With low false-positive rates and high false-negative rates, these regression methods are conservative tests of our hypotheses[3].

[2] Note that our intention is to provide a conservative test to build theory. If our purpose was to decide which structure to implement in an organization, we would also need to balance the costs of false negatives.

[3] We cannot say for certain what the probability is that there is a true result conditional on a finding of statistical significance without making strong assumptions about the underlying distribution of effects in the world.

Table 4: Estimated number of people correct under different levels of social influnce (S.I.)

Model	no S.I.	medium S.I.	max S.I.
ground truth	2.986	3.029	3.005
binomial	2.986	3.029	3.005
logit, Clustered se's	2.986	3.029	3.005
logit, random effects	2.986	2.191	0.333
quasibinomial	2.986	3.029	3.005
quasipoisson	2.986	3.029	3.005
beta binomial	3.012	2.753	2.381

Table 5: False positive rates for networks with same true mean

Model	$\alpha = 0.10$	$\alpha = 0.05$	$\alpha = 0.01$
binomial	0.255	0.198	0.118
logit, Clustered se's	0.068	0.036	0.009
logit, random effects	0.424	0.391	0.358
quasibinomial	0.044	0.021	0.004
quasipoisson	0.044	0.022	0.004
beta binomial	0.112	0.076	0.026

C.3 Conclusion

Both quasibinomial and quasipoisson regression provided unbiased estimates and conservative hypothesis tests. In the tables of results, we opted for quasipoisson regression so that the coefficients would be on the same scale as the models of conditional quantiles for count data.

Table 6: Estimated number of people correct under different levels of social influence (S.I.)

Model	no S.I.	medium S.I.	max S.I.
ground truth	3.007	4.742	3.015
binomial	3.007	4.742	3.015
logit, Clustered se's	3.007	4.742	3.015
logit, random effects	3.007	4.895	0.311
quasibinomial	3.007	4.742	3.015
quasipoisson	3.007	4.742	3.015
beta binomial	3.031	4.499	2.372

Table 7: False negative rates for trials with different true means

Model	$\alpha = 0.10$	$\alpha = 0.05$	$\alpha = 0.01$
binomial	0.521	0.558	0.619
logit, Clustered se's	0.675	0.717	0.796
logit, random effects	0.454	0.492	0.564
quasibinomial	0.728	0.808	0.933
quasipoisson	0.731	0.814	0.945
beta binomial	0.735	0.784	0.874