Network Centralization and Collective Intelligence: a Randomized Experiment

Jesse Shore, Boston University
Ethan Bernstein, Harvard Business School
Alice Jang, Boston University

1. SUMMARY AND OVERVIEW

Prior research shows that social network centralization—such as can be seen in online social media—harms the wisdom of the crowd phenomenon. Here, we present experimental evidence that the mechanisms that responsible for harming the wisdom of the crowd (ineffective integration of information currently held in the network) actually make the network more adaptable to new information that arises. Central nodes are exposed to the independent and thus diverse solutions of peripheral nodes, which makes them relatively unlikely to engage in deleterious herding behavior. These benefits of centrality lead central nodes to be more likely to find correct solutions, which are then spread easily to the rest of the network. However, we also present a major caveat: central nodes must not only influence, but also be influenced by sufficiently many relatively independent peripheral nodes to get the benefits of centralization.

2. CENTRALIZATION HARMS THE WISDOM OF THE CROWD

Starting with [Galton 1907] and increasing recently, 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. When people influence each other's judgement, it can sometimes lead to "herding" behavior, correlated errors and reduction of the wisdom of the crowd effect [Lorenz et al. 2011; Becker et al. 2017]. Exogenous network centralization interferes with crowd wisdom by subjecting many peripheral nodes to the estimate of the central node(s), which are not necessarily good.

2.1 Centralization may help networks be adaptable to new information

Prior literature on centralization has primarily been focused on issues of how to make best use of the information, knowledge, or perspectives currently within the network, rather than the ability of network members adapt to new information and thus consider moving from one potential solution to another. In other words, prior work on the wisdom of the crowds phenomenon has focused on problems of maximizing "exploitation" of existing ideas, rather than "exploration" for new ones.

We believe that the most consequential collective intelligence tasks are not as simple as effectively integrating information that is already present in the network, but include navigating rugged solution spaces (e.g. what combination of technologies leads to a useful innovation?) and occur in changing information environments (e.g. how should our policies adapt to unfolding and unpredictable events?).

Previous literature shows that network structures that limit exploitation – like centralization – can support greater exploration [Lazer and Friedman 2007]. Thus, we conduct a randomized experiment to test the effects of centralization on tasks that require exploration or adaptation rather than simple exploitation of existing information.

3. EXPERIMENTAL DESIGN

Using a new online lab platform, we conducted a randomized experiment of the effect of network centralization in a complex problem-solving setting (see Figure 1 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 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.

3.1 Network Treatments

We examine six different network treatments with nine nodes (individual participants) each (see Fig 1). The two centralized networks that are the focus of our study are the hub-and-spoke network and the core-periphery network. 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. Finally, we tested three uncentralized network structures as controls. The locally-clustered network consists of three interconnected cliques of three individuals. The complete clique is a group of nine people, all of whom are connected to each other. 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. These control represent what we believe are empirically plausible alternatives to centralization, with different levels of social influence.

Fig. 1. Visualizations of network treatments

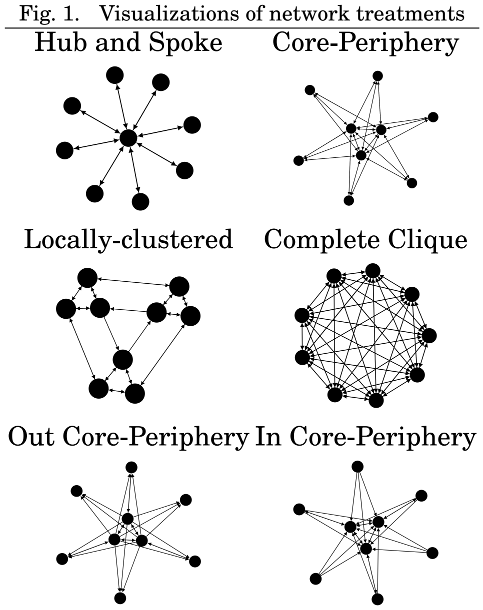

We collected 30 independent trials for each experimental condition, and 90 bootstrap-sampled synthetic trials of isolates. Network treatments were collected in randomly-permuted order to avoid systematic biases due to time effects.

4. RESULTS

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 complex problem-solving 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 of communication without the costs.

Figure 2 illustrates the beneficial effects of centralization in the form of empirical cumulative distribution functions (eCDFs) for three network treatments (space prevents reporting full results for all treatments). 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). Panel B 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).

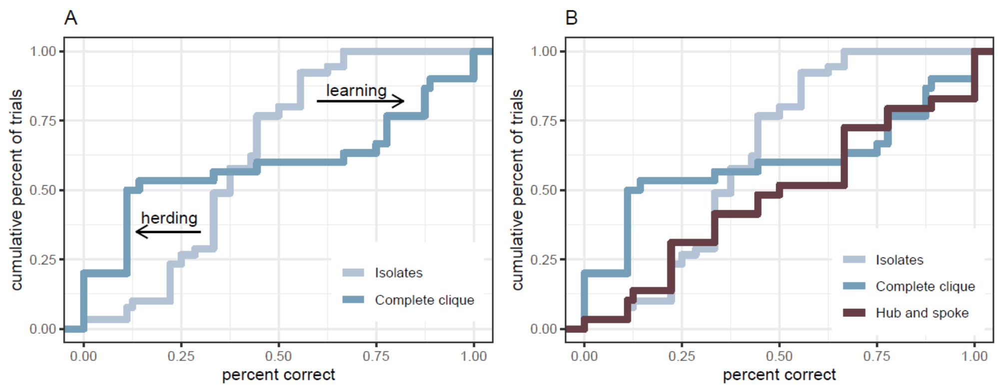

Fig. 2. Empirical cumulative distribution functions of number of individuals correct per trial. Panel A shows only the isolates (no social influence) and complete clique (maximal social influence) conditions. The complete clique condition has more trials with few people correct ("herding" on the wrong answer) and more people with many people correct ("learning" the right answer from peers). Panel B shows that the hub and spoke network caused the learning effect without the herding effect.

The hub and spoke network dominated the three controls: quantile regressions show that 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 was higher at 4.66 as compared to 3.19 for the isolates, 3.53 for the complete clique, and 3.13 for the locally clustered network.

4.1 Mechanisms

The full paper provides evidence for the following mechanisms underlying the success of centralized networks in adapting to new information. (1) Peripheral nodes are relatively independent, making them less likely to engage in deleterious herding than people in other network positions. (2) *In undirected networks*, core nodes can thus see relatively diverse answers, including more instances of the correct answer, even when they themselves do not have the correct answer registered. (3) The core is thus more likely to adopt and further the spread of the correct answer.

REFERENCES

Joshua Becker, Devon Brackbill, and Damon Centola. 2017. Network dynamics of social influence in the wisdom of crowds. *Proceedings of the national academy of sciences* (2017), 201615978.

Francis Galton. 1907. Vox Populi. *Nature* 75 (1907), 450–451.

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

Jan Lorenz, Heiko Rauhut, Frank Schweitzer, and Dirk Helbing. 2011. How social influence can undermine the wisdom of crowd effect. *Proceedings of the national academy of sciences* 108, 22 (2011), 9020–9025.