Rewiring the Wisdom of the Crowd

Jason W. Burton (jburto03@mail.bbk.ac.uk)
Department of Psychological Sciences, Birkbeck, University of London
London, WC1E 7HX, United Kingdom

Abdullah Almaatouq (amaatouq@mit.edu)
Sloan School of Management, Massachusetts Institute of Technology
Cambridge, MA 02142, United States of America

M. Amin Rahimian (rahimian@pitt.edu)
Department of Industrial Engineering, University of Pittsburgh
Pittsburgh, PA 15269, United States of America

Ulrike Hahn (u.hahn@bbk.ac.uk)
Department of Psychological Sciences, Birkbeck, University of London
London, WC1E 7HX, United Kingdom

Abstract

Digitally-enabled means for judgment aggregation have renewed interest in "wisdom of the crowd" effects and kick-started collective intelligence design as an emerging field in the cognitive and computational sciences. A keenly debated question here is whether social influence helps or hinders collective accuracy on estimation tasks, with recently introduced network theories offering a reconciliation of seemingly contradictory past results. Yet, despite a growing body of literature linking social network structure and the accuracy of collective beliefs, strategies for exploiting network structure to harness crowd wisdom are under-explored. In this paper, we introduce a potential new tool for collective intelligence design informed by such network theories: rewiring algorithms. We provide a proof of concept through agent-based modelling and simulation, showing that rewiring algorithms that dynamically manipulate the structure of communicating social networks can increase the accuracy of collective estimations in the absence of knowledge of the ground truth.

Keywords: collective intelligence; wisdom of crowds; social networks; agent-based model

Introduction

Researchers have long demonstrated so-called "wisdom of the crowd" effects, where the collective judgment of a group is more accurate than the judgments of individual experts or the individual group members themselves (Condorcet, 1785; Galton, 1907; Grofman, Owen, & Feld, 1983; Surowiecki, 2005). Yet recently, the impetus for crowd wisdom research has been rejuvenated as new digitally-enabled means for judgment aggregation have given rise to modern applications such as online prediction markets (Wolfers & Zitzewitz, 2004; Arrow et al., 2008), crowdsourcing (Howe, 2006), and digital democracy (Simon, Bass, Boelman, & Mulgan, 2017; Morgan, 2014). Following the successes of these applications, there is now an emerging field in the cognitive and computational sciences dedicated to *collective intelligence design*, whereby digital tools (e.g., algorithms, artificial intelligence, online forums) are developed so as to effectively extract wisdom from ever-present crowds (Mulgan, 2018). In the present work, we draw from extant literature on wisdom of the crowd effects and propose a new instrument for collective intelligence design familiar from network science: rewiring algorithms.

Social Influence, Network Structure, and Collective Estimation

The earliest results on wisdom of the crowd effects in collective estimation tasks assumed that individuals' judgments are made independently, meaning that their errors are uncorrelated and cancel out in aggregate (Condorcet, 1785). However, this independence assumption often goes unmet in the real world because people communicate with or otherwise influence one another. Past research on the effects of social influence in collective estimation tasks has produced seemingly contradictory findings. On one hand, there is evidence that social influence indeed undermines crowd wisdom by causing individuals' judgments to become correlated (Muchnik, Aral, & Taylor, 2013; Lorenz, Rauhut, Schweitzer, & Helbing, 2011; Hahn, von Sydow, & Merdes, 2019); while on the other, there are studies that report an increase in collective accuracy following social influence (Gürçay, Mellers, & Baron, 2015; Becker, Brackbill, & Centola, 2017; Becker, Porter, & Centola, 2019; Almaatouq, Noriega-Campero, et al., 2020).

Formal results that incorporate the possibility of non-independence provide a potential explanation of these seeming contradictions (e.g., Ladha, 1992; Page, 2008). Social influence is neither inherently beneficial nor inherently detrimental to crowd wisdom; instead its effects depend on whether the benefits of communication to individual accuracy outweigh the detrimental effects of non-independence on collective accuracy. The logic of this is made clear in the Diversity Prediction Theorem, which states that collective error squared is the difference between the average individual error squared and the diversity of the individuals' judgments (Page, 2008). While providing a mathematical guarantee that the collective estimate will always be more accurate, in terms of error squared, than the average individual's as long as there exists some diversity in the group, this theorem formalizes how social influence can be both good for collective accuracy (if it leads to an increase in average individual accuracy) and bad (if it leads to too much of a decrease in diversity). Whether social influence will increase or decrease collective accuracy for any given group thus depends on which one of these dueling effects is greater.

To provide predictions for when groups will benefit from social influence, recent research has turned towards studying how different social network structures affect collective accuracy (e.g., Jönsson, Hahn, & Olsson, 2015; Becker et al., 2017; Hahn, Hansen, & Olsson, 2018; Hahn et al., 2019; Almaatouq, Noriega-Campero, et al., 2020). Because social network structures delineate the paths through which social influence can be exerted in a group, it follows that different structural characteristics will feature in determining whether the net effect of social influence will be beneficial for collective accuracy. For example, high levels of connectivity and free-flowing information can lead to "excess correlation" (i.e., correlation between individuals that is not accuracy inducing) (Jönsson et al., 2015); high levels of centralization can lead to certain individuals wielding excessive influence over the network (Becker et al., 2017); and a lack of structural plasticity can prevent networks from effectively responding to feedback about individuals' performance (Almaatouq, Noriega-Campero, et al., 2020).

A reading of the literature linking network structure and collective accuracy begs the question: is there an optimal social network structure for eliciting the wisdom of the crowd? The answer, it seems, is that it depends on the "estimation context": the specific population of individuals and the specific estimation task (Almaatouq, Rahimian, & Alhajri, 2020). Different populations (e.g., computer scientists vs. MTurkers) faced with different tasks (e.g., estimating the number of jelly beans in a jar vs. predicting the time it will take for AGI to be successfully developed) might display different biases and be more or less dispersed in their initial estimates. To operationalize these differences, Almaatouq, Rahimian, and Alhajri (2020) introduce a feature of the estimation context, Omega (Ω), as a function of empirically measured bias (the relationship between a population's initial collective estimate and the ground truth) and dispersion (the shape of the population's initial estimate distribution) for a given task. Through simulation and a re-analysis of data from four prior studies (Lorenz et al., 2011; Gürçay et al., 2015; Becker et al., 2017, 2019), they demonstrate that Ω can predict the probability that a group will benefit from a centralized network structure, thereby emphasizing that optimal network structure is context dependent. For example, consider an estimation task that elicits a fat-tailed distribution of estimates where the extreme errors of a few individuals dominate the collective estimate. In such contexts, which display a high Ω value[1], a completely decentralized network would likely be worse off than a centralized one because the few erroneous individuals would be able to wield as much influence over the network as the accurate majority.

Rewiring Algorithms for Collective Accuracy

Despite the abundance of knowledge on the relationship between network structure and collective accuracy, strategies for exploiting network structure to increase collective accuracy remain under-explored. While there may be considerable difficulties in manipulating the structure of social networks in the analog world, the digital world provides new opportunities. Just as algorithms have already been used to mediate the information presented to online social networks (Lazer, 2015) and to identify influential nodes in social networks (Wei et al., 2018), it seems plausible that algorithms could be used to rewire the structure of online social networks to boost the wisdom of crowds.

In this paper, we use agent-based modelling and simulation to explore the viability of rewiring algorithms—programmable rules for manipulating who communicates with whom—as a tool for collective intelligence design. Specifically, we develop and test two candidate algorithms and evaluate their effects on collective accuracy in a range of estimation contexts: *Can rewiring the structure of communicating social networks produce more accurate collective estimates? How do the effects of the rewiring algorithms vary across estimation contexts? What cues are there for predicting which algorithm might be beneficial?*

Modelling and Simulations

Our simulations utilize an extended DeGroot (1974) model of opinion dynamics in which 16-agent networks estimate (or predict) some unknown positive number. Such a task maps onto classical crowd wisdom scenarios such as estimating the weight of an ox, as well as high-stakes, real-world scenarios like forecasting the number of ICU admissions per week during a pandemic.

To initialize our model, we randomly generate an undirected small-world network (Watts & Strogatz, 1998) and endow each agent with an initial estimate. As a way of both simulating different estimation tasks and forgoing assumptions about belief distributions, each agent's initial estimate is assigned by sampling from a compilation of empirical data from four previously published experiments (Lorenz et al., 2011; Becker et al., 2017; Gürçay et al., 2015; Becker et al., 2019). This compiled dataset spans a total of 57 estimation tasks on which 4,002 individuals provided independent estimates (Almaatouq, Rahimian, & Alhajri, 2020)[2]. Each task—or "estimation context"—in this dataset is represented by a distribution of independent estimates and a true value For example, one task contains 278 participants' estimates of the London population in July 2010, with the true value of 7,825,200 (Gürçay et al., 2015).

[1] Ω values range from 0 to 1, with $\Omega > 0.5$ indicating the initial estimates are better suited to a centralized network structure, and vice versa (Almaatouq, Rahimian, & Alhajri, 2020).

[2] We scale the estimates for each task to be between 0 and 1 in order to suit our belief updating rule and mean-extreme rewiring algorithm, while maintaining the distributions' shape.

Once their initial estimates are assigned, the agents communicate with one another over the course of five time points, $t = 0,1,2,3,4$. At each time point, each agent i revises their estimate in light of those communicated by their network neighbors according to a DeGroot (1974) belief updating rule (Becker et al., 2017):

$$R_{t+1,i} = \alpha_i \times R_{t,i} + (1-\alpha_i) \times \bar{R}_{t,j \in N_i}, \tag{1}$$

where $R_{t+1,i}$ is the agent's revised estimate following communication; $R_{t,i}$ is the agent's current estimate; $\bar{R}_{t,j \in N_i}$ is the average current estimate of the agent's network neighbors; and α_i and its complement $(1-\alpha_i)$ represent the weight that the agent places on its own estimate versus those of its peers, respectively. Following the empirical analysis of belief revision in Becker et al. (2017), each agents' α at any given time point is determined by the following regression equation:

$$\alpha_i = 0.75 - 0.05\varepsilon_i + \mathcal{N}, \tag{2}$$

where ε_i is the agent's absolute error, and $\mathcal{N}$ is Gaussian noise with $\mu = 0$ and $\sigma = 0.06$. This stochastic process means that there is a modest association ($r \approx 0.21$) between accuracy and resistance to social influence among our agents (Becker et al., 2017).

Network Conditions

Of particular interest to the present work is how different network conditions perform in the general modelling framework outlined above. Specifically, we consider collective accuracy in four conditions: non-communicating agents (i.e., aggregating agents' initial estimates), static networks (i.e., unchanging network structure), and networks to which we apply one of two candidate rewiring algorithms: the *mean-extreme algorithm* and the *polarize algorithm*. Crucially, both candidate algorithms operate solely on the distribution of current estimates in a network with no regard for the true value.

The mean-extreme algorithm aims to increase the average accuracy of individuals in a network by directing social influence towards individuals with potentially erroneous, outlying estimates. The algorithm first calculates the mean estimate in a network at a given time point and identifies which side of the scale midpoint[3] the network's mean estimate lies. If the network's mean estimate is less than the midpoint, the algorithm identifies the agent with the lowest estimate and adds directed, outgoing ties to the three agents with the highest estimates. If the network's mean estimate is greater than the midpoint, the algorithm identifies the agent with the highest estimate and adds directed, outgoing ties to the three agents with the lowest estimates. This procedure effectively brings the estimates of the outliers closer to the mean.

The polarize algorithm aims to maintain the diversity of estimates in a network and prevent a potentially biasing homogenization. It first identifies the two most extreme agents on either side of the current distribution of estimates (i.e., the agent with the highest estimate and the agent with the lowest estimate) and cuts all incoming ties to these agents so as to preserve their beliefs from social influence. Then, the influence of these extreme agents is increased by granting each of them two directed, outgoing ties to "core" agents. These core agents are the four individuals with the median estimates in the network (e.g., in a 16-agent network, the agent with the lowest estimate receives an outgoing tie to the agents with the 7th and 8th lowest estimates, and the agent with the highest estimate receives two outgoing ties to the agents with the 9th and 10th lowest estimates).

Results

Following 500 iterations of each of the 57 estimation tasks in which four matched networks are simulated (i.e., one of each network condition starting from an identical initial network), we assess collective accuracy by calculating the squared error of the mean estimate post-communication, henceforth referred to as collective error squared (*CES*). While other loss functions such as absolute error and square root error may be applicable in some task domains, our pattern of results is consistent across these loss functions and we thus focus on *CES* for the sake of this paper; also because of the theoretical link of *CES* to the Diversity Prediction Theorem.

Across all of the estimation tasks considered, the four network conditions' *CES* was nearly equal on average (non-communicating agents, M = 0.015, SD = 0.031; static networks, M = 0.015, SD = 0.031; mean-extreme networks, M = 0.016, SD = 0.033; polarize networks, M = 0.016, SD = 0.029). However, these averages overlook potential context-dependent effects. Indeed, an analysis of *CES* task-by-task, rather than in aggregate, reveals that mean-extreme networks achieved the highest accuracy on 35 tasks, polarize networks achieved the highest accuracy on 15 tasks, non-communicating agents achieved the highest accuracy on 6 tasks, and static networks achieved the highest accuracy on just 1 task. This observation that the two rewiring conditions outperformed non-communicating agents and static networks on 88% of the simulated estimation tasks suggests that rewiring algorithms may serve as a viable strategy for boosting collective accuracy in social networks.

Omega (Ω): Context-Dependent Network Performance

To better understand the context-dependent effects of the rewiring algorithms, we borrow from Almaatouq, Rahimian, and Alhajri (2020) and use the Ω feature to characterize each of the 57 estimation tasks. By calculating each network condition's average *CES* in each task, we examined how the rewiring algorithms influenced collective accuracy across the Ω parameter space.

In Figure 1, we plot the distribution of Ω (M = 0.542, SD = 0.158) in the estimation tasks considered in our simulations, and the distributions of Ω in the tasks where each network condition was the best performer (the static network condition is excluded because it was the best performer on only

[3]For our scaled data, this midpoint is 0.5 in every task.

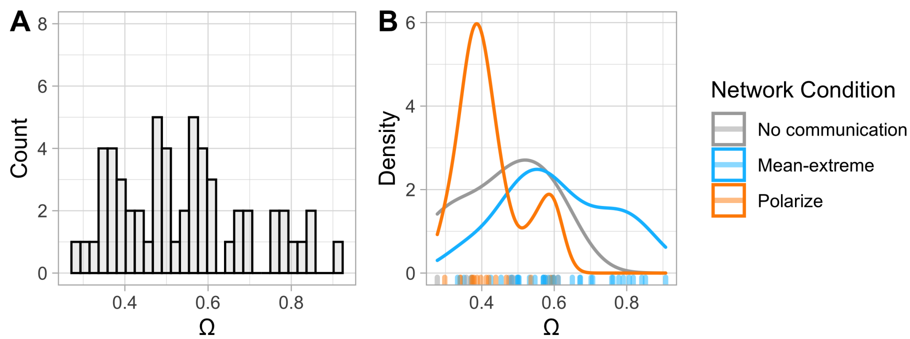

Figure 1: The Ω parameter space. **[A]** The distribution of Ω in the 57 estimation tasks considered. **[B]** The distribution of Ω where each network condition produced the lowest collective error as compared to the other conditions (the static network condition is excluded because it produced the lowest collective error on only one task; Ω = 0.390).

one task; Ω = 0.390). In doing so, we find that the mean-extreme algorithm displays a clear favoritism for tasks with higher Ω ($n = 35$, $M = 0.610$, $SD = 0.151$) as compared to the polarize algorithm, which favors tasks with lower Ω ($n = 15$, $M = 0.423$, $SD = 0.089$). While non-communicating agents were the most accurate on tasks with a mid-range Ω ($n = 6$, $M = 0.464$, $SD = 0.127$), between the borders of the rewiring algorithms' favored Ω values, this analysis encouragingly suggests that the two rewiring algorithms together account for a vast majority of the Ω parameter space.

In Figure 2, we further investigate how the effects on collective accuracy produced by the rewiring algorithms track over Ω. Using the *CES* of non-communicating agents as a baseline condition (i.e., an aggregation of the agents' initial estimates), we calculated three measures for each of the 57 estimation tasks for static, mean-extreme, and polarize networks: the average effect on *CES* (i.e., the average change in error), the average relative effect on *CES* (i.e., the average change in error divided by the average error of non-communicating agents), and the probability of improvement (i.e., the proportion of the 500 iterations of each task where a given network condition was more accurate than matched non-communicating agents). This analysis suggests not only that the different rewiring algorithms prefer different estimation contexts, but that there is an important interaction: the mean-extreme algorithm actively increases collective error on tasks with low Ω and the polarize algorithm actively increases collective error on tasks with high Ω. Therefore, the issue of algorithm selection—choosing one rewiring algorithm or another to apply to a network—seems particularly crucial.

Skewness: A Heuristic for Algorithm Selection

The Ω feature introduced by Almaatouq, Rahimian, and Alhajri (2020) provides useful insight into when different network structures, and thus different rewiring algorithms, are capable of eliciting accurate collective estimates. However, since calculation of Ω requires knowledge of the ground truth, Ω cannot be used for algorithm selection when it would be needed most. For example, consider an online prediction market tasked with estimating the probability that Scotland will hold a new sanctioned independence referendum before May 2024[4]. If one sought to apply a rewiring algorithm to the forecasters' communication network in hopes of boosting the accuracy of the prediction, Ω could not be used to select the appropriate algorithm because the truth is fundamentally uncertain and Ω is thus incalculable. In such circumstances, there is a need for a truth-agnostic cue that can be used for algorithm selection. For this reason, we explored the possibility of using the skewness of each task's initial estimate distribution as such a cue.

In the 57 estimation tasks simulated, Ω displays a significant correlation with the skewness of the initial, pre-communication estimate distribution ($r = 0.413$, $p = 0.001$). As shown in Figure 3, the rewiring algorithms display a clear favoritism for skewness in a manner similar to Ω: mean-extreme networks were the most accurate for tasks with highly skewed estimate distributions ($n = 35$, $M = 11.27$, $SD = 10.71$), polarize networks were the most accurate for tasks with estimate distributions that display low skewness ($n = 15$, $M = 1.56$, $SD = 1.38$), and non-communicating agents were the most accurate on tasks with mid-range skewness ($n = 6$, $M = 2.97$, $SD = 3.26$).

Although skewness may be less explanatory than Ω, this finding suggests that knowledge of the skewness of a task's estimate distribution alone could be useful as a heuristic for algorithm selection. To investigate this further, we re-ran the analyses that are displayed in Figure 2, but with skewness instead of Ω, and found that the same interaction effect is observed in both parameter spaces. Specifically, we note that while the mean-extreme algorithm is capable of producing desirable effects on *CES* for tasks with highly skewed estimate distributions (e.g., decreasing *CES* by 50% in some tasks; Figure 4, B), the polarize algorithm is exceptionally damaging for tasks with a highly skewed estimate distribution (e.g., increasing *CES* by 400% in those same tasks; Figure 4,

[4]https://www.metaculus.com/questions/6369/official-scottish-independence-referendum/

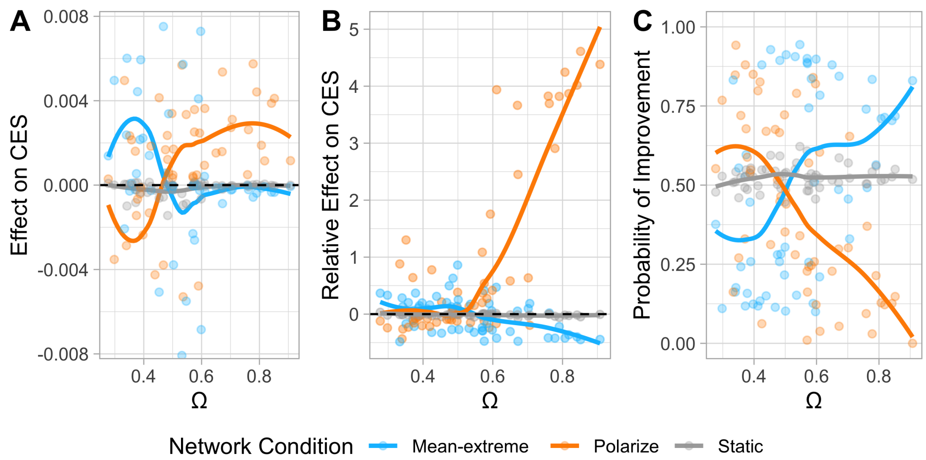

Figure 2: Network performance over Ω as compared to non-communicating agents. **[A]** The average effect on *CES* across Ω (i.e., the average change in *CES* compared to matched non-communicating agents). **[B]** The average relative effect on *CES* across Ω (i.e., the average change in error divided by the average error of matched non-communicating agents) **[C]** The probability of improvement across Ω (i.e., the proportion of the 500 iterations of each task where a given network condition was more accurate than matched non-communicating agents).

B). The reason for this is intuitive. When the estimate distribution is highly skewed and the egregious errors of a few individuals jeopardize collective accuracy, the mean-extreme algorithm will centralize these few individuals around a potentially accurate individual; meanwhile, the polarize algorithm would increase the influence of one of these erroneous individuals by granting them new outgoing edges, effectively broadcasting inaccurate information to the network.

Discussion

The results of our simulations provide a proof of concept that it may be possible to identify distributional characteristics of judgments that allow one to select a rewiring algorithm capable of increasing the accuracy of social networks' collective estimations. Crucially, these are applicable in contexts where there is no track record of individuals' predictive success and the truth or falsity of individual estimates is not (yet) known. Where sufficient ground truth data on accuracy exists, such as in expert judgments of medical scans, that data can unquestionably be used to fine-tune networks of judges (Kurvers et al., 2019). However, that leaves many of the most pressing real-world judgment tasks unaccounted for. In particular, we may want collective judgments to derive high-quality predictions for consequential unique events—and for these events, by definition, ground truth data will be unavailable. A method that enhances collective accuracy in such contexts would thus provide a valuable prediction tool for many domains. The simulation results presented here are encouraging; the next step will be to validate them with experimental studies.

Code Availability

All code used for the simulations and analyses reported in this paper are available on the accompanying OSF project page.

Acknowledgments

This research is supported by a grant from Nesta, the innovation foundation.

References

Almaatouq, A., Noriega-Campero, A., Alotaibi, A., Krafft, P., Moussaid, M., & Pentland, A. (2020). Adaptive social networks promote the wisdom of crowds. *Proceedings of the National Academy of Sciences*, *117*(21), 11379–11386.

Almaatouq, A., Rahimian, M. A., & Alhajri, A. (2020). When social influence promotes the wisdom of crowds. *PsyArXiv*, *version 1*.

Arrow, K. J., Forsythe, R., Gorham, M., Hahn, R., Hanson, R., Ledyard, J. O., ... Zitzewitz, E. (2008). The promise of prediction markets. *Science*, *320*(5878), 877–878.

Becker, J., Brackbill, D., & Centola, D. (2017). Network dynamics of social influence in the wisdom of crowds. *Proceedings of the National Academy of Sciences*, *114*(26), E5070–E5076.

Becker, J., Porter, E., & Centola, D. (2019). The wisdom of partisan crowds. *Proceedings of the National Academy of Sciences*, *116*(22), 10717–10722.

Condorcet, M. J. A. N. C. (1785). *Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix*. Chelsea.

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

Galton, F. (1907). *Vox populi.* Nature Publishing Group.

Grofman, B., Owen, G., & Feld, S. L. (1983). Thirteen theorems in search of the truth. *Theory and Decision*, *15*(3), 261–278.

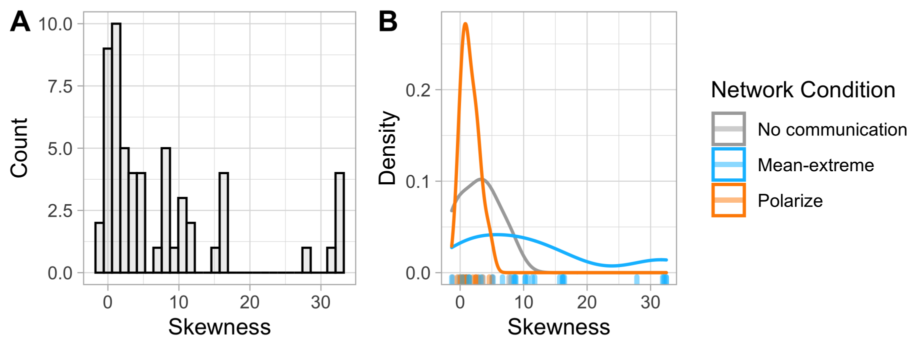

Figure 3: The skewness parameter space. **[A]** The distribution of skewness in the 57 estimation tasks considered. **[B]** The distribution of skewness where each network condition produced the lowest collective error as compared to the other conditions.

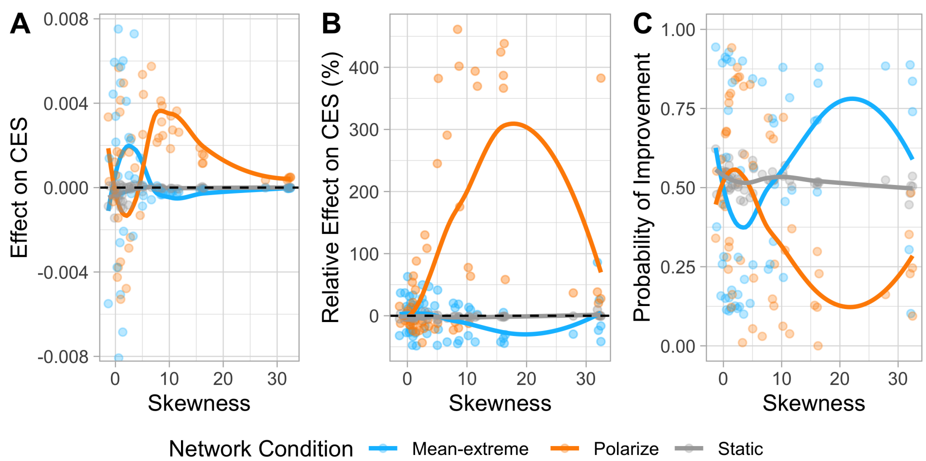

Figure 4: Network performance over skewness as compared to non-communicating agents. These plots display the same analysis as Figure 2, but with skewness instead of Ω. **[A]** The average effect on *CES* across skewness. **[B]** The average relative effect on *CES* across skewness. **[C]** The probability of improvement across skewness.

Gürçay, B., Mellers, B. A., & Baron, J. (2015). The power of social influence on estimation accuracy. *Journal of Behavioral Decision Making*, *28*(3), 250–261.

Hahn, U., Hansen, J. U., & Olsson, E. J. (2018). Truth tracking performance of social networks: how connectivity and clustering can make groups less competent. *Synthese*, 1–31.

Hahn, U., von Sydow, M., & Merdes, C. (2019). How communication can make voters choose less well. *Topics in Cognitive Science*, *11*(1), 194–206.

Howe, J. (2006). The rise of crowdsourcing. *Wired Magazine*, *14*(6), 1–4.

Jönsson, M. L., Hahn, U., & Olsson, E. J. (2015). The kind of group you want to belong to: Effects of group structure on group accuracy. *Cognition*, *142*, 191–204.

Kurvers, R. H., Herzog, S. M., Hertwig, R., Krause, J., Moussaid, M., Argenziano, G., ... Wolf, M. (2019). How to detect high-performing individuals and groups: Decision similarity predicts accuracy. *Science Advances*, *5*(11), eaaw9011.

Ladha, K. K. (1992). The condorcet jury theorem, free speech, and correlated votes. *American Journal of Political Science*, 617–634.

Lazer, D. (2015). The rise of the social algorithm. *Science*, *348*(6239), 1090–1091.

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.

Morgan, M. G. (2014). Use (and abuse) of expert elicitation in support of decision making for public policy. *Proceedings of the National academy of Sciences*, *111*(20), 7176–7184.

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

Mulgan, G. (2018). Artificial intelligence and collective intelligence: the emergence of a new field. *AI & SOCIETY*, *33*(4), 631–632.

Page, S. E. (2008). *The difference: How the power of diversity creates better groups, firms, schools, and societies-new edition*. Princeton University Press.

Simon, J., Bass, T., Boelman, V., & Mulgan, G. (2017). Digital democracy. *The Tools Transforming Political Engagement*, 87.

Surowiecki, J. (2005). *The wisdom of crowds*. Anchor.

Watts, D. J., & Strogatz, S. H. (1998). Collective dynamics of 'small-world'networks. *Nature*, *393*, 440-442.

Wei, H., Pan, Z., Hu, G., Zhang, L., Yang, H., Li, X., & Zhou, X. (2018). Identifying influential nodes based on network representation learning in complex networks. *PloS one*, *13*(7).

Wolfers, J., & Zitzewitz, E. (2004). Prediction markets. *Journal of Economic Perspectives*, *18*(2), 107–126.