

# The Wisdom of the Network: How Adaptive Networks Promote Collective Intelligence

Abdullah Almaatouq<sup>1†</sup>, Alejandro Noriega Campero<sup>1†</sup>, Peter Krafft<sup>1</sup>,  
Abdulrahman Alotaibi<sup>1</sup>, Mehdi Moussaid<sup>2</sup>, Alex (Sandy) Pentland<sup>1\*</sup>

<sup>1</sup>Massachusetts Institute of Technology, Cambridge, MA, USA

<sup>2</sup>Max Planck Institute for Human Development, Berlin, Germany

<sup>†</sup> Authors contributed equally to this work

\* To whom correspondence should be addressed; E-mail: pentland@mit.edu

**Social networks continuously change as people create new ties and break existing ones. It is widely noted that our social embedding exerts an influence on what information we receive, and on how we form beliefs and make decisions. However, studies on collective intelligence overlook the dynamic nature of social networks. It remains unknown (1) how network structures adapt to the performances of individuals, and (2) whether this adaptation promotes the accuracy of individual and collective decisions. Here, we answer these questions through a series of behavioral experiments and simulations. Our results reveal that groups of people embedded in dynamic social networks can adapt to non-stationary information environments. As a result, individual and collective accuracy is substantially improved over static networks and unconnected groups. Moreover, we show that groups in dynamic networks far outperform their best-performing member and that even the best member's judgment substantially benefits from group engagement. Thereby, our findings demonstrate the role of dynamic social networks as adaptive mechanisms for refining individual judgments and harnessing groups' collective intelligence.**

Intelligent systems, both natural and artificial, rely on feedback, empirical learning, and adaptation (1). Such systems are widespread, and can often be viewed as networks of interacting entities that dynamically evolve over time. Cell reproduction, for example, relies on protein networks to combine sensory inputs into gene expression choices adapted to environmental conditions (2). Neurons in the brain dynamically rewire for human learning (3). Animal swarms modify their connectivity to enhance their collective intelligence (4). Dynamic interaction networks have been shown to promote human cooperation (5), and culture transmission networks over generations enabled human groups to develop technologies above any individual's capabilities (6). In the artificial realm, prominent machine learning algorithms rely on similar logics, where dynamically updated networks integrate input signals into useful output (7, 8). Across the board, networks' dynamic properties embody key mechanisms that enable systems to adapt to environmental changes (9, 10).

In our view, the information processing capabilities of interacting human groups are no exception. People's behavior, opinion formation, and decision-making are deeply rooted in cumulative bodies of social information (11), accessed through social networks formed by choices of whom we friend (12, 13), follow (14), call (15, 16), imitate (17, 18), trust (19, 20), and cooperate with (5, 21, 22). Moreover, peer choices are frequently revised, most often based on notions of peer success and reliability, or approximations such as reputation, popularity, and socio-demographics (23–27).

It is widely noted, however, that social influence strongly correlates individuals' judgment in estimation tasks (28–31), compromising the first of two assumptions underlying common statistical accounts of 'wisdom-of-crowds' phenomena (32): namely, that (i) individual estimates are uncorrelated, or negatively correlated, and (ii) individuals are correct in mean expectation (31, 33). In recent years, numerous studies have offered conflicting findings, showing that social interaction can either significantly benefit group and individual estimates (30, 34, 35),

or, conversely, lead them astray by inducing social bias, herding, and group-think (28, 29, 31). Notably, both theoretical and experimental work has been limited mainly to frameworks where agents are randomly placed in static social structures—dyads (34, 36), groups (29, 35, 37), or networks (30, 31)—and has found that these divergent effects are a function of whether well-informed individuals are placed in prominent positions in the network structure (30, 31), and how self-confident they are (34, 36, 38).

However, unlike what is assumed in most existing work, the social networks we live in are not random or imposed by external forces, but emerge shaped by endogenous social processes and gradual evolution. The present study builds on the observation that agent characteristics, such as skill and information access, are not randomly located in network structures. Instead, their distribution is often the outcome of social heuristics that form and break ties influenced by social cues (23, 24, 39). Intuitively, groups can benefit from awarding centrality to and amplifying the influence of well-informed individuals. Hence we hypothesize that dynamic social influence networks may be central to human collective intelligence, acting as core mechanisms by which crowds, which may not initially be wise, evolve into wisdom, adapting to biased and potentially non-stationary information environments.

To test these hypotheses, we developed a web-based experiment that allows us to identify the role of dynamic networks in fostering an adaptive ‘wisdom of crowds.’ Participants ( $n = 719$ ) from Amazon Mechanical Turk (population details in SI) engaged in a sequence of 20 estimation tasks. Each task consisted of estimating the correlation of a scatter plot, and monetary prizes were awarded in proportion to performance. Participants were randomly allocated to groups of 12 ( $n = 60$  groups), and each group was randomized to one of three treatment conditions: a *solo* condition, where each individual solved the sequence of tasks in isolation; a *static* condition, in which participants were randomly placed in static communication networks; and a *dynamic* condition, in which participants at each round were allowed to select up to three neighbors to

communicate with. Fig. 1A illustrates the experiment flow according to each condition. At each round, participants were initially asked to submit an independent guess. Then those in *static* and *dynamic* conditions entered a social exposure stage, where they could observe the answers of their network peers, update their own, and see peers' updates in real time. After submitting a final guess, participants in all conditions were shown feedback on their score and the correct answer. Lastly, those in the *dynamic* condition were shown scores of all participants and allowed to revise which peers to follow in subsequent rounds.

The estimation task was designed so that it allowed manipulation of the quality of information provided to participants. Scatter plots with three levels of difficulty were used (varying the number of points, and adding outliers or non-linearities; see Figure 1B). At every round, all plots seen by participants shared an identical true correlation, but difficulty levels could differ among them (40). The design allowed the simulation of a shock to the distribution of information among participants. Specifically, each participant experienced a constant difficulty level across the first ten rounds; then, at round eleven, we introduced shocks by reshuffling difficulties to new levels that remained constant thereafter. Participants were not informed about the difficulty levels they or their peers faced.

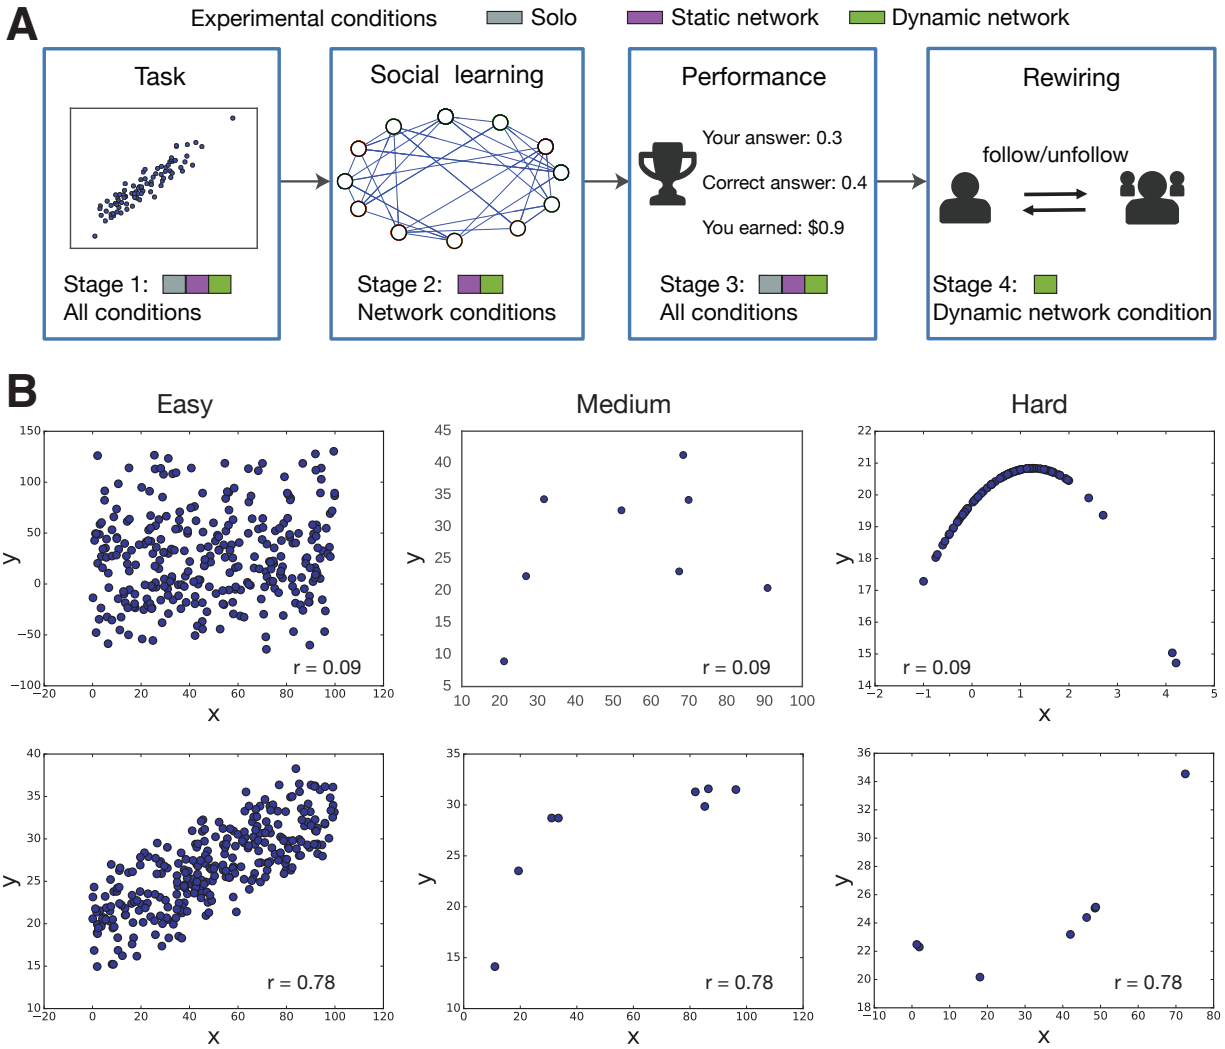


Figure 1: **Experimental design.** Panel (A) shows an illustration of the experiment design. Panel (B) illustrates examples of the scatter plots used in the experiment. For any given round, all participants saw plots that shared an identical true correlation, but difficulty levels could differ among them. Participants were not informed about the difficulty level they or other participants were facing.

## Results

### Individual and collective outcomes

We first compared individual- and group-level errors across conditions. Evolutionary reasoning suggests that people's propensity to imitate follows from its direct benefits to the individual, but it may, nonetheless, induce benefits to the population as a whole (39). Fig. 2 shows the individual and group error rates of the static and dynamic conditions—using the arithmetic mean as group estimate—normalized with respect to baseline errors in the solo condition. Overall, we find that participants in dynamic networks achieved the lowest error rates, averaging 33% lower individual error ( $P < 10^{-5}$ ), and 34% lower group error, compared to participants in static networks ( $P < 10^{-4}$ ). In particular, the performance edge of groups in dynamic networks was larger in periods where networks had adapted to their information environment (rounds  $[6, 10] \cup [16, 20]$ , the 'adapted periods'), reducing individual error by 45% ( $P < 10^{-10}$ ) and group error by 47% ( $P < 10^{-10}$ ) compared to groups connected by static networks. Hence, these results support our primary hypothesis that dynamic networks' adaptiveness can benefit both individual and collective judgment.

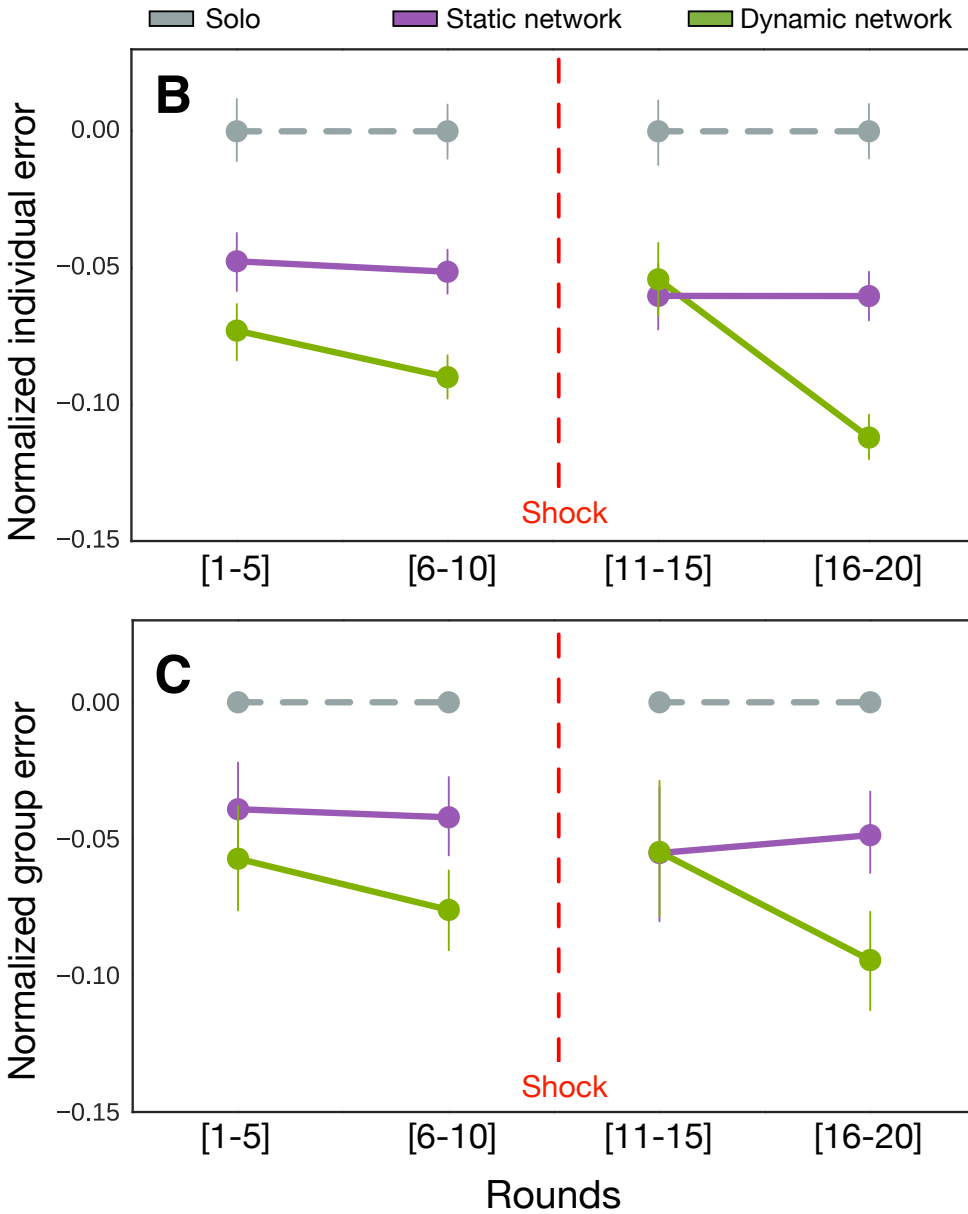


Figure 2: **Individual and collective outcomes.** Groups connected by dynamic influence networks incur substantially lower individual errors ( $P < 10^{-5}$ ) as shown in Panel (A) and collective errors ( $P < 0.001$ ) in Panel (B). The reduction is notably larger and more significant in periods where networks had adapted to the information environment (i.e., rounds [6, 10] and [16, 20]). Errors are normalized with respect to average errors in the *solo* condition. Statistics on individual-level outcomes are based on White cluster-robust standard errors. Error bars indicate 95% confidence intervals.

## **Adaptive Mechanisms**

Two social mechanisms underlie the favorable performance of networked groups. First, dynamic networks adaptively centralized over high-performing individuals. This behavior was predicted by abundant evidence from cognitive science and evolutionary anthropology, which indicate that people naturally engage in selective social learning (6, 23, 39)—i.e., the use of social cues related to peer competence and reliability to choose whom we pay attention to and learn from selectively. Figs. 3A and 3B show that participants in dynamic networks consistently used peers' past performance information as success cues to guide their peer choices. As rounds elapsed, performance information accrued, and social networks evolved from fully distributed into networks that amplified the influence of well-informed individuals. Upon receiving an information shock, the networks slightly decentralized, entering a transient exploration stage before finding a configuration adapted to the new distribution of information among participants (see Fig. 4).



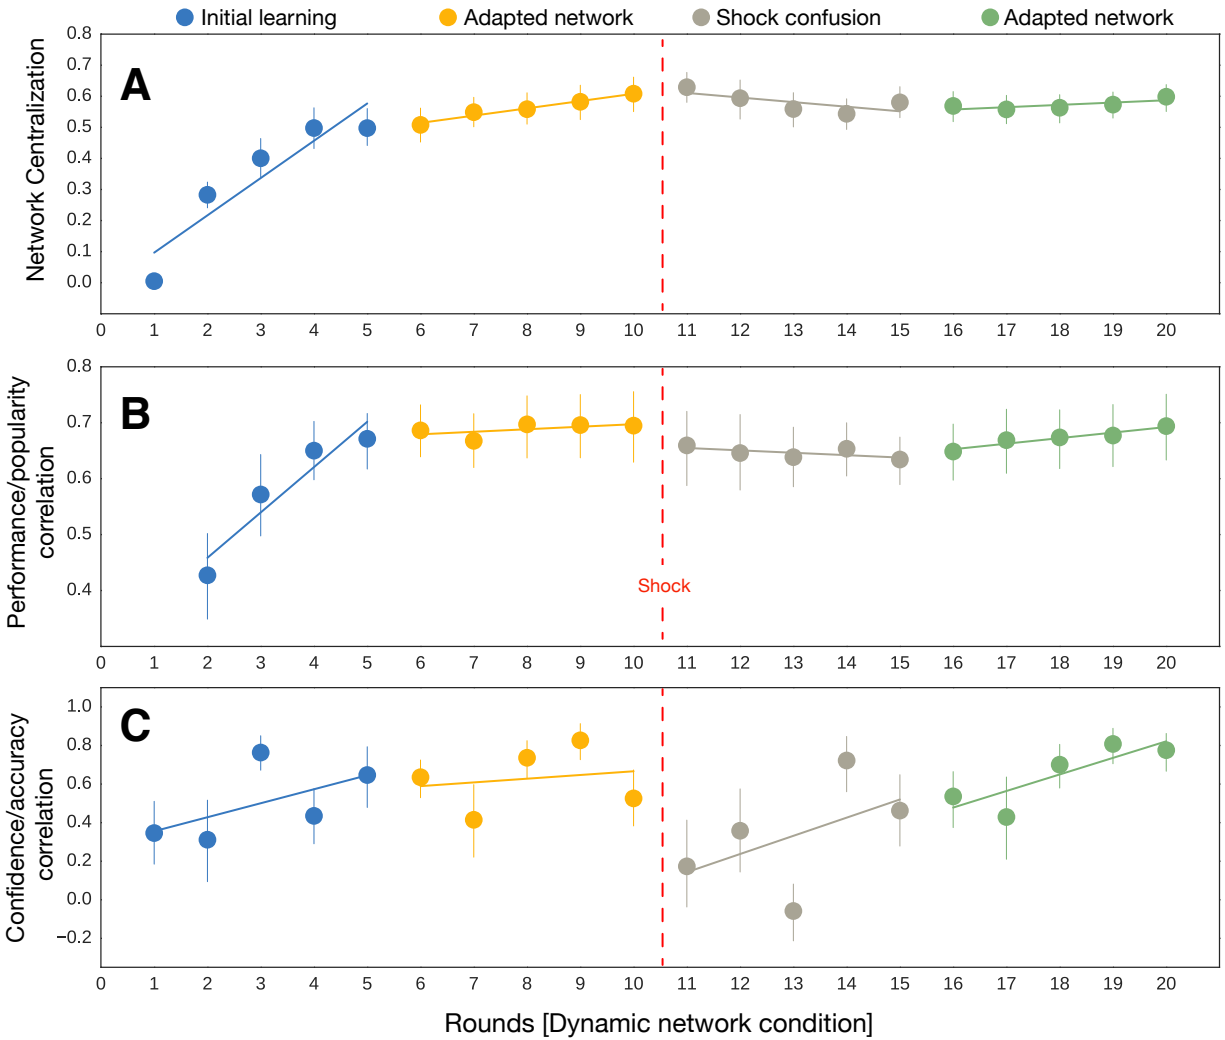
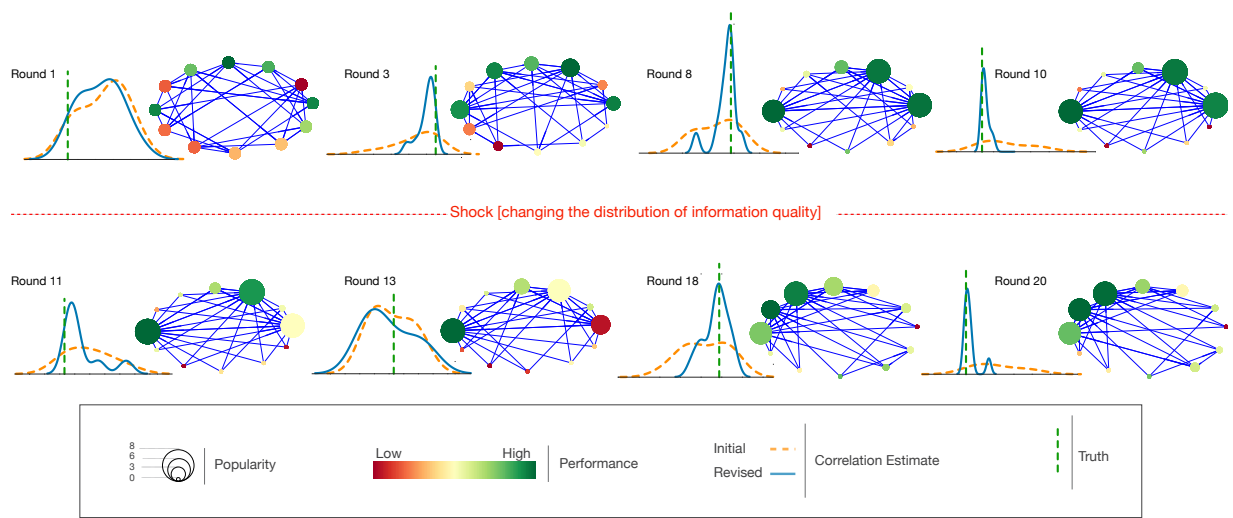


Figure 3: **Mechanisms promoting collective intelligence in dynamic networks.** Panel (A) shows that the network becomes more centralized with time (Freeman global centralization—i.e., how far the network is from a star network). Panel (B) depicts the relation between performance (i.e., average error) and popularity (i.e., number of followers). Panel (c) shows the relationship between accuracy of initial estimate and confidence (i.e., resistance to social influence). Error bars indicate 95% confidence intervals.

A centralization mechanism alone could suggest that group members may merely follow and imitate the best individual among them, hence bounding collective performance by that of the group’s top performer. However, research on the *two-heads-better-than-one* effect indicates that, in the simpler case of dyads, even the best individual can benefit from social inter-

action (34, 36); and that the critical mechanism enabling this effect is a positive relationship between individuals' accuracy and their confidence. Fig. 4C shows that participants in dynamic networks had, overall, a positive correlation between the accuracy of their initial estimates and their self-confidence (measured in terms of resistance to social influence). Participants were likely to rely on private judgments whenever these were accurate and likely to rely on social information otherwise. Fig. 4C also shows that, as rounds elapsed, participants used task feedback to calibrate their accuracy-confidence relation gradually, and were able to readapt gradually upon the shock. Consistent with prior literature (30, 38), a positive correlation of confidence and accuracy was found in static networks too (see S18 in SI), explaining their favorable performance compared to unconnected groups.



**Figure 4: An illustrative example of the network dynamics.** The circle color represents performance. The size of each circle represents the number of followers (i.e., popularity). The dashed orange line is the distribution of estimates prior to social influence, the blue solid line is the distribution of post-social influence estimates, while the dashed vertical line is the true correlation.

## Mean-Variance Trade-off

The joint effect of centralization and confidence mechanisms explains the adaptive advantage of dynamic networks. Moreover, it suggests that their collective performance may not be bounded by that of the best individual, and that even the best individual may benefit from network interaction. To test these implications, we generalize the use of group means as collective estimates, common in ‘wisdom of crowds’ studies, and analyze the performance of *top-k* estimates—that is, collective estimates where only the guesses of the  $k$  best-performing group members are averaged. Top- $k$  subsets within each group were computed based on post-hoc individual performances across all rounds. In particular, *top-12* estimates correspond to the group mean, and *top-1* to estimates of the group’s best-performing individual. Fig. 5 reports the mean and standard deviation of estimation errors incurred by *top-k* estimates during the adapted periods. Ideal estimates would minimize both mean error and variability, approaching the lower left end of the trade-off space.

The shape of *top-k* curves reveals that, as we remove low-performing individuals (from  $k = 12$  to  $k = 1$ ), estimates initially improve in both mean and standard deviation. Then, as we further curate the crowd beyond  $k = 6$ , *top-k* estimates trade off between decreasing mean error and increasing variability, and finally regress in both objectives as  $k \rightarrow 1$ . Comparison across conditions shows that, for any  $k \in [1, 12]$ , dynamic influence networks improved estimation errors in terms of both mean and standard deviation. In particular, Fig. 5 shows that the full-group average in dynamic networks got 28% lower error and 48% less variability than the best individual in the *solo* groups (*dynamic top-12* vs. *solo top-1*;  $P < 10^{-2}$ ). Moreover, even the best individual derived substantial benefits from social interaction, averaging 32% lower error and 38% less variability when forming and revising social connections rather than working in isolation (*dynamic top-1* vs. *solo top-1*;  $P < 10^{-2}$ ).

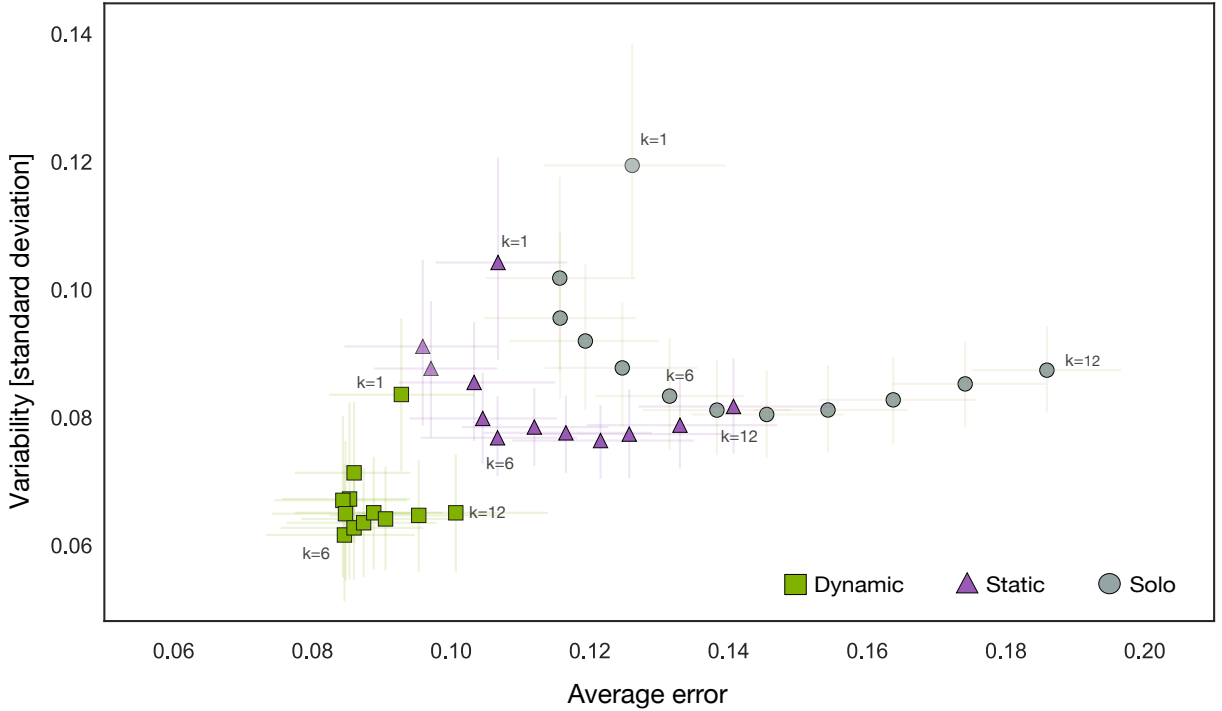


Figure 5: **Mean-variance trade-off.** Mean and standard deviation of absolute errors incurred by  $top-k$  estimates during the adapted periods (rounds  $\in [6, 10] \cup [16, 20]$ ).  $Top-12$  estimates correspond to the full-group mean, and  $top-1$  to the group’s best individual. Within each condition,  $top-k$  trade-off curves first gain in both objectives, then trade off lower average error for higher variability, and finally regress in both objectives as  $k \rightarrow 1$ . Across conditions, for any  $k \in [1, 12]$ , groups in the *dynamic* condition outperformed groups in the *static* and *solo* conditions. Moreover, the full-group mean of *dynamic* networks averaged 28% lower error and 48% less variability than the best individual playing *solo* (*dynamic top-12* vs. *solo top-1*;  $P < 10^{-2}$ ); and the best individual in the *dynamic* condition averaged 32% lower error and 38% less variability than her analogue in *solo* (*dynamic top-1* vs. *solo top-1*;  $P < 10^{-2}$ ). Bars indicate 95% confidence intervals.

We implemented numerical simulations to further explore the effects of dynamic influence networks in collective estimation. We simulated interacting agents that update beliefs according to a DeGroot model (41), and rewire social connections according to a performance-based preferential attachment process (see Section S7). Results of these simulations corroborate our experimental results. Dynamic networks adapted to shocks by shifting influence weight to agents with better information, substantially decreasing individual and group error. Results also indic-

ate a relationship between the network learning rate—the network’s sensitivity to changes in agents’ performance—and the arrival rate of environmental shocks (Fig. S21). Networks with faster learning rates could adapt to environments with frequent information shocks. Conversely, networks with slower learning rates could leverage longer learning periods, achieving lower error rates when in environments characterized by infrequent shocks. This short-term versus long-term accuracy trade-off implies that optimal network learning rates depend on the pace at which the information environment changes, analogous to notions of regularization and learning rates in artificial intelligence algorithms (7, 8).

## **Discussion**

Existing work on collective intelligence and social influence has not considered the rewiring of influence networks, a phenomenon that is widespread in natural situations. We have shown that dynamic influence networks can adapt to biased and non-stationary environments, inducing individual and collective beliefs even far more accurate than the independent beliefs of the best-performing individual. We expect this adaptive systems view on collective intelligence to further sprout connections with fields such as evolutionary biology and artificial intelligence. Our findings also emphasize the relevance of high-quality and accessible social information, and the relation between the stability of social influence ties and how fast the information environment changes. The insights here provided suggest design guidelines germane to real-world collective intelligence mechanisms, in contexts such as commodity markets, social trading platforms, crowdfunding, crowd work, prediction markets, and online education (e.g., MOOCs).

## **Methods**

We recruited 720 participants using Amazon Mechanical Turk (AMT; see Supplementary Information). Following a between-subjects design, each subject participated only once, was

randomly assigned to one of the three conditions, and did not know that other conditions existed. Participants interacted anonymously over the Internet using web-application playable in a browser window. Consistent with standard payment rates on AMT, participants received a show-up fee of \$2.0, and then had the opportunity to earn an additional bonus of up to \$8 based on their performance in the game—average hourly pay of \$15. We excluded one participant from our analysis who was assigned to the solo (control) condition and dropped from the game without completing any of the tasks. Consequently, the analysis of this study is based on 719 valid participants. No statistical methods were used to predetermine sample size. The investigators were not blinded to allocation during analysis.

Participants were randomly allocated into groups of size 12, and each group was randomized to one of three treatment conditions: a *solo* condition, where each individual solved the sequence of tasks in isolation; a *static* condition, in which participants were randomly placed in static communication networks; and a *dynamic* condition, in which participants at each round were allowed to select up to three neighbors to communicate with. The initial communication network for the *static* and *dynamic* conditions were the same (i.e., each participant had three random incoming edges and three random outgoing ones). At each round, participants were initially asked to submit an independent guess. Then those in *static* and *dynamic* conditions entered a social exposure stage, where they could observe the answers of their network peers, update their own and see peers' updates in real time. After submitting a final guess, participants in all conditions were shown feedback on their score and the correct answer. Lastly, those in the *dynamic* condition were shown scores of all participants and allowed to revise which peers to follow in subsequent rounds.

Our experimental task consisted of estimating the correlation from a scatter plot. Participants engaged in 20 correlation estimation tasks. The correlation estimation task was designed so that it allowed manipulation of the task difficulty. We manipulated the difficulty of a scatter

plot by varying the number of points, adding outliers, or non-linear relationship between the two variables (see Supplementary Information for visual examples). At every round, all plots seen by participants shared an identical true correlation (i.e., the correct answer for everyone is the same), but difficulty levels could differ among them. The design allowed the simulation of a shock to the participants. Specifically, each participant experienced a constant difficulty level across the first ten rounds; then, at round eleven, we introduced shocks by reshuffling difficulties to new levels that remained constant after that. Participants were not informed about the difficulty levels they or their peers faced.

To evaluate the difference in performance between the different treatment conditions, we used a treatment dummy to estimate a linear regression for the average normalized error (i.e., the average across the 20 rounds). We examined the effects at two levels of analyses: individual outcomes (i.e., one data point per user;  $n = 719$ ); and 2) group-level outcomes (i.e., one data point per-group). Statistics on individual-level outcomes are based on White cluster-robust standard errors to account for correlated error terms within each group. Statistics on group-level outcomes are based on two-tailed t-test. We also performed various statistical robustness checks (see Supplementary Information).

The study was reviewed and approved by the Committee on the Use of Humans as Experimental participants (COUHES) at MIT. All participants provided an explicit consent to participants in this study and COUHES approved the consent procedure.

The data for the experiment will be publicly available on Harvard Dataverse upon acceptance.

## References

1. N. Wiener, *Cybernetics or Control and Communication in the Animal and the Machine*, vol. 25 (MIT press, 1961).

2. D. H. Erwin, E. H. Davidson, *Nature Reviews Genetics* **10**, 141 (2009).
3. D. A. Gutnisky, V. Dragoi, *Nature* **452**, 220 (2008).
4. Y. Katz, K. Tunstrøm, C. C. Ioannou, C. Huepe, I. D. Couzin, *Proceedings of the National Academy of Sciences* **108**, 18720 (2011).
5. D. G. Rand, S. Arbesman, N. A. Christakis, *Proceedings of the National Academy of Sciences* **108**, 19193 (2011).
6. J. Henrich, *The secret of our success: how culture is driving human evolution, domesticating our species, and making us smarter* (Princeton University Press, 2015).
7. L. Bottou, *On-line learning in neural networks* **17**, 142 (1998).
8. K. P. Murphy, Dynamic bayesian networks, Ph.D. thesis, University of California, Berkeley (2002).
9. P. Holme, J. Saramäki, *Physics reports* **519**, 97 (2012).
10. A. Li, S. P. Cornelius, Y.-Y. Liu, L. Wang, A.-L. Barabási, *Science* **358**, 1042 (2017).
11. A. Bandura, R. H. Walters (1977).
12. A. Almaatouq, L. Radaelli, A. Pentland, E. Shmueli, *PloS one* **11**, e0151588 (2016).
13. H. Wang, B. Wellman, *American Behavioral Scientist* **53**, 1148 (2010).
14. J. Sutton, *et al.*, *Proceedings of the National Academy of Sciences* **112**, 14793 (2015).
15. N. Eagle, A. S. Pentland, D. Lazer, *Proceedings of the national academy of sciences* **106**, 15274 (2009).
16. J.-P. Onnela, *et al.*, *Proceedings of the national academy of sciences* **104**, 7332 (2007).



17. M. G. Zimmermann, V. M. Eguíluz, M. San Miguel, *Physical Review E* **69**, 065102 (2004).
18. P. Shih, *et al.*, *Neuropsychologia* **48**, 2931 (2010).
19. W. S. Chow, L. S. Chan, *Information & management* **45**, 458 (2008).
20. S. Valenzuela, N. Park, K. F. Kee, *Journal of Computer-Mediated Communication* **14**, 875 (2009).
21. J. Wang, S. Suri, D. J. Watts, *Proceedings of the National Academy of Sciences* **109**, 14363 (2012).
22. B. S. Frey, S. Meier, *The American Economic Review* **94**, 1717 (2004).
23. T. N. Wisdom, X. Song, R. L. Goldstone, *Cognitive science* **37**, 1383 (2013).
24. J. Henrich, M. Chudek, R. Boyd, *Phil. Trans. R. Soc. B* **370**, 20150013 (2015).
25. J. A. Cuesta, C. Gracia-Lázaro, A. Ferrer, Y. Moreno, A. Sánchez, *Scientific reports* **5** (2015).
26. E. Gallo, C. Yan, *Proceedings of the National Academy of Sciences* **112**, 3647 (2015).
27. G. Miritello, E. Moro, R. Lara, *Physical Review E* **83**, 045102 (2011).
28. L. Muchnik, S. Aral, S. J. Taylor, *Science* **341**, 647 (2013).
29. J. Lorenz, H. Rauhut, F. Schweitzer, D. Helbing, *Proceedings of the National Academy of Sciences* **108**, 9020 (2011).
30. J. Becker, D. Brackbill, D. Centola, *Proceedings of the National Academy of Sciences* p. 201615978 (2017).
31. B. Golub, M. O. Jackson, *American Economic Journal: Microeconomics* **2**, 112 (2010).

32. J. Surowiecki, *The wisdom of crowds* (Anchor, 2005).
33. F. Galton, *Nature* **75**, 450 (1907).
34. B. Bahrami, *et al.*, *Science* **329**, 1081 (2010).
35. J. Navajas, T. Niella, G. Garbulsky, B. Bahrami, M. Sigman, *Nature Human Behaviour* p. 1 (2018).
36. A. Koriat, *Science* **336**, 360 (2012).
37. A. W. Woolley, C. F. Chabris, A. Pentland, N. Hashmi, T. W. Malone, *science* **330**, 686 (2010).
38. G. Madirolas, G. G. de Polavieja, *PLoS Comput Biol* **11**, e1004594 (2015).
39. R. Boyd, P. J. Richerson, J. Henrich, *Proceedings of the National Academy of Sciences* **108**, 10918 (2011).
40. M. Moussaïd, S. M. Herzog, J. E. Kämmer, R. Hertwig, *Proceedings of the National Academy of Sciences* p. 201611998 (2017).
41. M. H. DeGroot, *Journal of the American Statistical Association* **69**, 118 (1974).

## **Acknowledgments**

All data files will be uploaded as supplementary materials. We thank Elaine Lin for her help in developing the user interface for the experiment. Abdullah Almaatouq thanks King Abdulaziz City for Science & Technology, and Abdulrahman Alotaibi thanks Kuwait University, for their financial support.

## **Author contributions**

A.N.C., A.A., and A.P. conceived the project. A.A., A.N.C., P.K., and M.M. designed the experiment. A.S.A. and A.A. developed the experimentation software. A.N.C. and A.A. collected the data, performed calculations, analyzed results, and produced figures. A.N.C, A.A., P.K, A.S.A., M.M, and A.P. wrote the manuscript.

Supplementary Materials for:

# The Wisdom of the Network: How Adaptive Networks Promote Collective Intelligence

Abdullah Almaatouq<sup>1\*</sup>, Alejandro Noriega-Campero<sup>1\*</sup>, Peter Krafft<sup>1</sup>,  
Abdulrahman Alotaibi<sup>1</sup>, Mehdi Moussaid<sup>2</sup>, Alex (Sandy) Pentland<sup>1,a</sup>

<sup>1</sup>Massachusetts Institute of Technology,  
77 Massachusetts Ave, Cambridge, MA, USA  
<sup>2</sup>Max Planck Institute for Human Development  
14195 Berlin, Germany

\* Authors contributed equally to this work

# 1 Experimental Setup

We aim at causal identification of the role of dynamic networks in fostering an adaptive ‘wisdom of crowds’. For this, we conducted an experiment with real monetary stakes, where:

- The core estimation task of the experiment is a ‘Guess the correlation’ game, where participants are shown a scatter plot of two variables and asked to estimate their correlation (See Section 1.2.2 in the SI).
- We randomly form groups (three conditions) and ask participants to engage in a sequence of 20 correlation tasks.
- The experiment tests three treatment conditions, building up from the ‘wisdom of the crowd’ to the ‘wisdom of dynamic networks’ (See Figure S3) :
  - Independent Estimates [Wisdom of Crowds]: Participants make independent estimates about the correlation of two variables. This condition corresponds to the baseline wisdom of the crowd context (more details in section 1.2.3).
  - Social Learning [Wisdom of Static Networks]: Participants make independent estimates, but then engage in a stage of active social learning, where they are exposed to the estimates of their ego-network peers in real time. This context is analogous to that studied by work at the intersection of the ‘wisdom of crowds’ and social learning (more details in section 1.2.4).
  - Selective Social Learning [Wisdom of Dynamic Networks]: This condition builds on condition (2), and adds the possibility for participants to choose and periodically reevaluate who to follow and be influenced by. After each task round, participants are shown performance information of all peers in the group, and are allowed to follow/unfollow a bounded number of peers (more details in section 1.2.5).

- The experiment is designed to expose the role of dynamic social learning networks as fundamental mechanisms that allow collective intelligent systems to adapt to changes in their environment:
  - In guess the correlation tasks, for any given correlation value we can regulate the amount and structure of information provided to each participant and manipulate it over time by implementing three difficulty levels: easy, medium, hard. Additional details in section 1.2.2.
  - For example, the number of points, existence of outliers, non-linear structures
  - In every round, the correct answer of the correlation was always identical, but the difficulty level could differ between participants.
  - We provide a change in the environment at round 10 (we call it a shock) and thereby we simulate non-stationary distributions of information among participants.
- We ran the experiment on Amazon Mechanical Turk ( $N = 720$ ) where participants were rewarded based on their performance:
  - Participants received noncompetitive monetary incentives for good individual guesses.
  - The rewards applied to all rounds to make sure that individuals took all of their decisions seriously.
  - The incentives trigger the ability to use information of others only for improving own estimates to get closer to the truth and not for aligning with others for the sake of conformity.
  - Mentioned the numerical simulation

## **1.1 Amazon’s Mechanical Turk**

Amazon’s Mechanical Turk (<http://mturk.com>) is an online labor market with a large and diverse pool of people ready to promptly perform tasks for pay (called human intelligence tasks, or HITs). Typical tasks include image labeling, sentiment analysis, or classification of URLs. Additionally, Mechanical Turk is increasingly becoming a popular tool for behavioral scientists as well. Studies from across the social sciences have systematically replicated classic results from psychology and economics with data obtained from such online labor markets and deemed online experiments to be as reliable as that obtained via traditional methods (1–5). Accordingly, we posted each of our experimental sessions as an external HIT (a URL of our web application, is displayed in a frame in the Worker’s web browser) and recruited workers to participate in the experiment. Workers could choose to accept the HIT, at which point the work was officially assigned to them and they could begin participating in the study.

## **1.2 Experiment design**

In this section, we present the design of our experiments conducted online with a sample of participants recruited via MTurk. Our experiment follows the between-group design and implements a repeated ‘Guess the Correlation Game’, where participants are shown a scatter plot of two variables and asked to estimate their correlation.

### **1.2.1 Participant Recruitment**

All participants were recruited on MTurk by posting a HIT for the experiment, entitled “Guess the correlation and win up to \$10”, a neutral title that was accurate without disclosing the purpose of the experiment. The study (Approval#: 1509172301) was reviewed and approved by the Committee on the Use of Humans as Experimental participants (COUHES) at MIT. All participants provided an explicit consent to participate in this study and COUHES approved

the consent procedure. All data collected in the experiment could be associated only with participant’s Amazon Worker ID on MTurk, not with any personally-identifiable information. All players remain anonymous for the entire study. Descriptive statistics on the participants is shown in Figure S1. At the beginning of a session, participants read on-screen instructions for the condition they are randomly assigned to (more details on conditions later) and could start the experiment only once they have completed a set of comprehension questions.

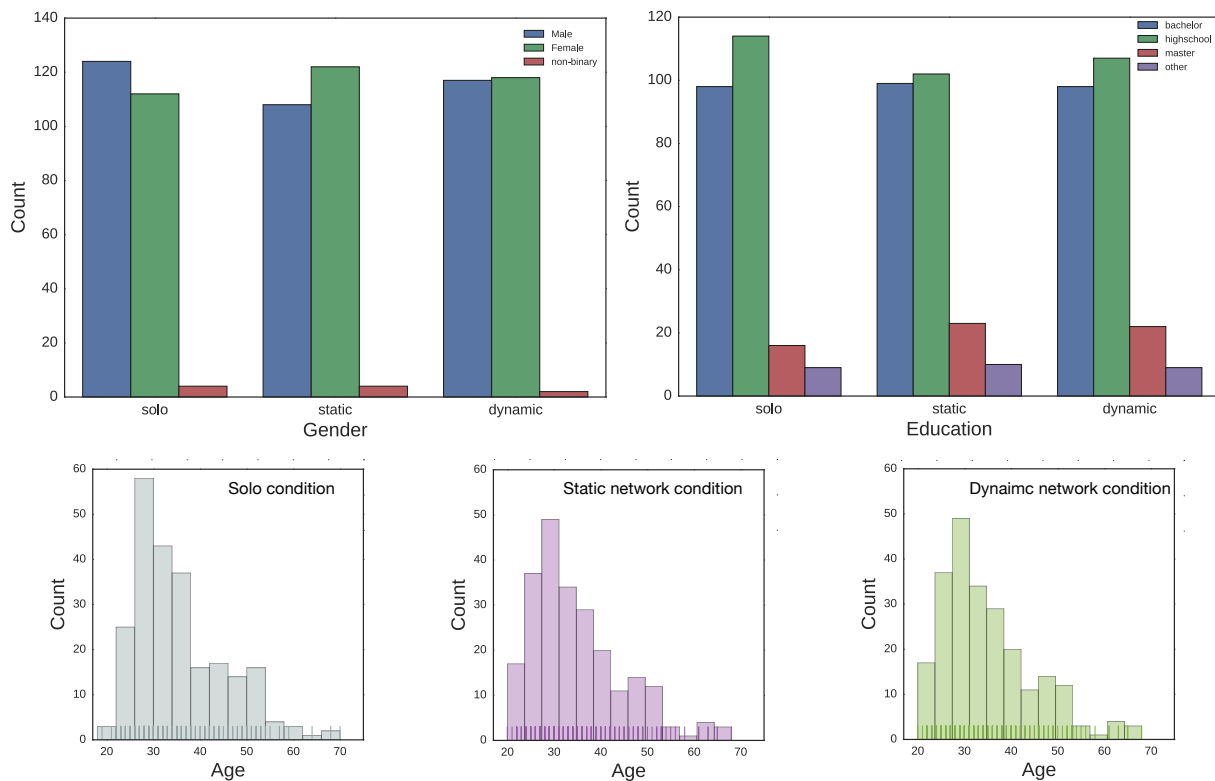


Figure S1: The gender, education, and age breakdowns for our recruited participants via Amazon Mechanical Turk. participants were assigned to their condition randomly.

### 1.2.2 Guess the Correlation Task

participants were prompted to estimate the correlation from a scatter plot (namely, ‘Guess the correlation’ game) and were awarded a monetary prize based on the accuracy of their final estimate. This estimation task is designed to expose the mechanisms that allow intelligent



systems to adapt to changes in their information environment. We can influence the performance level of participants by implementing three difficulty levels (e.g., adjusting the number of points or linearity): easy, medium, and hard. In every round, the correct answer of the correlation was always identical, but the difficulty level could differ between participants. The difficulty level that each member experienced remained constant across the first 10 rounds. Then we simulate a shock after round 10 by changing the difficulty levels for the remaining of the experiment. Participants were not informed about the difficulty level they or other participants were facing. Example of the scatter plots used in the experiment are shown in Figure S2.

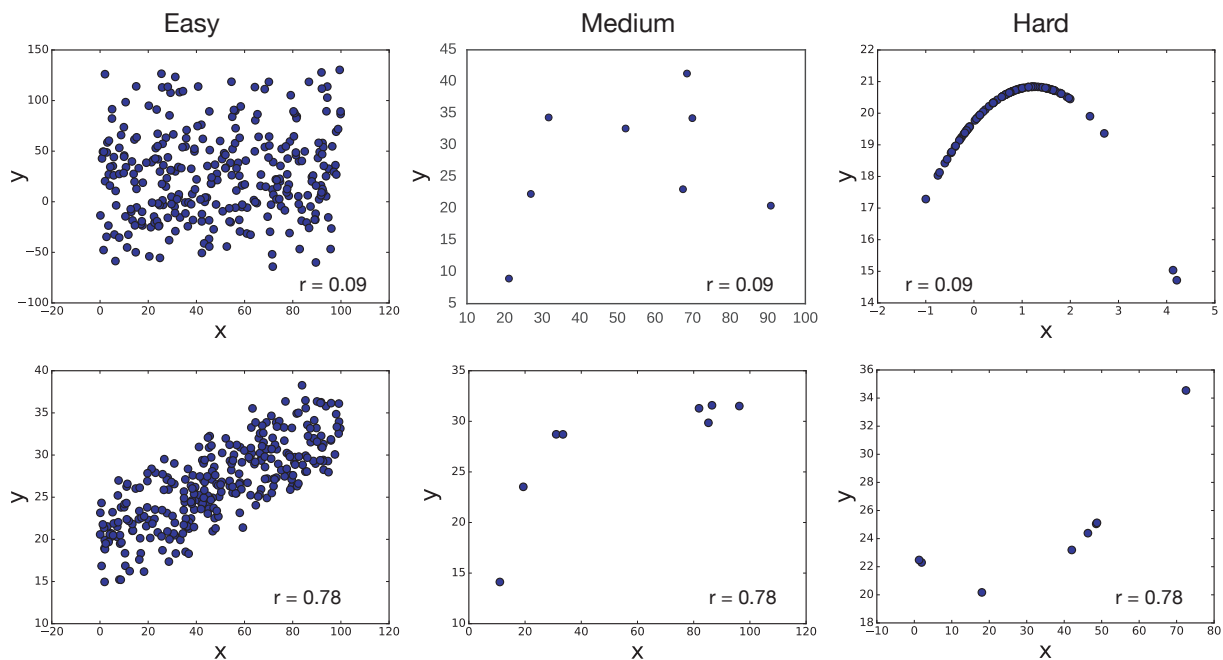


Figure S2: Example of the scatter plots used in the experiment.

### 1.2.3 Condition 1: Solo [Wisdom of Crowds]

In this condition participants make independent guesses about the correlation of two variables (See figure S3). Once an independent guess had been made, the participant learned the result of that round (i.e., how much they earned, what was the correct correlation). In this condition, there was no social information about the performance and estimates of the others.

Collective estimates are then formed by taking averages of individual estimates. This condition corresponds to the traditional ‘wisdom of the crowds’ context (6–8).

#### **1.2.4 Condition 2: Social Learning [Wisdom of Static Networks]**

In this condition participants make an independent guess, but then engage in a stage of active social learning, where they are exposed to their ego-network’s estimates in real time. In particular, in this condition each round consists of three stages (as illustrated in Figure S3):

1. **Initial stage:** participants make an initial independent guess.
2. **Interactive stage:** Participants observe the guesses of 3 of their randomly selected neighbors (fixed for the entire experiment) and can update their guess to accommodate the information and experiences, the opinions and judgments, the stubbornness and confidence, of the other players. Note that participants observe in real-time the estimates of the other participants that they are connected to and can update their estimates multiple times before they submit their final estimate.
3. **Round Outcome:** In the round outcome, the participant is shown the correct answer and how much they earned in the last round and cumulatively. participants do not see any performance information regarding their neighbors.

This context is analogous to that studied by work at the intersection of the ‘wisdom of crowds’ and social learning, such as (9–13).

#### **1.2.5 Condition 3: Selective Social Learning [Wisdom of Dynamic Networks]**

This condition builds on the static network condition, and adds the possibility for participants to choose and periodically reevaluate who to follow and be influenced by (see Figure S3). After each task round, participants are shown performance information of all peers in the group, and

are allowed to follow/unfollow a bounded number them. In this experiment, the maximum number of outgoing connections is three.

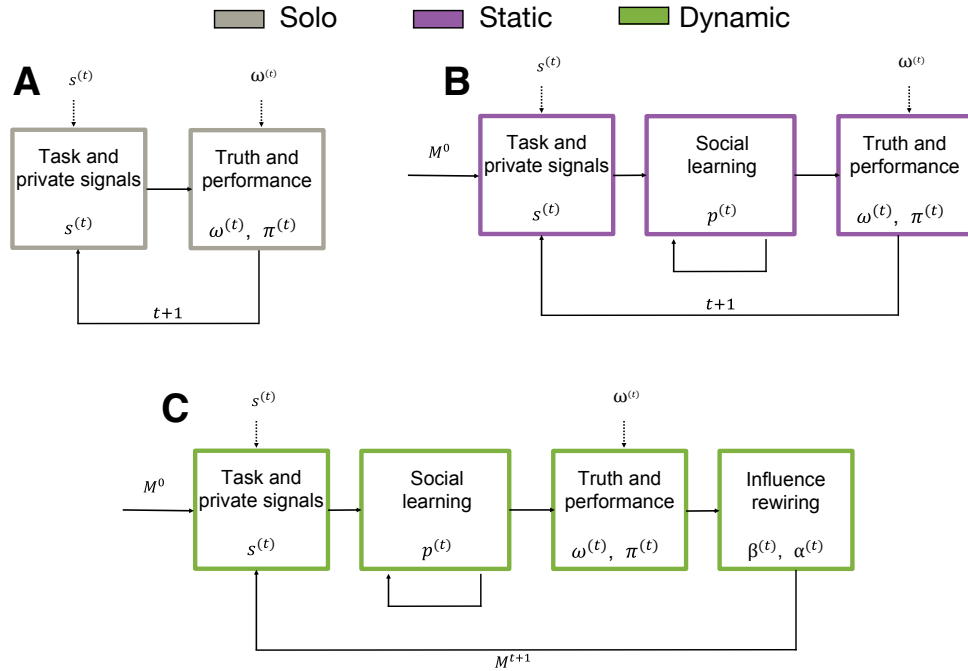


Figure S3: **Illustration of the experimental conditions.** Panel (A) depicts the Solo condition (i.e., no social information) where participants make independent estimates. This condition corresponds to the baseline wisdom of the crowd context. Panel (B) describes the Static network condition (i.e., social learning) where participants engage in a stage of interactive social learning, where they are exposed to the estimates of a fixed set of peers in real time. Panel (C) describes the Dynamic network (i.e., selective social learning) condition that adds the possibility for participants to choose who to follow and be influenced by in the next round.

## 2 Game Screenshots

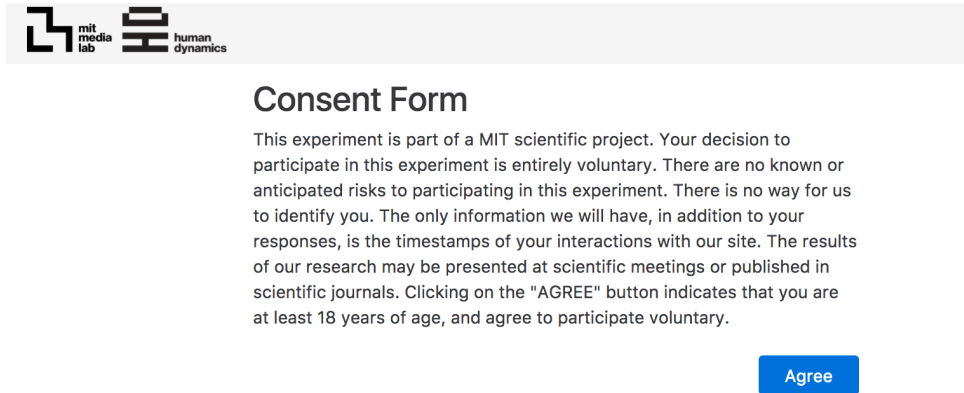


Figure S4: Illustration of the consent form presented to the participants. The study (Approval#: 1509172301) was reviewed and approved by the Committee on the Use of Humans as Experimental participants (COUHES) at MIT. All participants provided an explicit consent to participants in this study and COUHES approved the consent procedure.

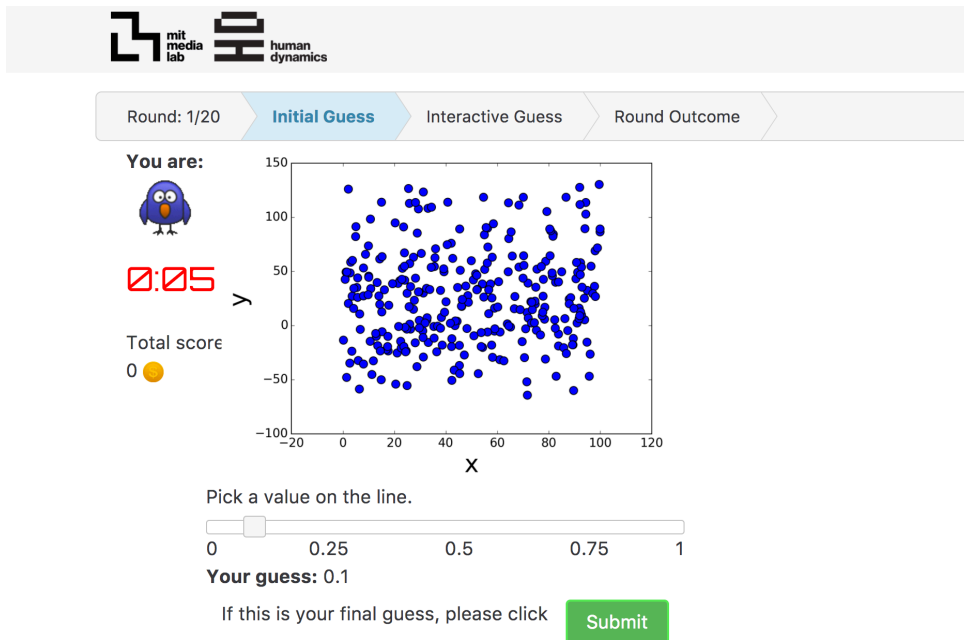


Figure S5: Participants in all conditions make independent guesses about the correlation of two variables independently.

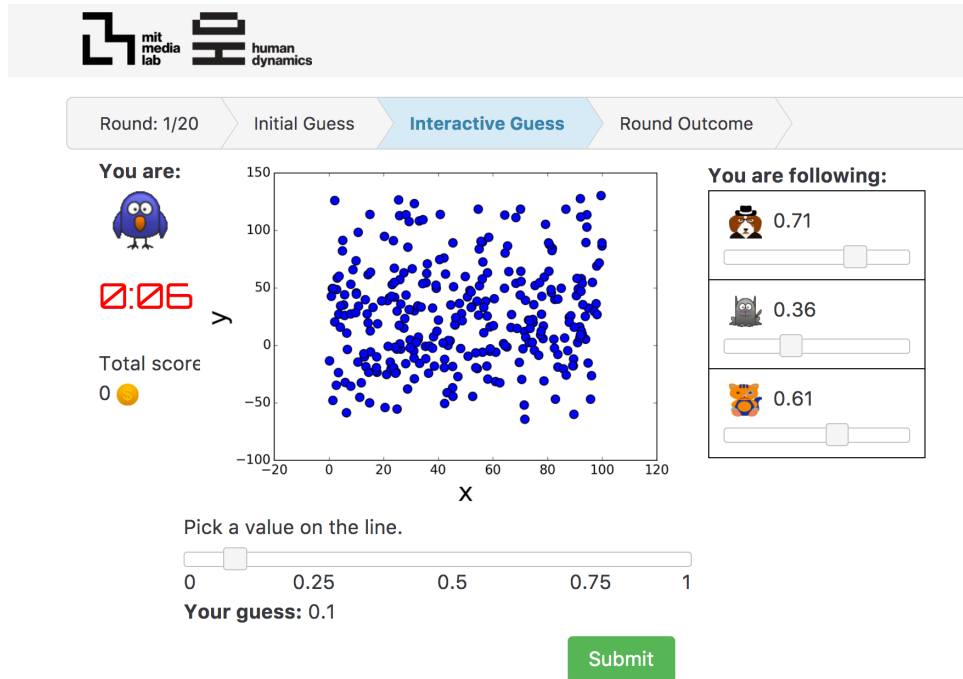


Figure S6: Participants in the Static and Dynamic network condition engage in an active social learning phase, where they are exposed to their ego-network's estimates in real time.

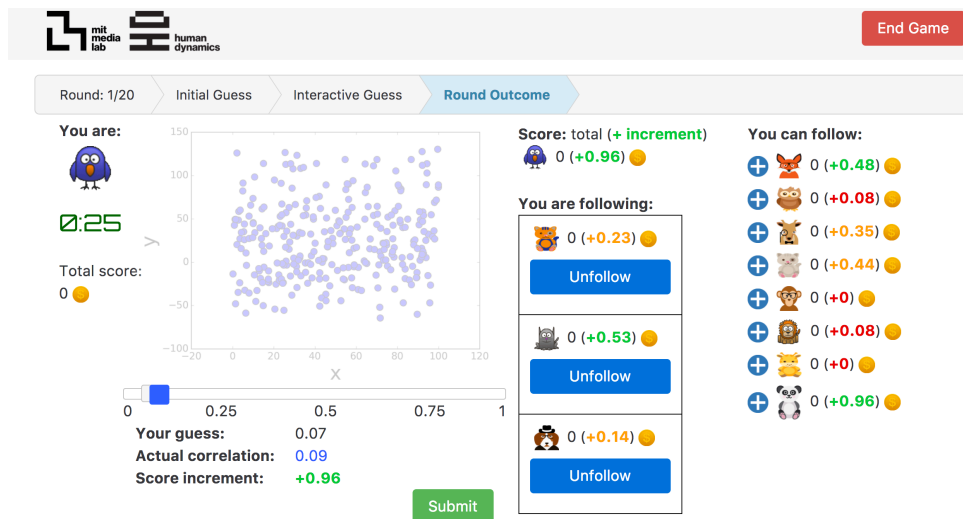


Figure S7: After each task round, participants in the Dynamic network conditions are shown performance information of all peers in the group, and are allowed to follow/unfollow a bounded number them. In this experiment, the maximum number of outgoing connections is three.

### 3 Individual-Level Accuracy

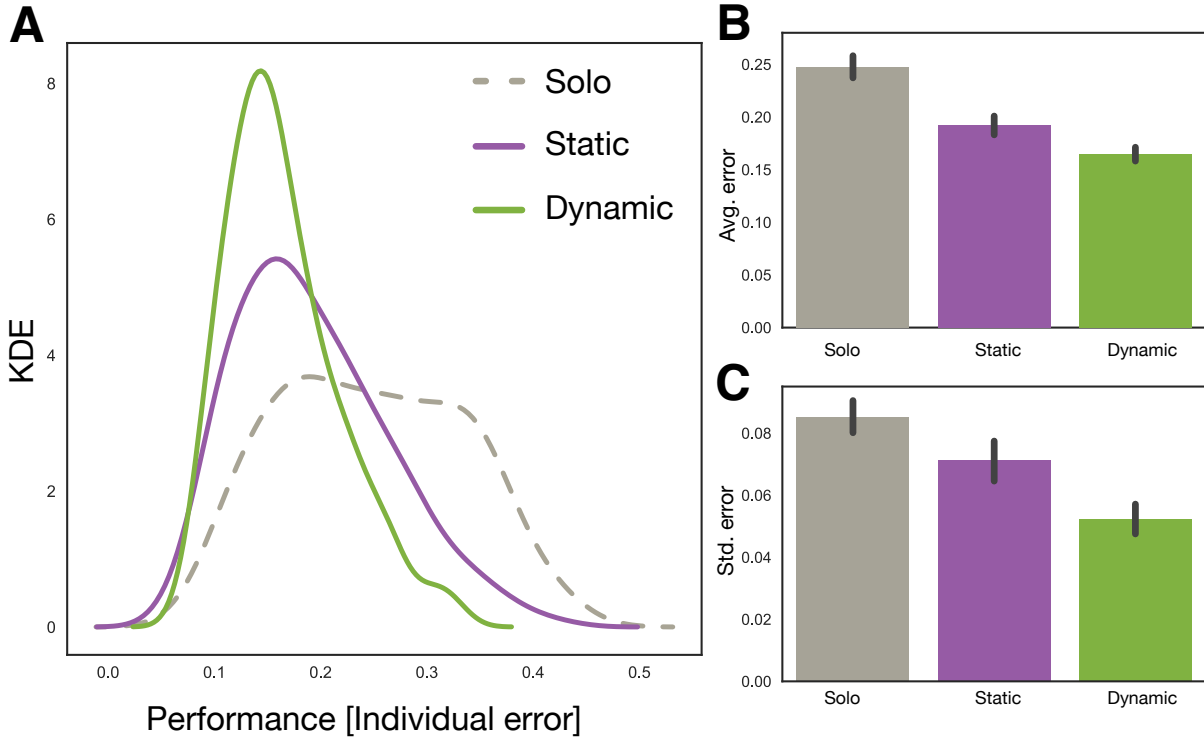


Figure S8: **Dynamic social influence benefits the performance of individuals in the crowd.** (A) Kernel Density Estimate (KDE) of participants' individual performance (i.e., average error across all rounds) for the three experimental conditions. We find that participants in groups connected by dynamic influence networks (Dynamic condition) achieved 38% reduction in average error compared to participants in unconnected groups (Solo condition), and 12% reduction in average error compared to participants in groups connected by static influence networks (Static condition). Panel (B) compares the average performance of individuals across conditions. Two-sample t-tests show a significant difference between the average individual error of participants in the Solo and Static conditions ( $P < 0.0001$ ), as well as between participants in the Static and Dynamic conditions ( $P < 0.001$ ). Panel (C) compares the standard deviation of participant's individual performance across conditions, and shows that individual performance in groups connected by dynamic influence networks was, not only better on average, but also substantially more equal on its distribution among group members.

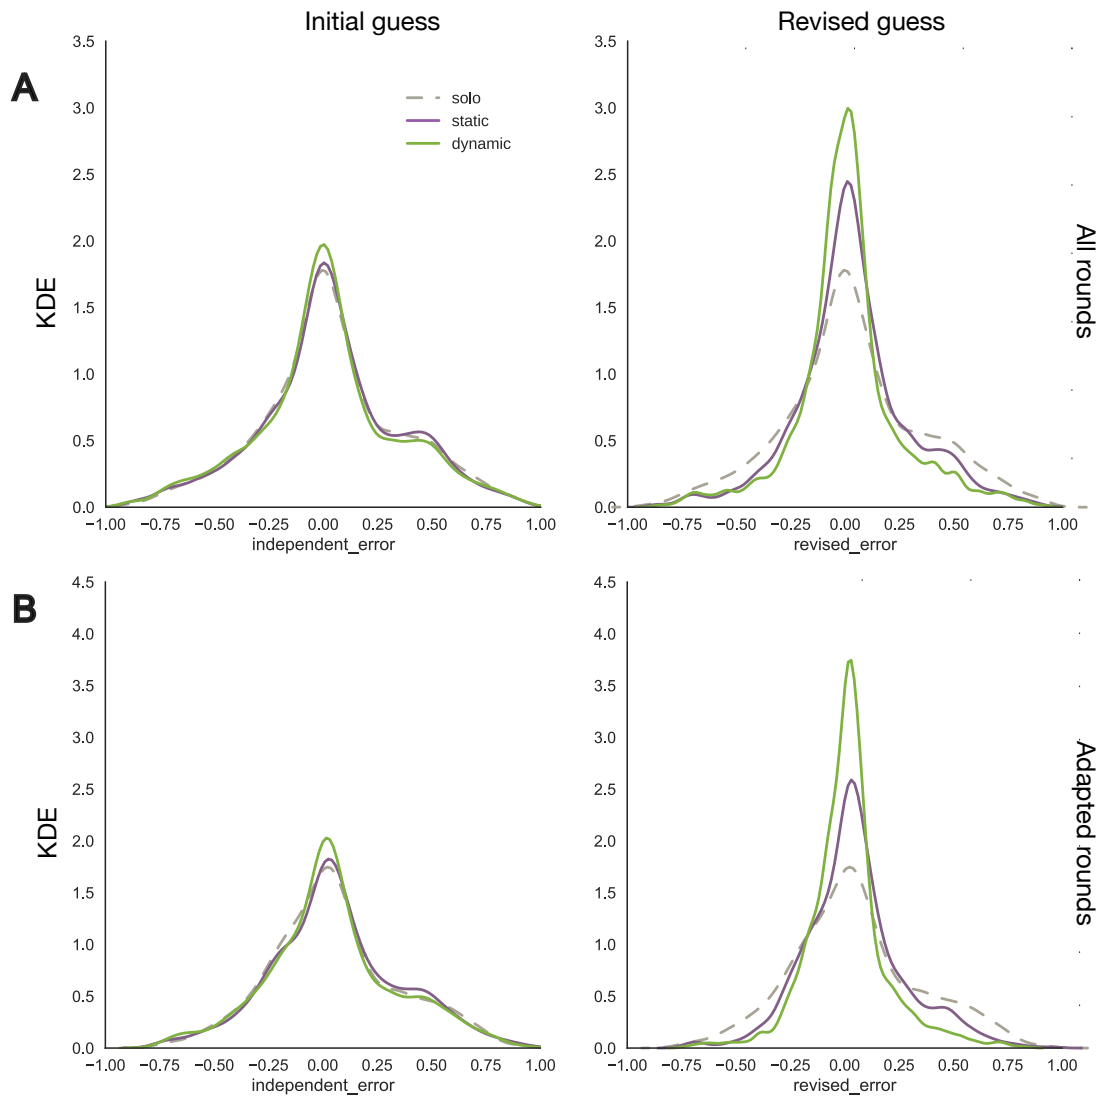


Figure S9: Panel (A) shows individual errors in the full game and Panel (B) shows the error in the adapted period (i.e., periods [6-10] and [16-20]). The error for the initial guess in both panels is the same across conditions, however, the dynamic network condition incurs much lower errors in the adapted periods (as in Panel B).

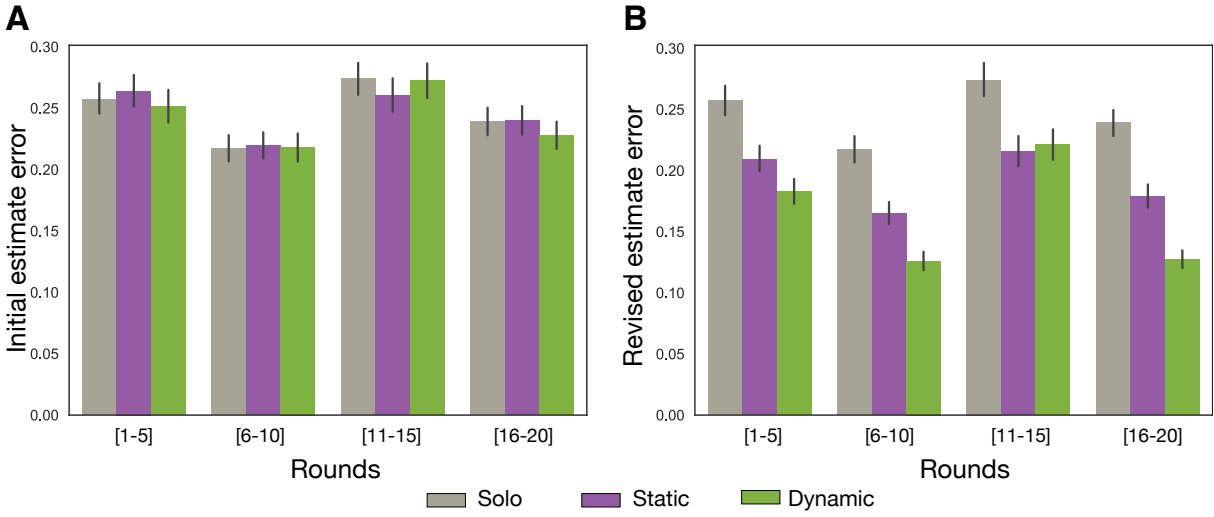


Figure S10: Panel (A) shows the errors before the interactive estimation phase (i.e., pre-social learning). Panel (B) shows the errors after the participants revised their estimates in the static and dynamic network conditions (i.e., post-social learning).

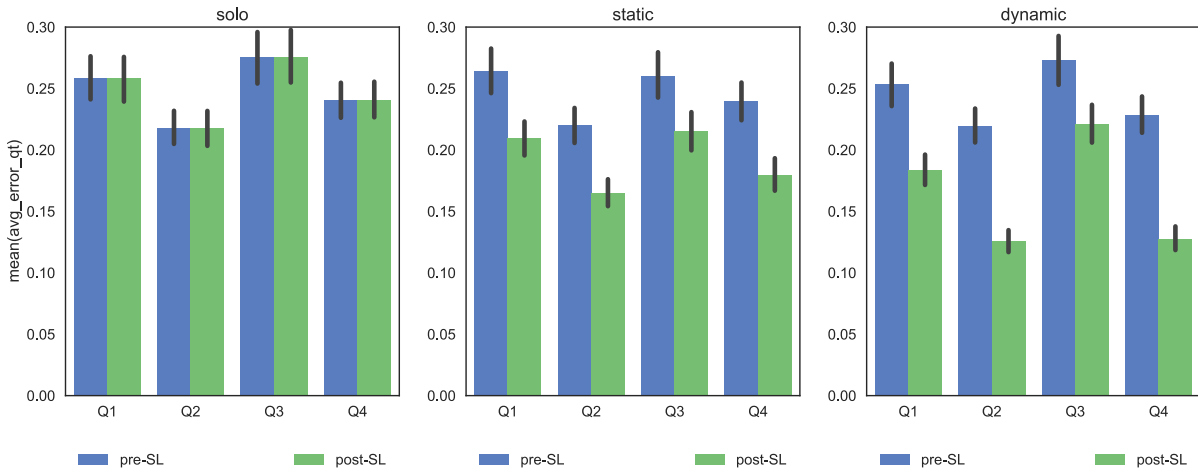


Figure S11: The distribution of pre-social learning and post-social learning for the three conditions.



# 4 Collective Accuracy

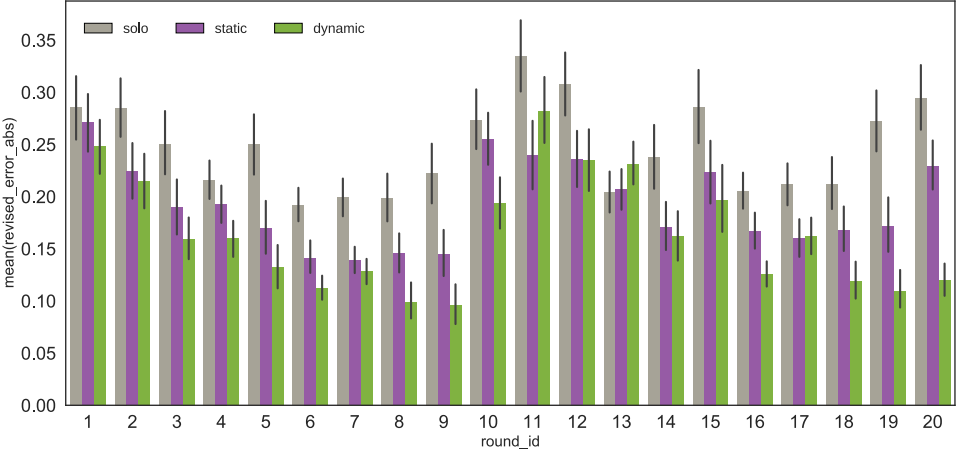


Figure S12: **Dynamic social influence effect in individual rounds: Adaptive with time and reduces individual error.** All error rates are post-social learning errors.

## 5 Network Analysis

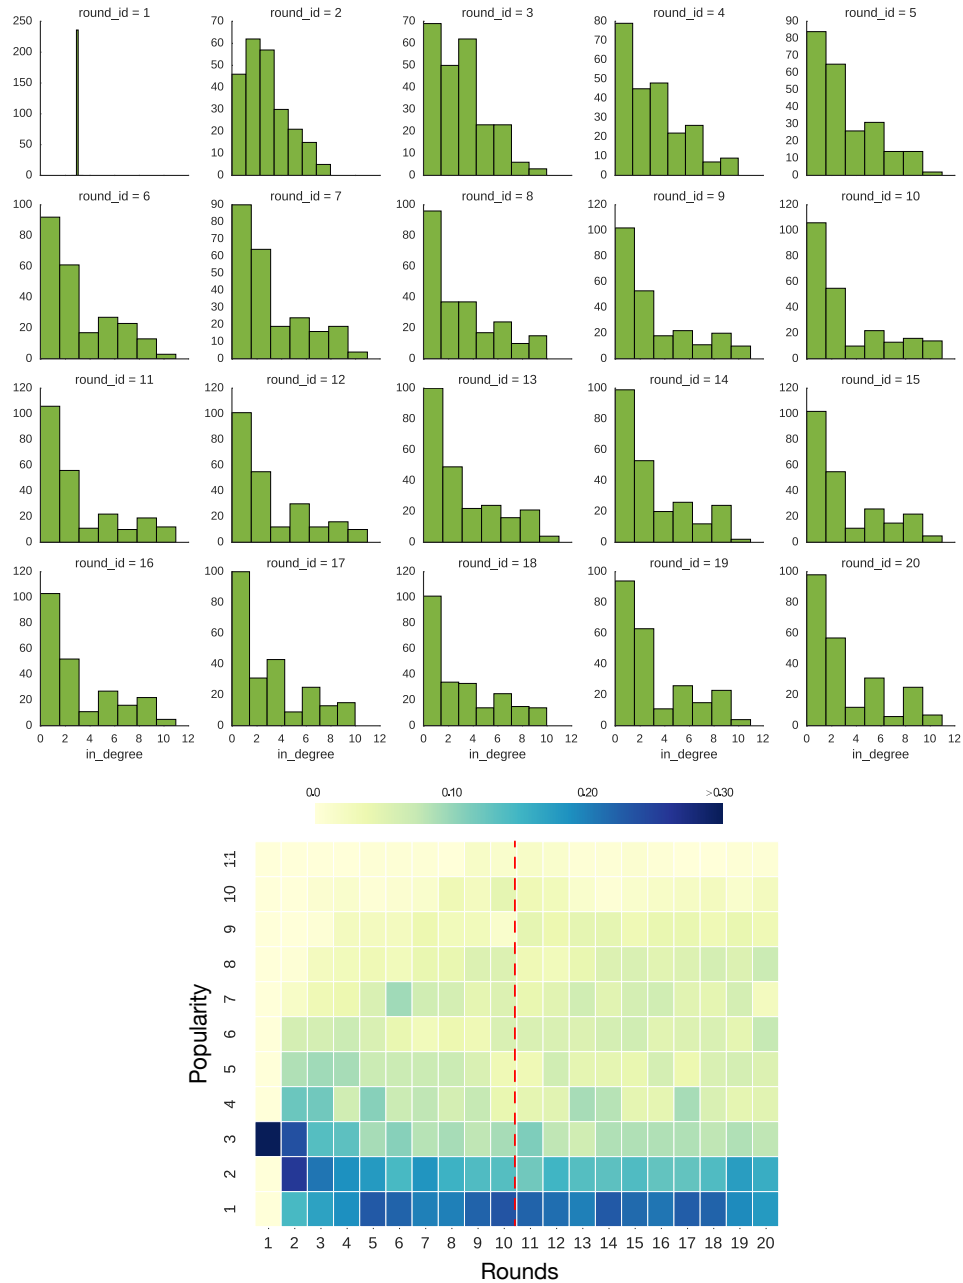


Figure S13: The in-degree distribution at each round (dynamic condition). We see that starting from a fully distributed network (everyone having the same in-degree) the network becomes more centralized over the course of the experiment.

Network (global) Centralization measure: We measure network centralization by calculating the sum in differences in centrality between the most central node in the network and all other nodes; and then divide this quantity by the theoretically largest such sum of differences in a star network of the same size

$$C_x = \frac{\sum_{i=1}^N C_x(p_*) - C_x(p_i)}{\max \sum_{i=1}^N C_x(p_*) - C_x(p_i)}$$

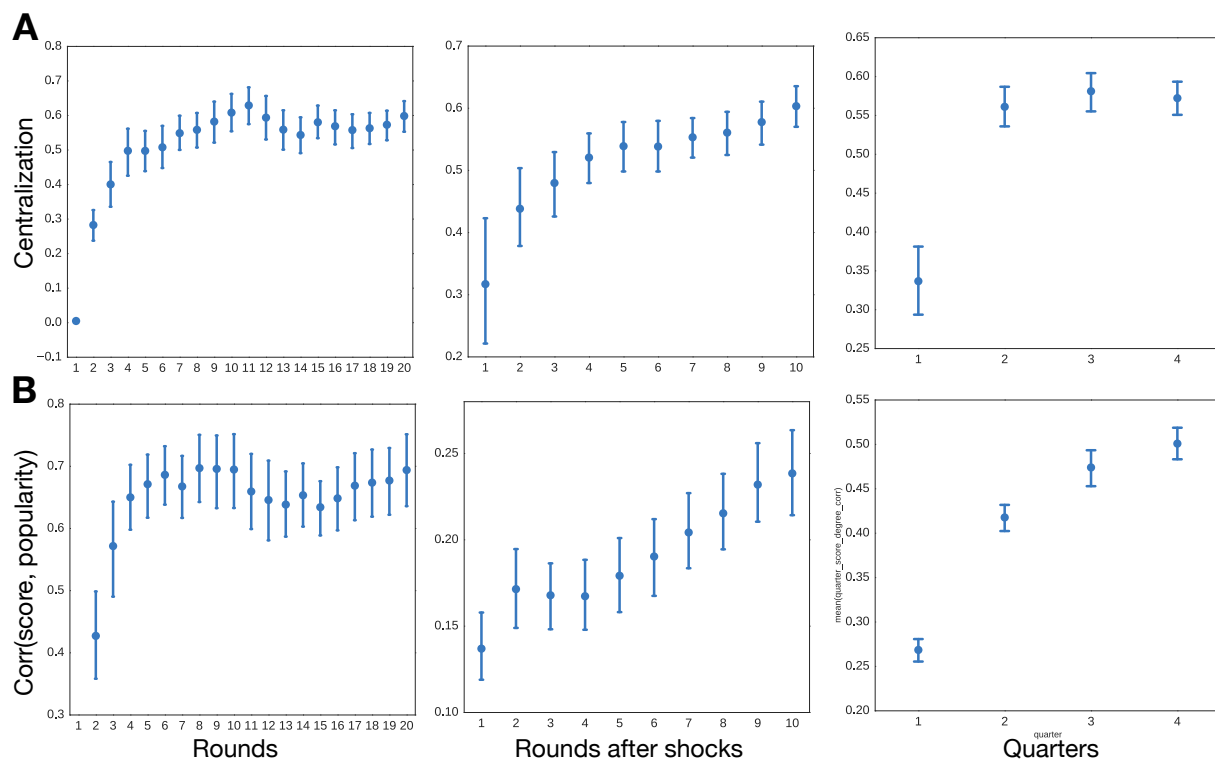


Figure S14: (A) The network becomes more centralized with time, and (B) the correlation between the node performance and popularity increases as the game progresses.

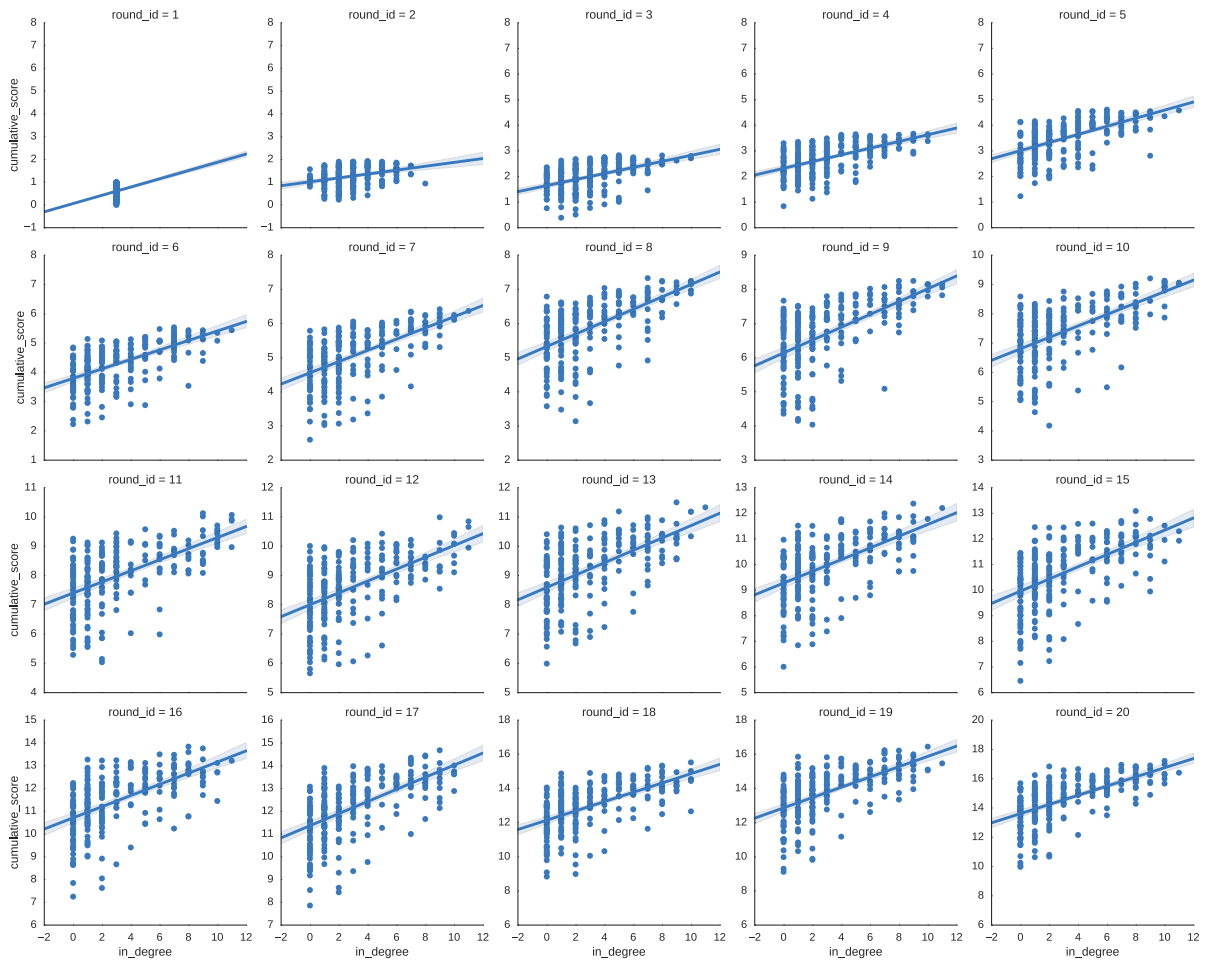


Figure S15: The relationship between the number of followers (i.e., popularity) and the player’s cumulative score (i.e., performance).

## 6 Confidence and Accuracy Correlation

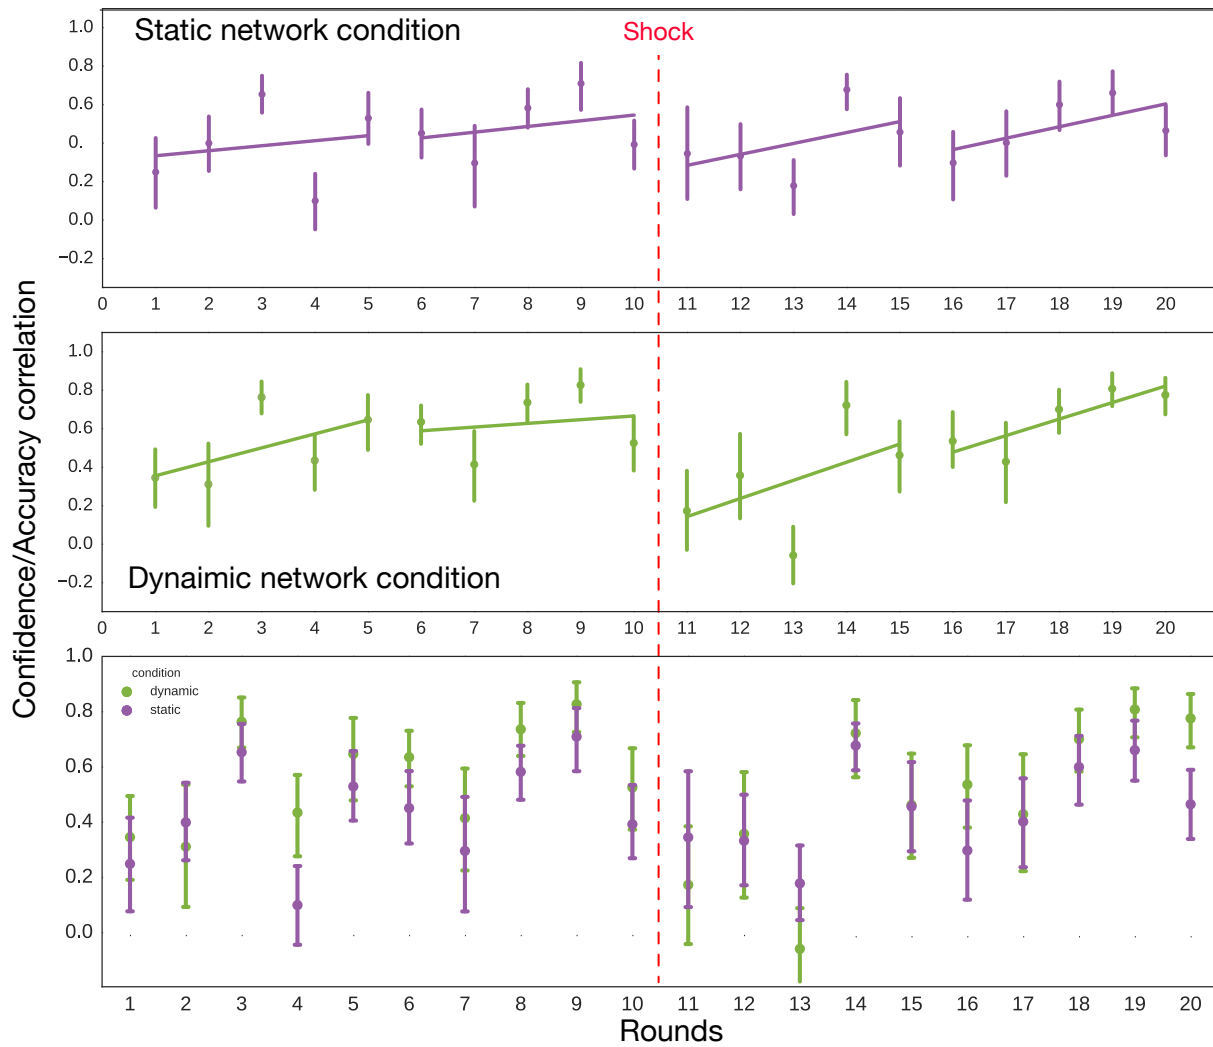


Figure S16: The relationship between confidence (i.e., resistance to social influence) and accuracy of initial estimates for the static and dynamic conditions. Error bars in the dynamic/static comparison (the bottom figure) are standard errors from the mean. All other errors are 95% confidence intervals.

## 6.1 Summary of Results and Main Findings

- Groups connected by dynamic influence networks (dynamic network condition) improve participants' individual performance over unconnected groups (solo condition), and groups connected by static influence networks. Dynamic networks achieve better performance, and more equally distributed. Figure S8:
  - Participants in the Dynamic network achieved 38% reduction in average error compared to participants in the Solo condition.
  - Participants in the Dynamic network achieved 12% reduction in average error compared to participants in the Static Network condition.
  - Both differences are statistically significant at the 95% confidence levels while accounting for within group dependencies.
  - Not only better individual performance on average, but also substantially **more equal** on its distribution among group members.
- Dynamic influence networks improve groups' collective estimates (Wisdom of Dynamic Networks) over the other conditions (Wisdom of Static Networks & Wisdom of Crowds).
  - Collective estimates of the dynamic network condition incur substantially lower errors in periods where the network has adapted ([6-10] and [16-20]):
    - \* 54% lower than unconnected groups.
    - \* 33% lower than groups connected by static influence networks
    - \* 28% lower than the errors of the best performer in the solo condition.
  - Even the top-performer in each group gets substantial benefits by joining a group connected by dynamic networks, rather than working in isolation.

- **Dynamics of the Network:** The adaptive mechanism in the dynamic network exploits the higher performer individuals by centralizing the network around them. The popular and high performer individuals before the shock are not the same as the ones after the shock.
- **Dynamics on the Network:** Another adaptive mechanism is the correlation between accuracy of initial estimate and resistance to social influence. This allows even the best individual to perform better (i.e., copy others when in doubt). This mechanism plays a role in the static network condition S16.
- **Mean-variance tradeoff.** We study top-k estimates across conditions, which generalize the traditional notion of group wisdom estimates, containing both group wisdom estimates and best individual's estimates as special cases (top-12 and top-1 respectively; see Figure 5 in the main text of this paper).
  - Results show that, as we remove low-performing individuals (decreasing  $k$  from  $k = 12$ ), top-k estimates initially improve in both mean and standard deviation; then, as we further curate the crowd, top-k estimates tradeoff between decreasing mean error and increasing its standard deviation; finally, top-k estimates disimprove in both objectives as  $k \rightarrow 1$  and top-k approaches the estimates of the single best individual.
  - Comparison of top-k curves across conditions show that, for any  $k \in [1, 12]$ , dynamic influence networks improved estimation errors in terms of both mean and standard deviation over estimation errors of groups in the unconnected and static network conditions.
- These results open the floor for **discussion** on:
  - Adaptive systems as paradigm for collective intelligence systems: distributed and

non-stationary information environments, learning rates, availability of performance information and social cues, etc.

- Fundamental difference and complementarity between aggregation strategies and intervening the social system itself, towards improving collective intelligence systems.



## 7 Numerical Model & Simulation

Analysis of the dynamic properties of Figure S3's framework requires models of its human components, i.e., the social learning and network rewiring heuristics. We model the former as a DeGroot process (14), and propose a performance-based preferential detachment and attachment model for the latter. This section describes such models and agents' input private signals.

**Notation.** Let  $N = \{1, 2, \dots, n\}$  represent a group of agents that participate in a sequence of tasks, indexed by discrete time  $t$ . Let  $G(N, E^{(t)})$  be a sequence of directed graphs representing the influence network at each period  $t$ . Let  $e_{ij}^{(t)} \in [0, 1]$  denote the edge weight of  $(i, j)$  at time  $t$ , and  $M^{(t)}$  the row-normalized stochastic matrix associated with  $E^{(t)}$ , i.e.,  $M_{ij}^{(t)} = \frac{e_{ij}^{(t)}}{\sum_{h \in N} e_{ih}^{(t)}}$ .

Agents receive private signals  $s_i^{(t)} \in [0, 1]$ , for  $i \in N$ , regarding the true state of the world  $\omega^{(t)} \in [0, 1]$ . Similarly, we denote agents' post-social learning beliefs by  $p_i^{(t)} \in [0, 1]$ , for  $i \in N$ .

**Private Signals.** We depart from the commonly made assumption of collective unbiasedness of agents' private signals (7, 9, 11, 12, 15), allowing agents' signals to be distributed with arbitrary means and skewness. Let  $\mu_i = E[s_i]$  denote the mean of agent  $i$ 's signal, and  $\bar{\mu} = \frac{1}{n} \sum_i \mu_i$  be the collective mean of private signals; we are interested on the more general setting of information environments where  $\bar{\mu} \neq \omega$ , i.e., where the collective distribution of initial signals is not centered on the truth. Figure S17 illustrates the difference between unbiased and biased information environments.

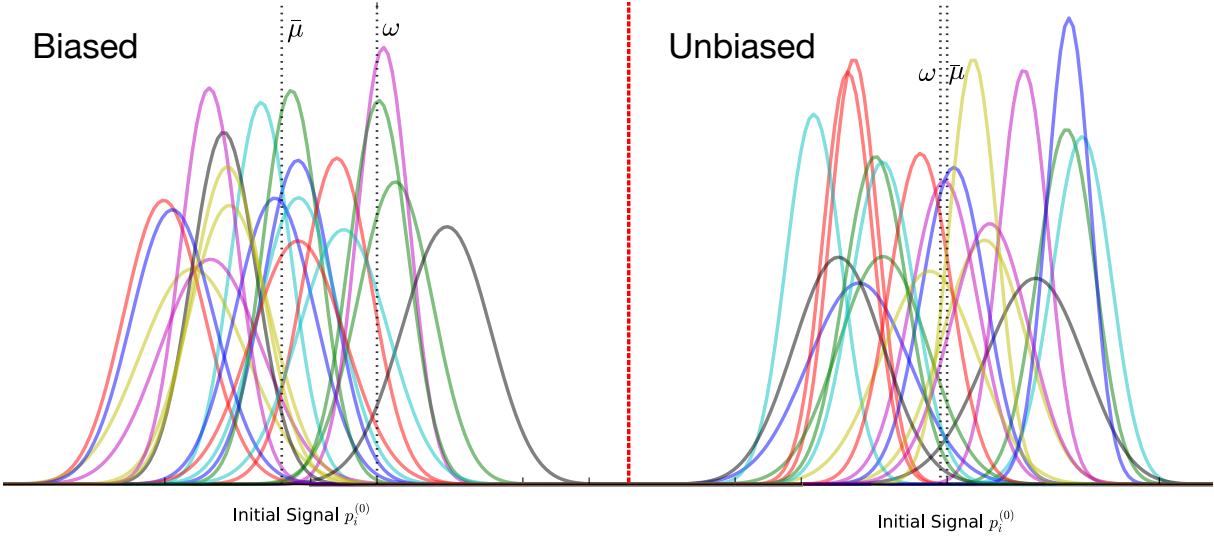


Figure S17: Traditional accounts of ‘wisdom of crowds’ phenomena assume unbiased and statistically independent signals among agents. In our model, we assume arbitrary (potentially biased) initial signals.

## 7.1 Social Learning Process

Social learning is modeled as a DeGroot process (14), where each agent updates her belief by taking weighted averages of her own belief (i.e., private signal) and the beliefs of neighboring agents. DeGroot averaging as social learning heuristic has been well studied empirically and theoretically (11, 12), and shown to robustly describe real-world belief updating better than more optimal rational Bayesian models (16). In particular, we model post-social learning beliefs as the result of a two-stage DeGroot process on private signals, given by

$$p^{(t)} = (M^{(t)})^2 s^{(t)} \quad (1)$$

## 7.2 Individual and collective performance

**Individual performance** is evaluated based on the errors of post-social influence estimates.

Individual cumulative error is defined by:

$$\epsilon_i^{(t)} = \frac{1}{\lambda + 1} \sum_{r \in [0, \lambda]} \left| p_i^{(t-r)} - \omega^{(t-r)} \right|,$$

where  $\lambda$  controls the number of retrospective periods that performance information is averaged across.

Agents assess performance of other agents relative to the performance of the best agent in the group. We define relative error of agent  $i$  as  $\pi_i^{(t)} = \epsilon_i^{(t)} - \epsilon_{min}^{(t)}$ , and denote the set of performance information available to agent  $i$  at time  $t$  by vector  $\Pi_i^{(t)} \in [0, 1]^n$  with elements

$$\pi_{ij}^{(t)} = \begin{cases} \pi_j^{(t)} & \text{for } j \neq i \\ \pi_{s_i}^{(t)} & \text{for } j = i \end{cases}$$

where  $\pi_{s_i}^{(t)}$  is the relative error of agent  $i$ 's private signal.

**Collective error.** We are interested on the wisdom of the dynamic network (WDN) error,  $\epsilon_{wdn}$ , which captures collective error after selective social learning according to the interaction network. We compare  $\epsilon_{wdn}$  against the wisdom of the crowd (WC) baseline,  $\epsilon_{wc}$ , which captures collective error of the simple averaging of agents' initial signals.

$$\epsilon_{wdn}^{(t)} = \left| \omega^{(t)} - \frac{1}{n} \sum_i p_i^{(t)} \right| \quad (2)$$

$$\epsilon_{wc}^{(t)} = \left| \omega^{(t)} - \frac{1}{n} \sum_i s_i^{(t)} \right| \quad (3)$$

## 7.3 Influence Rewiring Process

Individuals connect by weighted influence links that are revised over time. We model influence rewiring heuristics that strengthen links when a neighbor exhibits high performance and weaken

or break links when a neighbor performs poorly. Agents can distribute attention among a limited number of peers, captured by parameter  $\kappa$ , which represents cognitive or infrastructure constraints (e.g., limits on our ability to keep track of social information and relations (17)). In particular, agents dynamically allocate  $\kappa \in \mathbb{N}$  shares of their attention to other agents. Let  $e_{ijk} \in \{0, 1\}$ , for  $k \in \{1, 2, \dots, \kappa\}$ , indicate that  $i$  places attention share  $k$  on  $j$ , then  $e_{ij} = \sum_k e_{ijk}$  and  $e_{ij} \in \{0, 1, 2, \dots, \kappa\}$ .

**Probability of detachment.** Probability that agent  $i$  detaches from  $j$  is a positive function of  $i$  and  $j$ 's errors, and given by equation 4. For example, if  $i$ 's error is among the lowest of the group ( $\pi_i^{(t)} \approx 0$ ),  $i$  is unlikely to rewire her local network. Conversely, if  $i$ 's error is significant (e.g.,  $\pi_i^{(t)} \approx 1$ ),  $i$  detaches from  $j$  with probability proportional to  $j$ 's error.

$$\beta_{ij}^{(t)} = \left( \pi_i^{(t)} \pi_{ij}^{(t)} \right)^{\frac{1}{2}} \quad (4)$$

**Probability of Attachment.** High-performing agents are more likely to be followed. Analogous to generalized preferential attachment (18), probability that agent  $i$  attaches to  $j$  is inversely proportional to  $j$ 's error, and given by

$$\alpha_{ij}^{(t)} = \left( \frac{1 - \pi_{ij}^{(t)}}{n - \sum_j \pi_{ij}^{(t)}} \right)^2 c \quad (5)$$

where  $c$  is a normalization constant.

**Network Evolution.** Define i.i.d. random variables  $b_{ijk}^{(t)} \sim \text{Bernoulli}(\beta_{ij}^{(t)}) \quad \forall (i, j, k)$ , then random variables  $b_{ij}^{(t)} = \sum_k b_{ijk}^{(t)} e_{ijk}^{(t)} \sim \text{Binomial}(e_{ij}^{(t)}, \beta_{ij}^{(t)})$  indicate the amount of attention shares that  $i$  detaches from  $j$  in period  $t$ . Define  $n$ -dimensional random vectors  $a_i^{(t)} \sim \text{Multinomial}(\sum_j b_{ij}^{(t)}, \alpha_i^{(t)})$ , where  $\alpha_i^{(t)}$  is  $i$ 's vector of attachment probabilities. Elements  $a_{ij}^{(t)} \in \{0, 1, \dots, \kappa\}$  indicate the amount of shares that  $i$  attaches to  $j$  in period  $t$ , and network evolution is given by

$$e_{ij}^{(t+1)} = e_{ij}^{(t)} - b_{ij}^{(t)} + a_{ij}^{(t)} \quad \forall i, j \quad (6)$$

## 7.4 Results

### 7.4.1 Collective Error and Adaptability

Using the above-described model, we performed Monte Carlo simulations of a group of twenty agents who participate in a sequence of estimation tasks, where agents can follow and be influenced by a maximum number of five peers ( $\kappa = 5$ ).

Figure S18A shows the evolution of dynamic networks' estimates along the sequence of tasks, and compares them against simple wisdom of the crowd estimates. Starting at  $t = 0$ , WC and WDN estimates incur the same average error, equal to the bias of initial signals against the true state of the world:  $E[\epsilon_{wc}^{(0)}] = E[\epsilon_{wdn}^{(0)}] = |\bar{\mu} - w^{(0)}|$ . As agents observe each others' past performance, and rewire local influence networks accordingly, those with better information gain influence in the network, decreasing  $\epsilon_{wdn}^{(t)}$  and stabilizing it substantially below  $\epsilon_{wc}^{(t)}$ . Difference in means tests for each  $t$  show that  $E[\epsilon_{wdn}^{(t)}] < E[\epsilon_{wc}^{(t)}]$  for all  $t > 3$  with a 95% confidence level.

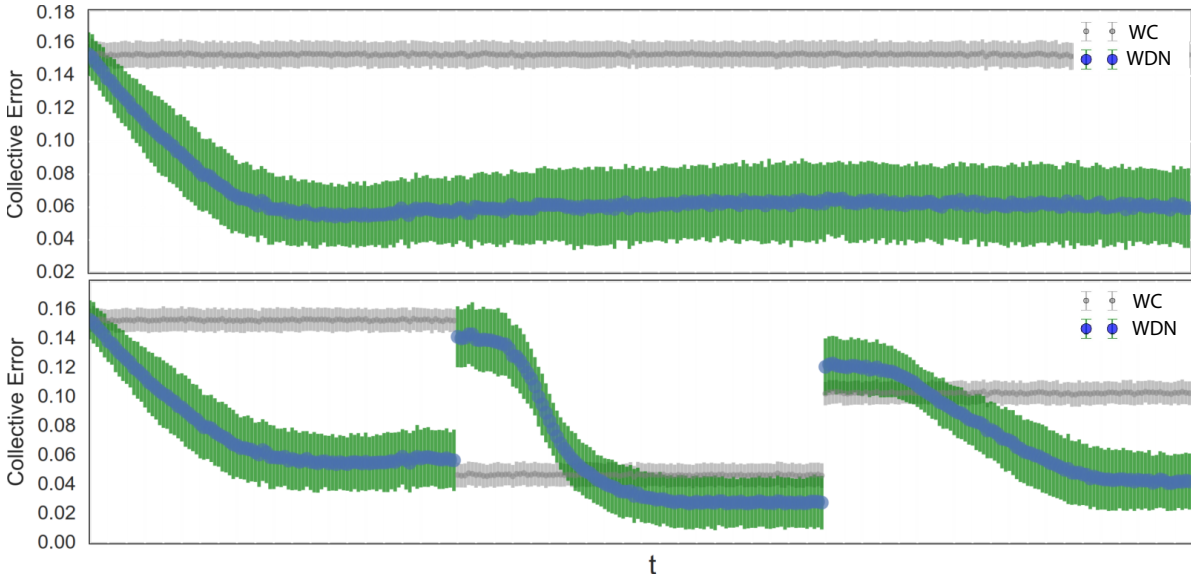


Figure S18: Evolution of collective error: wisdom of the crowd (WC) and wisdom of the dynamic network (WDN). Panel A) stationary distribution of information among agents. Panel B) non-stationary information environment, shocks to the information distribution introduced at  $t = \{100, 200\}$

We study as well how the dynamic agent network responds to environmental shocks. Figure S18B compares the dynamics of collective intelligence for WC and WDN under a non-stationary environment, where shocks to the joint distribution of private signals  $p_s^{(t)}$  and truth  $\omega_s^{(t)}$  are introduced at  $t = \{100, 200\}$ . The dynamic network adapts to post-shock distributions by shifting influence to agents with better information in the post-shock environment, driving collective error  $\epsilon_{wdn}^{(t)}$  significantly below  $\epsilon_{wc}^{(t)}$ . Difference in means tests showed that  $E[\epsilon_{wdn}^{(t)}] < E[\epsilon_{wc}^{(t)}]$  for all  $t \in \{[3, 99] \cup [140, 199] \cup [225, 300]\}$  with a 95% confidence level.

### 7.4.2 Feedback Quality and Performance

This section focuses on the role of feedback quality on the performance of the collective. This is a realistic assumption as usually the environmental cues about performance can be noisy in many cases. This leads to increasing the chance of picking the wrong individual to follow. Indeed, we see in Figure S19 that as we increase the feedback noise level, the collective performance degrades until it converges to the performance of the independent crowd.

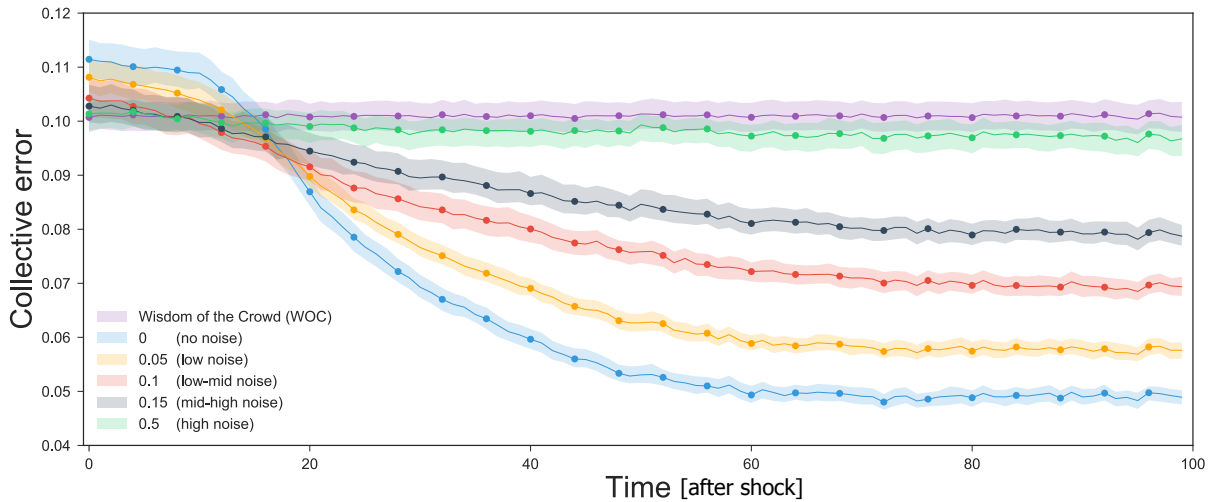


Figure S19

### 7.4.3 Rates of Learning and Environment Change

This section focuses on how the retrospective length of cumulative performance information, parameterized by  $\lambda$ , affects the learning and adaptation rate of the network. Moreover, we study the interaction effects that the network learning rate and the environment change rate have on collective error.

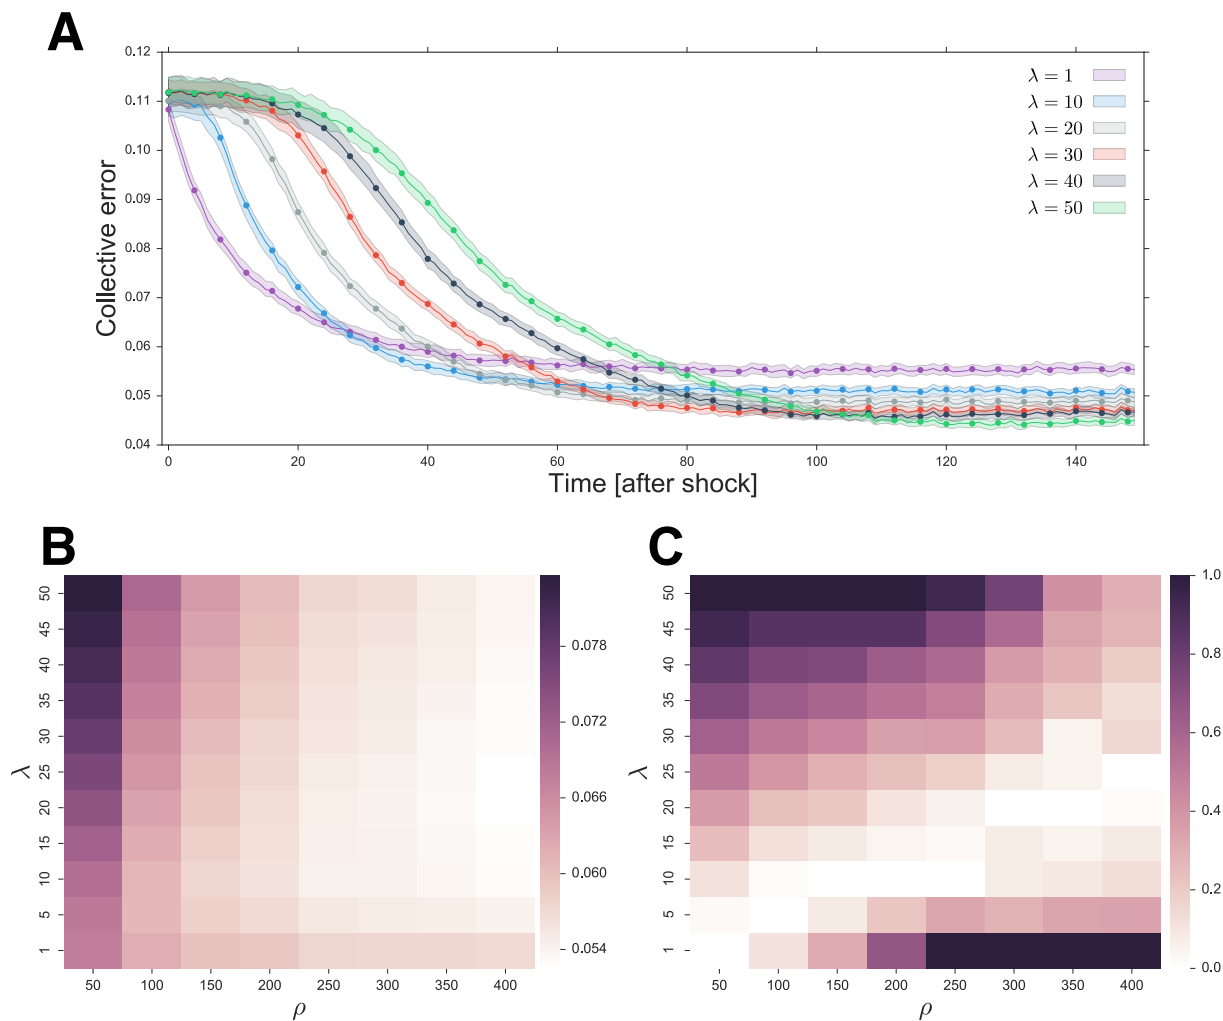


Figure S20: Panel (A): Learning rates associated to different  $\lambda$ 's, where colored bands show 95% confidence intervals. Panel (B): Effects of  $\lambda$  and  $\rho$  on collective error, where shades of orange indicate time-averaged collective error. Panel (C): Effects of  $\lambda$  and  $\rho$  on collective error, normalized per type of information environment ( $\rho$  column).

Figure S20A shows the after-shock evolution of collective errors for a set of  $\lambda$ 's. On average,

small  $\lambda$ 's adapt significantly faster to new information environments, but eventually stabilize at higher collective error levels compared to larger  $\lambda$ 's. This short-term versus long-term accuracy tradeoff associated to the network learning rate implies that the optimal choice of  $\lambda$  depends on the frequency of environment change.

We study such interaction by simulating non-stationary environments where random shocks to the distribution of information among agents arrive according to a Bernoulli process with average interarrival time  $\rho \in \mathbb{N}$ . Figure S20B shows the time-averaged collective error of the network for a grid of  $(\rho, \lambda)$  combinations. Figure S20C shows collective errors for the same parameter grid, but normalized per type of information environment ( $\rho$  column).

We observe two main effects. First, for any given  $\lambda$ , the less frequent the environment changes the better accuracy the network attains. Second, for any particular  $\rho$ , average collective error is a convex function of  $\lambda$ , where optimal  $\lambda$ 's balance short and long term accuracy in accordance with the rate of environment change. Intuitively, networks facing stationary environments are best off accumulating performance information over the entire available history. However, in the context of non-stationary environments with frequent shocks, networks are better off restricting attention to performance information in the near past; while, similar to stationary environments, networks facing moderately stationary environments benefit from accumulating performance over longer periods of time.

## References

1. A. A. Arechar, S. Gächter, L. Molleman, *Experimental Economics* pp. 1–33 (2017).
2. T. Chesney, S.-H. Chuah, R. Hoffmann, *Journal of Economic Behavior & Organization* **72**, 618 (2009).
3. J. Hergueux, N. Jacquemet, *Experimental Economics* **18**, 251 (2015).



4. J. J. Horton, D. G. Rand, R. J. Zeckhauser, *Experimental economics* **14**, 399 (2011).
5. A. J. Berinsky, G. A. Huber, G. S. Lenz, *Political Analysis* **20**, 351 (2012).
6. F. Galton, *Nature* **75**, 450 (1907).
7. J. Surowiecki, *The wisdom of crowds* (Anchor, 2005).
8. D. Prelec, S. Seung, J. McCoy, *Nature* (2017).
9. J. Lorenz, H. Rauhut, F. Schweitzer, D. Helbing, *Proceedings of the National Academy of Sciences* **108**, 9020 (2011).
10. G. Madirolas, G. G. de Polavieja, *PLoS Comput Biol* **11**, e1004594 (2015).
11. B. Golub, M. O. Jackson, *American Economic Journal: Microeconomics* **2**, 112 (2010).
12. P. M. DeMarzo, J. Zwiebel, D. Vayanos, *Social Influence, and Uni-Dimensional Opinions (November 2001). MIT Sloan Working Paper* (2001).
13. W. Mason, D. J. Watts, *Proceedings of the National Academy of Sciences* **109**, 764 (2012).
14. M. H. DeGroot, *Journal of the American Statistical Association* **69**, 118 (1974).
15. D. Acemoglu, A. Nedic, A. Ozdaglar, *Decision and Control, 2008. CDC 2008. 47th IEEE Conference on (IEEE, 2008)*, pp. 1714–1720.
16. A. G. Chandrasekhar, H. Larreguy, J. P. Xandri, *Work. Pap., Mass. Inst. Technol., Cambridge, MA* (2012).
17. R. Dunbar, *brain* **9**, 178 (1998).
18. A.-L. Barabási, R. Albert, *science* **286**, 509 (1999).