

Local Signals, Global Stakes: Groups with Aligned Interests Struggle to Coordinate when Information is Noisy and Shared Locally

Short title: Coordination under local and noisy information

Isabelle Brocas *

Juan D. Carrillo †

May 2025

Abstract

We perform a laboratory experiment to study the challenges of collective action in groups that share goals but are subject to local and noisy information. We show that consensus on the correct action declines in larger networks with fewer local connections and an asymmetric distribution of signals, where misinformation forms small stubborn pockets of incorrect beliefs. Individual characteristics also affect success: undergraduate students—who actively explore different alternatives and are more willing to accept that their choice is potentially wrong—achieve coordination more effectively than their less adventurous and less malleable high school peers. Our findings explain persistent societal misconceptions such as vaccine hesitancy or climate denial and highlight the importance of designing better-informed communication networks to improve collective decision-making in a variety of contexts, including committee deliberations, team projects and public goods. Educational interventions that encourage exploration may also help communities better navigate complex decisions and achieve successful cooperation.

Keywords: social learning, collective action, social network, consensus.

Teaser: Exploration and adaptability help overcome stubborn misinformation.

*Department of Economics, University of Southern California, Los Angeles, CA, USA. Institute for Advanced Study in Toulouse, Toulouse, France. Center for Economic Policy Research, London, United Kingdom. ORCID ID: 0000-0001-5847-1428. Email: brocas@usc.edu. Both authors contributed equally all aspects of the paper.

†Department of Economics, University of Southern California, Los Angeles, CA, USA. Institute for Advanced Study in Toulouse, Toulouse, France. Center for Economic Policy Research, London, United Kingdom. ORCID ID: 0000-0001-8616-1259. Email: juandc@usc.edu

1 Introduction

Consensus building is typically challenging because individuals frequently have different objectives or diverging views on how best to achieve shared objectives. Classic examples include government intervention vs. laissez-faire, gun control vs. gun rights, and drug enforcement vs. legalization. Yet, disagreement over objectives is not the only barrier to collective action. Even when individuals share common interests, coordination can fail due to informational frictions. In this paper, we argue that multi-person coordination among individuals with common interests may fail *if information is noisy and exchanged only locally*. This may explain why substantial portions of the population hold inaccurate beliefs, such as the idea that cell phone radiation significantly increases cancer risk, climate change is not real, or vaccines cause autism.¹

To test this hypothesis, we design a controlled laboratory experiment in a social network setting [12, 14, 15], using populations of high school and undergraduate students. Our experiment combines and extends insights from the literatures on *collective action* and *social learning*. The study of collective action has been central in the social computer science literature. The pioneering experimental work by Kearns and colleagues demonstrates that individuals, while observing the actions of only a limited set of neighbors, can solve computationally complex global problems, such as graph coloring [17] and consensus-building [16], among other tasks. In parallel, social learning has been extensively studied in economics. It refers to the process by which individuals, each possessing private but imperfect information about the state of the world, use the observed behavior of their peers to infer additional information and adjust their own decisions accordingly [2]. Research has shown that informational cascades — where individuals rationally ignore their private signal to follow the observed choices of others — occur naturally [1, 4], but also break easily [9]. More recently, social learning has been studied in network environments, where the structure of connections plays a critical role in determining how information spreads and aggregates [6, 11]).

Unlike the collective action literature, we study a setting with a correct state (e.g., the polio vaccine does not cause autism), where each individual receives a private signal. While the majority of participants receive correct information (e.g., through medical professionals), some are exposed to misinformation (e.g., through social media or pseudo-science).

¹Despite extensive scientific evidence to the contrary, two surveys by the Annenberg Public Policy Center (April 2024) and Gallup (July 2024) found that 24% and 13% of U.S. adults believed the MMR vaccine causes autism (<https://vaccinateyourfamily.org/wp-content/uploads/2025/03/The-Truth-About-Autism-and-Vaccines-VYF-ASF-Guide-2025-Color.pdf> and <https://news.gallup.com/poll/648308/far-fewer-regard-childhood-vaccinations-important.aspx>).

Success, however, requires coordination on the correct action (vaccinating). Unlike the social learning literature, where individuals typically aim to make the correct decision for themselves, our setting features network-dependent payoffs: an individual benefits only if the entire population — or a large majority — chooses the correct action, as is the case with herd immunity in vaccination campaigns. Moreover, participants engage in informal communication: they choose tentative actions that are visible to their neighbors but can be revised over time. This feature also distinguishes our design from the experimental repeated game literature [7] and from social learning in networks [11], where actions both reveal intentions and have immediate payoff consequences.

We then explore how different network structures—in which participants observe only the actions of a limited number of neighbors—affect the ability of the group to aggregate dispersed information and coordinate on the correct action. Our goal is to identify which features of the environment facilitate or hinder collective success in the presence of local information frictions. Furthermore, studying both high school and undergraduate students allows us to capture behavioral variations across populations with different levels of maturity, experience, and strategic sophistication. This design enables us to explore whether coordination failures arise from network structure, individual traits, or both.

Our experiment shows that successful coordination depends on a combination of structural and behavioral factors. First, and perhaps unsurprisingly, network size matters. Because the game exhibits a weakest-link feature, where a single deviation is sufficient to prevent success, coordination naturally becomes more fragile as the number of participants increases. Second, neighborhood size plays a critical role. Networks with more interconnected participants create greater opportunities for information exchange and diffusion, making consensus more likely. Third, the distribution of signals shapes outcomes. When individuals with incorrect signals are concentrated within small, tightly-knit neighborhoods, it is much harder to overturn local misinformation than when incorrect signals are evenly dispersed across the population. Finally, we find striking behavioral differences across populations. Undergraduate students are more likely to act initially in line with their private signal and adjust their behavior quickly in response to their peers. Compared to high school students, they switch actions more frequently, adapt faster to changes in their local environment, and exhibit greater responsiveness to their neighbors' choices, behaviors that facilitate coordination in dynamic settings.

Together, our results show that even when individuals share common interests, coordination is fragile in locally connected networks, and success depends critically not only on the design of the network but also on the behavioral characteristics of its participants. This highlights the core challenge of building environments that support information diffusion and enable effective collective action.

2 Results

2.1 Design

We conducted the experiment with networks of 10, 15 or 20 high schoolers (ages 14 to 18) from Oscar de la Hoya Animo Charter High School (ODLH) and undergraduates (ages 18 to 23) from the University of Southern California (USC) distributed as follows:

Population	ODLH	ODLH	USC	USC	USC
Network size	10	15	10	15	20
# of networks	19	15	6	7	5
total individuals	190	225	60	105	100

For ODLH, we set up a portable laboratory in a classroom using PC tablets interconnected via a closed wireless network. Participants were seated in groups of 10 and 15 at individual stations with physical separations to preserve anonymity. The experiment was programmed in oTree [5]. Procedures at USC were identical, except that sessions were conducted at LABEL, our experimental laboratory in the Department of Economics, and included networks of 20 participants in addition to those of size 10 and 15.

Participants in a network played a game in which nature first selected a state of the world and sent a private signal to each individual. In each round, 60% of participants received the correct signal, while the remaining 40% received the wrong signal (see Figures 1(a), 1(b) and 1(c)).

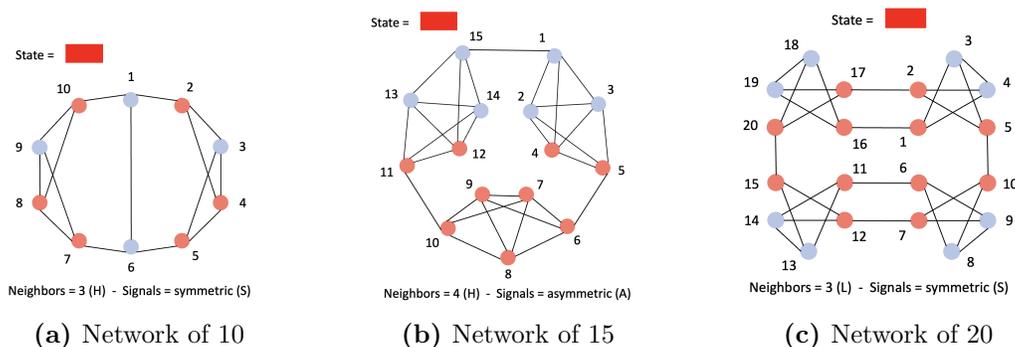


Figure 1: Panels 1(a), 1(b) and 1(c) show three examples of the full network structure, *visible only to the experimenter* for networks of 10 (*HS*), 15 (*HA*) and 20 (*LS*), respectively. The solid color represents the true state (red in all three cases). Each node corresponds to one player ID, and the light color on each node represents the signal received at the beginning of the round by that player: 60% correct (light red) and 40% wrong (light blue). Each line represents a direct connection, that is, a neighbor whose choice is observed by the participant. In networks of 15, there are 3 clusters of players (IDs 1 to 5, 6 to 10 and 11 to 15) and in networks of 20, there are 4 clusters of players (IDs 1 to 5, 6 to 10, 11 to 15 and 16 to 20).

Participants had 60 seconds (for networks of 10), 75 seconds (for networks of 15), or 90 seconds (for networks of 20) to update their actions (i.e., their color choices). During this period, they could observe in real time the choices made by a limited set of neighbors—their direct links (see Figure 2(a)). If, at any point within the allotted time, all participants in the network simultaneously coordinated on the correct color (the one matching the true state), the round was stopped, marked as a success, and each participant received a positive payoff (see Figure 2(b)). If the time expired without full coordination, the round was considered a failure and no payoff was earned (see Figure 2(c)).

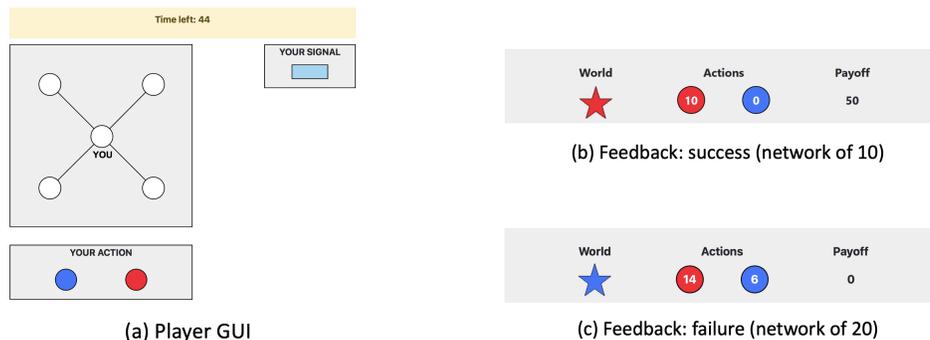


Figure 2: Graphical User Interface. Panel 2(a) shows a screenshot of the game as observed by a player. The player is labeled “you” and always placed in the center with the neighbors (four in this example) arranged equidistantly to prevent framing. The signal appears in the top right corner in light color at the beginning of each round. Players activate and change their choice by tapping on a color under “your action”. Choices are immediately observed by neighbors. The timer at the top of the screen runs backwards. Panels 2(b) and 2(c) show the feedback at the end of the round, including correct state (color of the “World”), number of players on each color, and payoff obtained by all players in the network.

In each round, participants were assigned either a high (H) or a low (L) number of neighbors with whom they could share information. Additionally, signals were distributed either symmetrically (S) or asymmetrically (A) across the network. A symmetric distribution meant that wrong signals were evenly dispersed across the network (see e.g. Figures 1a and 1c) whereas an asymmetric distribution meant that wrong signals were concentrated in one or two small groups of individuals (see e.g. Figure 1b). We employed a 2×2 within-subject design, in which each participant played 4 rounds of each condition (HS , LS , HA , LA), presented in counterbalanced blocks of two. Participants played all rounds with the same set of partners but were randomly and anonymously reassigned to a different position within the network in each round, for a total of 16 rounds. At the end of each round, participants learned the correct state, the final number of participants selecting each color, and the payoffs received by all members of the network.

All sessions followed the same structure: (i) read aloud instructions; (ii) one practice

round with time for questions; (iii) 16 paid rounds with random anonymous rematching; (iv) cash payment. Sessions lasted approximately 40 minutes, and earnings consisted of a \$10 show-up fee plus \$0.50 per successful round for an average of \$11.86 at ODLH and \$14.34 at USC.

We hypothesized that coordination success would be lower when the network size was larger (due to a weakest-link problem), when participants had fewer neighbors (limiting information flow), and when signals were distributed asymmetrically (increasing the risk of local clusters with self-fulfilling misinformation). Consequently, we expected the lowest success rates in condition *LA* and the highest in *HS*. When comparing across populations, we expected higher success rates among undergraduates from a highly ranked university (USC) than among high school students from a historically underprivileged neighborhood (ODLH). We provide the full set of instructions in Appendix SI1. The datasets generated and analyzed in this study are available at this [link](#).

2.2 Network outcomes

Our primary research question is whether and how networks can successfully aggregate dispersed information to reach consensus on the correct state. In principle, the task would be trivial in a world with global information—where individuals could observe everyone else’s choices. In such a scenario, individuals could start by selecting the color of their private signal, then observe the decisions of all others, and deduce the correct state by identifying the majority choice. At the opposite extreme, in the absence of any connections or opportunities to share information, successful aggregation would be nearly impossible, as individuals would have no basis for revising their initial belief beyond their own signal.² With access to local information and the possibility of sharing, the problem becomes more complex: information spreads gradually and imperfectly through the network. However, as we will show, successful coordination remains feasible under these conditions.

The determinants of success

[Figure 3](#) reports the proportion of successful rounds. [Table 1](#) reports a Probit regression of the likelihood of network convergence on the correct action.

There is a large gap in the proportion of successful rounds between ODLH and USC participants (26.0% v. 82.3% and 21.3% v. 69.6% in networks of 10 and 15, χ^2 -test of differences in proportions, $p < 0.0001$). Within ODLH, performance does not significantly

²If participants rely solely on their private signals, the optimal strategy is to follow their signal with the same probability that it is correct—in our case, 0.6. Under this strategy, the probability that all participants independently choose the correct state (i.e., successful coordination) is extremely low. Specifically, the probabilities of success are approximately $S_{10} = (0.6)^6 \times (0.4)^4 \simeq 0.0012$, $S_{15} = (0.6)^9 \times (0.4)^6 \simeq 4.1 \times 10^{-5}$ and $S_{20} = (0.6)^{12} \times (0.4)^8 \simeq 1.4 \times 10^{-6}$ in networks of 10, 15 and 20, respectively.

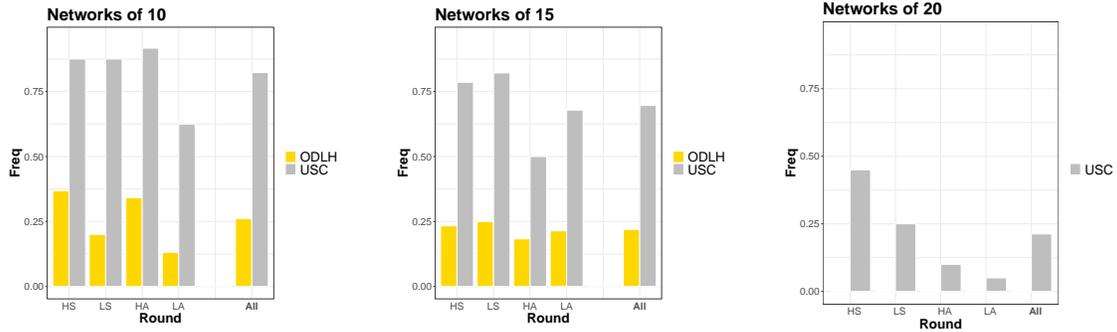


Figure 3: Proportion of successful rounds in each population (ODLH, USC) and network size (10, 15, 20), as a function of neighborhood size (H or L), and signal distribution (S or A), including the average for each case (All).

	ODLH	USC
Network of 15	-0.128 (0.192)	-0.432 (0.330)
Network of 20		-1.835*** (0.318)
A round	-0.149 (0.112)	-0.667*** (0.157)
L round	-0.309* (0.131)	-0.147 (0.183)
2nd half	0.089 (0.127)	0.146 (0.192)
Constant	-0.470* (0.192)	1.311*** (0.255)

*** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$; $p < 0.1$

Table 1: Probit regression of the likelihood of network convergence on the correct action in each population with dummies for network size (reference is 10), neighborhood size (reference is H), signal distribution (reference is S) and first v. second half of the experiment (reference is first).

differ across network sizes ($p = 0.236$). In contrast, performance at USC varies with network size, showing a modest decline from networks of 10 to 15 and a substantial drop in networks of 20 (pairwise comparisons with FDR correction between 10 and 15, $p = 0.051$; between 10 and 20, $p < 0.0001$; between 15 and 20, $p < 0.0001$).

The difference in performance between the allegedly easiest (HS) and most difficult (LA) treatments is significant in USC 10 (0.875 v. 0.625, $p = 0.046$), USC 20 (0.45 v. 0.05, $p = 0.003$), and ODLH 10 (0.36 v. 0.13, $p = 0.0007$), but not in USC 15 (0.79 v. 0.68, $p = 0.36$) and ODLH 15 (0.23 v. 0.21, $p = 0.81$). This pattern is consistent with the effects of asymmetry and the number of neighbors on performance, as reported in Table 1.

Two simple explanations for why some networks fail to achieve convergence could be that participants do not put in sufficient effort or that network failure is driven by a small minority of unskilled, stubborn, or inattentive individuals. However, we find evidence that contradicts both explanations. In [Table 2](#), we analyze unsuccessful trials by tracking, for each round, the maximum and minimum fractions of correct actions reached within the network over time, as well as the final fraction of correct actions. We then compute averages across all unsuccessful rounds, broken down by population and network size.

	ODLH 10	ODLH 15	USC 10	USC 15	USC 20
Max. correct	.68	.50	.65	.70	.70
Min. correct	.27	.52	.32	.28	.25
Final correct	.45	.47	.45	.47	.51
# obs.	304	232	96	112	80

Table 2: Maximum and minimum number of correct actions (on average) over the course of an unsuccessful round and final number of correct actions (on average) at the end of an unsuccessful round, computed separately for each population and network size.

With the exception of ODLH 15, we observe substantial variation in the number of correct actions within a round, ranging from one-quarter (minimum) to two-thirds (maximum). This spread indicates that participants are actively experimenting with different strategies: switching actions, taking the lead, or mimicking peers, all in an effort to coordinate on the true state. Regarding the final choices, the average proportion of correct actions in unsuccessful rounds consistently hovers around one-half. This suggests that failure cannot be attributed to a small subset of participants who persistently misunderstand the task and undermine the group outcome. Rather, the lack of convergence appears to be a robust phenomenon driven by the inherent difficulties of aggregating dispersed information within a network.³

What behavioral differences might then explain the gap in success rates across the two populations? The answer lies in differences in exploration behavior. First, the average number of actions per participant in successful and unsuccessful rounds combining networks of 10 and 15 is 1.89 and 2.54 in ODLH, compared to 2.26 and 3.43 in USC. This corresponds to 20% and 35% more switches in USC than in ODLH in successful and unsuccessful rounds respectively (test of differences between populations, $p < 0.0001$ in both cases). Second, the proportion of rounds in which all participants simultaneously coordi-

³To illustrate the difficulty of achieving consensus, consider a 20-player example where 10 participants (IDs 6 to 15 in [Figure 1c](#)) choose red and the other 10 choose blue. Among red players, 2 individuals (IDs 10 and 15) have two neighbors choosing red and one choosing blue, while the remaining 8 have all three neighbors choosing red. The same symmetric structure holds for blue players. This example demonstrates that even with sophisticated individuals, local homogeneity within clusters can coexist with global heterogeneity across clusters.

nate on the wrong state and subsequently change their actions is significantly higher in USC than in ODLH: 32.3% vs. 18.4% in networks of 10, and 27.7% vs. 14.6% in networks of 15 (χ^2 -tests of differences in proportions, $p = 0.006$ and $p < 0.0001$, respectively). Conversely, the proportion of rounds in which participants coordinate on the wrong state and remain there is significantly lower in USC than in ODLH: 1.0% vs. 8.2% in networks of 10, and 1.8% vs. 4.8% in networks of 15 (χ^2 -tests, $p < 0.0001$ in both cases).

Taken together, these results suggest that USC participants explore more — sometimes converging prematurely on the wrong state — but crucially, they continue experimenting as long as time permits. By contrast, ODLH participants explore less — reducing the likelihood of coordinating on the wrong state — but once they and their immediate neighbors align on a common action, they appear more inclined to believe that failure to converge is due to errors in other parts of the network, and thus tend to stay put.

Choice dynamics

Given the large shifts in the number of correct actions within a round (Table 2), it is informative to examine the dynamics of choice in more detail. Among successful rounds, the average time (in seconds) it takes participants to converge on the true state is: 32.2 (ODLH 10), 45.2 (ODLH 15), 26.8 (USC 10), 42.1 (USC 15) and 51.9 (USC 20). Not surprisingly, convergence is slower in networks of 15 compared to networks of 10 in both populations (t-tests, $p = 0.001$ for ODLH and $p < 0.0001$ for USC). Convergence is also significantly slower for USC in networks of 20 compared to networks of 10 (t-test, $p = 0.003$) though not significantly different from networks of 15 ($p = 0.211$). Finally, while convergence is slower in ODLH than in USC overall (Figure 3), the difference is only statistically significant in networks of 10 (t-test, $p = 0.033$).

Next, we examine the within-round evolution of the number of correct actions taken by participants in the network. Figure 4 illustrates this dynamic for selected treatments. We focus on successful rounds, as they contain richer information, including both the time to convergence and the magnitude of fluctuations in choices along the way. We also restrict the analysis to USC, which offers sufficient observations to enable comparisons across treatments (*HS* v. *LA*). In contrast, the number of successful rounds in ODLH is too small to allow for meaningful comparisons.

Even though we focus only on successful rounds within a single population, we observe dramatic differences across network sizes and treatments. In smaller networks (10 and 15), choice dynamics are characterized by frequent switching (“chattering”), whereas in larger networks (20), transitions are slower and smoother. This pattern leads to faster convergence on the true state in smaller networks but also increases the likelihood that all participants temporarily coordinate on the wrong state before correcting their actions.

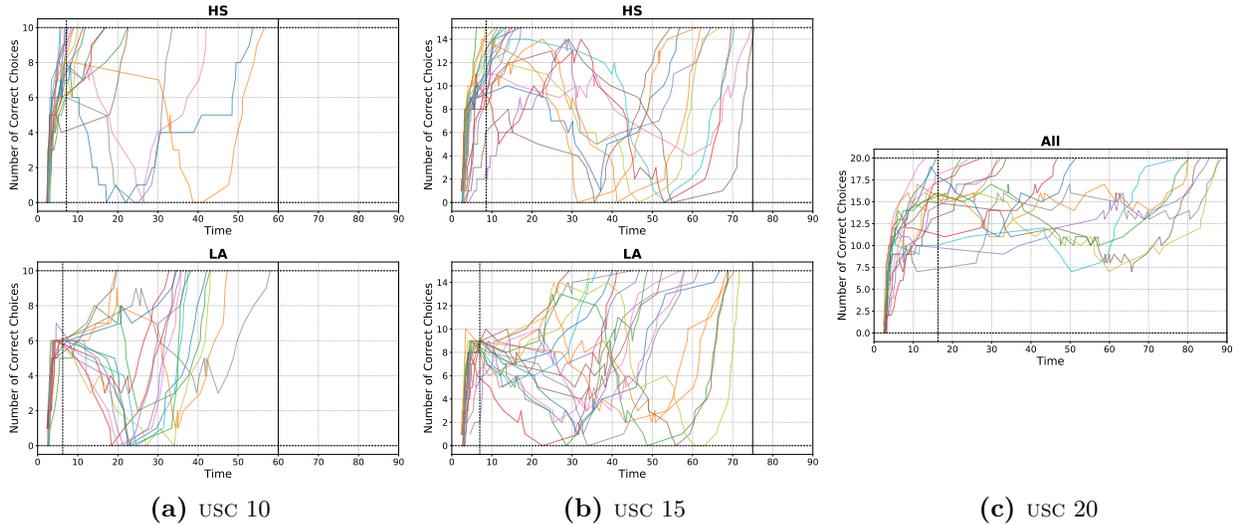


Figure 4: Each line shows the evolution of the number of participants at USC choosing the correct action in a successful round as a function of elapsed time (in seconds). The vertical dotted line indicates the average time at which all participants in that treatment have made at least one choice. The vertical solid line marks the time limit for the round (60s, 75s or 90s). The lower and upper horizontal dotted lines represent the bounds where no participant or all participants, respectively, are simultaneously choosing the correct action. Panels (a) and (b) display the *HS* (top) and *LA* (bottom) treatments for networks of 10 (left) and 15 (center). Panel (c) presents all successful treatments for networks of 20 (right).

Within each network size, we also find significant differences between the treatments we identified as the easiest (*HS*) and the hardest (*LA*). While the overall proportion of successful rounds does not differ dramatically between these treatments (Figure 3), the speed of convergence does. Fast convergence (within 20 seconds) occurs frequently in *HS* (71.4% and 54.6% of successful rounds for networks of 10 and 15, respectively) but is rare in *LA* (13.3% of successful rounds for networks of 10 and never for networks of 15).

Finally, motion animations illustrating the dynamics of choices in representative networks are available at <https://labelinstitute.github.io/NetworkMovies/>. In these animations, colored rings indicate the private signal, while the right-hand panel displays the sequence of choices, including the participant’s ID, selected color, and time of action. For ease of visualization, the animations proceed one choice at a time at uniform time intervals, without reflecting the actual timing delays between choices.

2.3 Individual analysis

While the network analysis accounts for aggregate performance, it is also important to examine behavioral differences at the individual level. In Figure 5, we report the initial

choice of ODLH and USC participants in the network sizes common to both populations (10 and 15), using this measure as an indicator of their initial intentions.

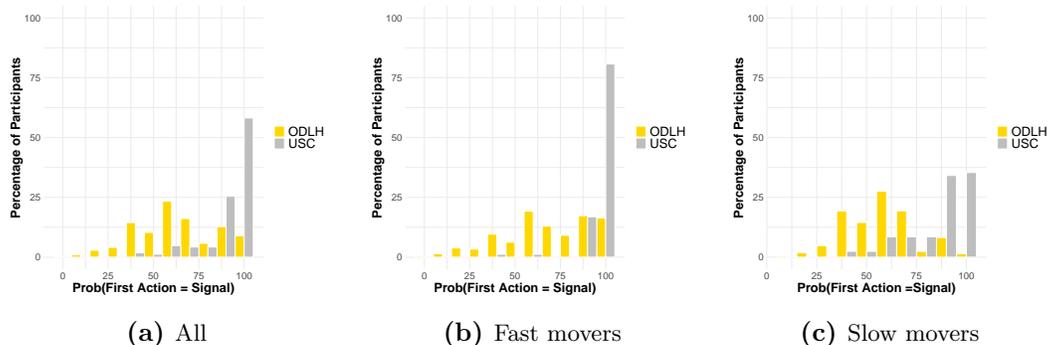


Figure 5: Panel (a). Distribution of individuals in each population for networks of 10 and 15 as a function of the proportion of rounds where first action coincides with signal. Panels (b) and (c). We split the sample into individuals whose average delay until first action is below median (fast movers) and above median (slow movers), and perform in each subsample the same analysis as in Panel (a).

We observe a substantial difference in initial behavior across populations (Figure 5a). USC participants make their first decision very quickly (median time: 2.7 seconds) and almost systematically follow their private signal (91.8%). By contrast, ODLH participants are significantly slower in making their initial choice (median time: 4.8 seconds) and are less likely to follow their signal (63.8%). The distribution of decision times differs between ODLH and USC participants (Wilcoxon rank-sum test, $p < 0.0001$), and the average probability of following one’s own signal is also significantly higher at USC (t-test, $p < 0.0001$). These findings suggest fundamentally different strategies across populations — not only in how participants react to the choices of others (as previously observed), but also in how they approach the decision problem from the outset.

When we divide the USC sample based on response time, we find that fast movers (Figure 5b) always follow their signal, whereas some slow movers (Figure 5c) occasionally do not (97.4% vs. 87.0% on average). A similar pattern emerges in ODLH, despite their lower overall tendency to follow signals (69.6% vs. 58.0% on average). Differences are statistically significant in both populations (t-tests, $p < 0.0001$). Thus, although the two populations differ in how they approach the initial choice, in both cases fast movers are more likely to follow their signal and take the lead, while slower movers are more likely to weigh their private signal against the observed actions of others before deciding.

Next, we present in Figure 6 the individual-level performance in the game.

Not surprisingly in view of Figure 3, individual performance is substantially higher

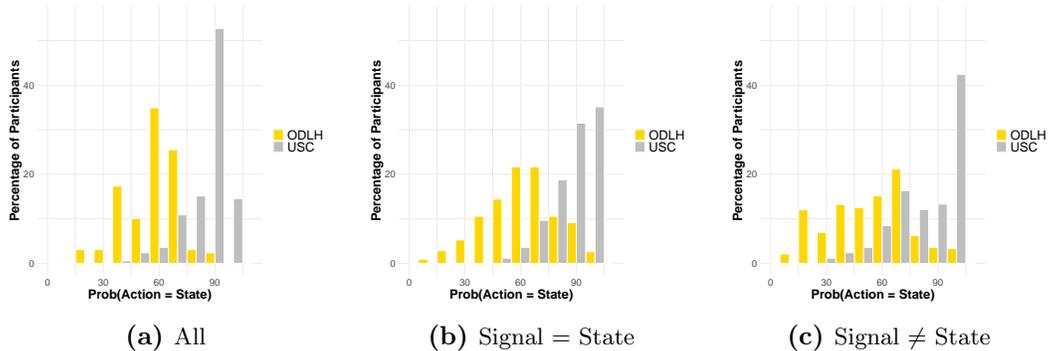


Figure 6: Panel (a). Distribution of individuals in each population for networks of 10 and 15 as a function of the proportion of rounds where the final action of the individual is correct. Panels (b) and (c). We split the sample into cases where the individual received the correct and the wrong signal, and perform in each subsample the same analysis as in Panel (a).

in USC than in ODLH (KS test of stochastic dominance, $p < 0.0001$ for all three pairs of distributions), with an impressive 67% of USC participants finishing at least 14 out of 16 rounds in the correct action (Figure 6a).

Interestingly, within each population, the difference in the number of correct actions between rounds where participants received the correct signal (Figure 6b) and rounds where they received the wrong signal (Figure 6c) is statistically significant but modest in magnitude: 88.2% vs. 83.7% in USC, and 61.5% vs. 52.2% in ODLH (paired t-tests, $p = 0.002$ and $p < 0.0001$, respectively). This suggests that participants in both populations are influenced by their private signal, but not dramatically so. Instead, they are (rightly) willing to adapt their choices to align with the actions of their neighbors. This finding contrasts with results from the traditional social learning literature, which often reports strong overweighting of private information [9].

Finally, we focus on networks of size 15 and identify situations within a round where a cluster of interconnected individuals (IDs 1 to 5, IDs 6 to 10, or IDs 11 to 15 in Figure 1b) all simultaneously choose the same action. These situations are particularly prone to generating stagnation — either on the correct or incorrect action — as individual choices are reinforced by those of immediate neighbors. Consistent with the differences in action switching across populations, clustering events occur, on average, 3.21 times per round in ODLH and 4.01 times per round in USC. We then compute the number of instances in which a participant breaks the consensus of their cluster despite observing the same action among all of their neighbors, as opposed to switching because they observe a different action from one individual in another cluster. We find that deviations without any supporting evidence

from neighbors occur with probability 0.69 in ODLH and 0.51 in USC (test of difference of proportions, $p < 0.0001$). This suggests that although high school participants are generally less likely to switch actions than undergraduates, conditional on switching, they are more prone to breaking local consensus without any observable justification, that is, without information from their immediate environment to support the change.

3 Discussion

Consensus is far from guaranteed even when everyone stands to benefit from the result. Local observation of neighbors' actions often leads to pockets of incorrect beliefs that stall global coordination, especially in larger networks, networks with fewer connections, and networks with asymmetric distribution of signals. Participants frequently switch their choices in an effort to coordinate, but final outcomes nonetheless vary greatly by treatment and population. The overall results underscore the fragility of coordination under imperfect information and local sharing.

USC *v.* ODLH. A stark contrast arises between undergraduate and high school participants in initial choice, frequency of switching, and likelihood of convergence. USC participants adopt a more exploratory approach. They appear to reason that if their entire neighborhood has coordinated on one action but the network remains stuck, their pocket is likely in error. By contrast, many ODLH participants seem to trust their local consensus once it is reached, believing that if global consensus has not been achieved, it is because other neighborhoods are making the mistake. A potential explanation for these behavioral differences (along the lines of repeated game experiments among individuals with different IQ [18, 19] or different academic achievements [3]) is that USC participants are more sophisticated strategic thinkers. Alternatively, with more maturity they may simply be more willing to acknowledge that their decision might be wrong and thus be more open to reverse it.

Designing networks to help coordination. Natural extensions could help understand the elements that facilitate or hinder the aggregation of information. First, requiring supermajority rather than unanimity might mitigate the weakest link problem. Second, some players could receive no private signal, forcing them to rely on the actions of their better informed neighbors, as in the political economy literature on voting [8]. This variant would reveal if confusion spreads globally or those with private information become pivotal influencers. Third, networks with skewed connectivity structures — for example, a star topology — could highlight the role of central hubs in accelerating or impeding convergence, as in [20]. Finally, adding a third possible state might increase the complexity

of the inference problem, but also reduce large-scale miscoordination, since each incorrect signal would reach a smaller fraction of the network. These variations directly relate to insights from percolation theory [22]. Our findings suggest that coordination failure is not simply due to unwillingness to cooperate but reflects structural limitations that prevent correct information from propagating through the network.

Individual traits and group composition. The composition of the population—psychological traits, preferences, demographic characteristics or prior experience—is likely to play a role in coordination dynamics. The literature on leadership and critical mass in networks [13, 21] emphasizes that a small number of influential individuals can play a pivotal role in steering behavior adoption, especially in environments with information frictions. In our context, introducing “partisan” players who care only about promoting a specific action, as often seen in political or ideological contexts [10], would test the resilience of collective judgment in the face of misinformation. Alternatively, grouping individuals based on personality traits such as agreeableness might facilitate coordination due to greater cohesion, but could also increase the risk of premature convergence on incorrect actions if no one is willing to challenge the local consensus. Similarly, same-gender networks might exhibit more consistency in risk attitudes, trust, or communication style, all of which could influence behavior. Comparing behavior across homogeneous networks could shed light on which social or cognitive factors most directly promote information aggregation. Another relevant extension concerns mixing populations of different backgrounds or strategic sophistication. For example, combining USC undergraduates and ODLH high schoolers in the same network could reveal whether dynamic leaders from the first group induce hesitant participants from the second to switch actions, demonstrating positive peer influence. This type of hybrid design would clarify how much critical mass of dynamic players is needed to steer a network—a question with direct implications for leadership, organizational behavior, and the design of interventions in real-world environments where ability levels, education, and risk attitudes vary.

Acknowledgements. We thank the members of the Los Angeles Behavioral Economics Laboratory (LABEL), especially Daniel Hermosillo, Prianca Nadkarni, Ulysses Rios and Andrea Victoria, for their insights and help running the experiment and to Sergio Sapiña from www.playstudies.com for his amazing programming skills. We are grateful to the staff at Oscar de la Hoya Animo Charter High School—in particular Kim Kawaratani and Robert Pambello—for their help and support. The study was conducted with the University of Southern California IRB approval UP-12-00528. We gratefully acknowledge the financial support of the National Science Foundation (NSF) grant SES-2315770 and the French National Research Agency (ANR) grant ANR-17-EURE-0010.

References

- [1] Lisa R Anderson and Charles A Holt. Information cascades in the laboratory. *The American economic review*, pages 847–862, 1997.
- [2] Sushil Bikhchandani, David Hirshleifer, Omer Tamuz, and Ivo Welch. Information cascades and social learning. *Journal of Economic Literature*, 62(3):1040–1093, 2024.
- [3] Isabelle Brocas and Juan D Carrillo. Academic achievement helps coordination on mutually advantageous outcomes. *Proceedings of the National Academy of Sciences*, 122(14), 2025.
- [4] Boğaçhan Çelen and Shachar Kariv. Distinguishing informational cascades from herd behavior in the laboratory. *American Economic Review*, 94(3):484–498, 2004.
- [5] Daniel L Chen, Martin Schonger, and Chris Wickens. otree—an open-source platform for laboratory, online, and field experiments. *Journal of Behavioral and Experimental Finance*, 9:88–97, 2016.
- [6] S Choi, E Gallo, and S Kariv. "Networks in the laboratory". The Oxford Handbook of the Economics of Networks, Y. Bramoullé, A. Galeotti and B. Rogers eds. 2016.
- [7] Pedro Dal Bó and Guillaume R Fréchette. The evolution of cooperation in infinitely repeated games: Experimental evidence. *American Economic Review*, 101(1):411–429, 2011.
- [8] Timothy J Feddersen and Wolfgang Pesendorfer. The swing voter’s curse. *The American economic review*, pages 408–424, 1996.
- [9] Jacob K Goeree, Thomas R Palfrey, Brian W Rogers, and Richard D McKelvey. Self-correcting information cascades. *The Review of Economic Studies*, 74(3):733–762, 2007.
- [10] Kenneth F Greene, Erin L Rossiter, Enrique Seira, and Alberto Simpser. Interacting as equals reduces partisan polarization in mexico. *Nature Human Behaviour*, 9(1): 147–155, 2025.
- [11] Veronika Grimm and Friederike Mengel. Experiments on belief formation in networks. *Journal of the European Economic Association*, 18(1):49–82, 2020.
- [12] Matthew O. Jackson. *Social and economic networks*. Princeton University Press Princeton, 2008.

- [13] Matthew O Jackson and Leeat Yariv. Diffusion, strategic interaction, and social structure. In *Handbook of social economics*, volume 1, pages 645–678. Elsevier, 2011.
- [14] Matthew O Jackson, Brian W Rogers, and Yves Zenou. Networks: An economic perspective. In *Oxford Handbook of Social Networks*. Oxford University Press, 2020.
- [15] Matthew O Jackson, Suraj Malladi, and David McAdams. Learning through the grapevine and the impact of the breadth and depth of social networks. *Proceedings of the National Academy of Sciences*, 119(34):e2205549119, 2022.
- [16] Stephen Judd, Michael Kearns, and Yevgeniy Vorobeychik. Behavioral dynamics and influence in networked coloring and consensus. *Proceedings of the National Academy of Sciences*, 107(34):14978–14982, 2010.
- [17] Michael Kearns, Siddharth Suri, and Nick Montfort. An experimental study of the coloring problem on human subject networks. *science*, 313(5788):824–827, 2006.
- [18] Eugenio Proto, Aldo Rustichini, and Andis Sofianos. Intelligence, personality, and gains from cooperation in repeated interactions. *Journal of Political Economy*, 127(3):1351–1390, 2019.
- [19] Eugenio Proto, Aldo Rustichini, and Andis Sofianos. Intelligence, errors, and cooperation in repeated interactions. *The Review of Economic Studies*, 89(5):2723–2767, 2022.
- [20] Rong Rong and Daniel Houser. Growing stars: A laboratory analysis of network formation. *Journal of Economic Behavior & Organization*, 117:380–394, 2015.
- [21] Tanya S Rosenblat and Markus M Mobius. Getting closer or drifting apart? *The Quarterly Journal of Economics*, 119(3):971–1009, 2004.
- [22] Dietrich Stauffer and Amnon Aharony. *Introduction to percolation theory*. Taylor & Francis, 2018.

Supplementary Information

SI1. Instructions

Hello, thank you for coming. Today, we are going to play a few computerized games and make money. In all the games, you are going to win points. At the end, you will be paid \$1 for every 100 points accumulated. You will also be paid \$10 just for playing with us.

The computer is going to form a network with many people. Each of you will be directly connected only to a few people but the people you are connected to are also themselves connected to other people and so on. This means that everyone is indirectly connected to everyone else in the network. Think of your connections as being the people you talk to, who themselves talk to other people, etc. Now, what do you have to do?

The goal of the game is to guess the color of the world. The world can be RED or BLUE. **[SLIDE 1]** - [see [Figure SI1](#) for the slides projected on the screen at the front of the classroom]

The computer will choose one color of the world with 50-50 chance, but it will not tell you. Instead, it will give each of you a signal.

- If the world is “RED”, more people will see a signal “light red”, and fewer people will see a signal “light blue”.

- If the world is “BLUE”, more people will see a signal “light blue”, and fewer people will see a signal “light red”.

Note that more people get the signal that corresponds to the true color of the world but not everyone, otherwise the game would be too easy. We are using light colors to make clear that it is an indication, and therefore it may be wrong. In your computer, you will see a screen like this. **[SLIDE 2]**

Let me walk you through the different parts of the screen. **[SLIDE 3]**

First, you can see your neighborhood. These are the people you are directly linked to, in this case 4 people but in your session, it may be a different number. Visually, you are always at the center of your neighborhood, and these are your neighbors [point to the neighbors in the screen]. This is because you see your neighborhood from your own perspective. Other people will be at the center of their own neighborhood.

If you are linked to one person, that person is linked to you. But remember, they will be also linked to people you are not linked to and you will be also linked to people they are not linked to. **[SLIDE 4]**

Second, you can see your signal. In this case, your signal provides support that the world is BLUE. But, of course, you cannot know for sure. **[SLIDE 5]**

Third, you can see your possible actions. That’s where you make your decisions. You have to choose between the RED action and the BLUE action. As soon as you choose an action, you and your neighbors will see it [point at area (1)]. When your neighbors choose their actions, it will also appear here. After you see their actions, you can change yours... or not. It is totally up to you.

The action is important because it determines your payoff. If at some point during the round everyone in the entire network, that is everyone in all the neighborhoods, chooses the action that matches the color of the world, you will all earn 50 points, that is 50 cents. If at least one person in the network chooses the action that does not match the color of the world, you will all get 0 points. Of course, if you all choose the same action but it does not match the color of the world,

you will still get 0 points. Remember you get 50 points only if everyone chooses the action that matches the color of the world. Is this clear? [SLIDE 6]

Finally, you will see a timer at the top of the screen. This tells you the amount of time left to play. You start with a certain amount of time and the clock runs backwards. You can change your action as many or as few times as you want and the same is true for your neighbors and everyone else in the network. Every time you change your action it appears immediately on the screen of your neighbors and every time your neighbors change their actions, it appears on your screen. If you want to change your action multiple times, you have to wait at least two seconds between changes. Now this is very important. If at any point you all chose the correct action, the clock stops, and you all get your 50 points. If you run out of time and at least one person has the wrong action, you all get 0 points.

At the end of the round, you will see a screen like this. [SLIDE 7]

It tells you what the actual color of the world was (in this case, blue), how many people chose the RED and BLUE actions (in this case 13 and 7) and your payoff (in this case 0, since not everyone chose blue).

You are going to play this same game multiple times. Each time, the computer will randomly choose the color of the world with a 50-50 chance. Each time you will get a new signal and a new set of neighbors, but the rules of the game are always the same: if you all collectively choose the color of the world you all get 50 points.

Before we start, we are going to play a practice round. This is only to familiarize yourself with the screen, the software, and the commands. The results of this round do not count for your final payment so feel free to play around, experiment and change the actions multiple times. Any questions?

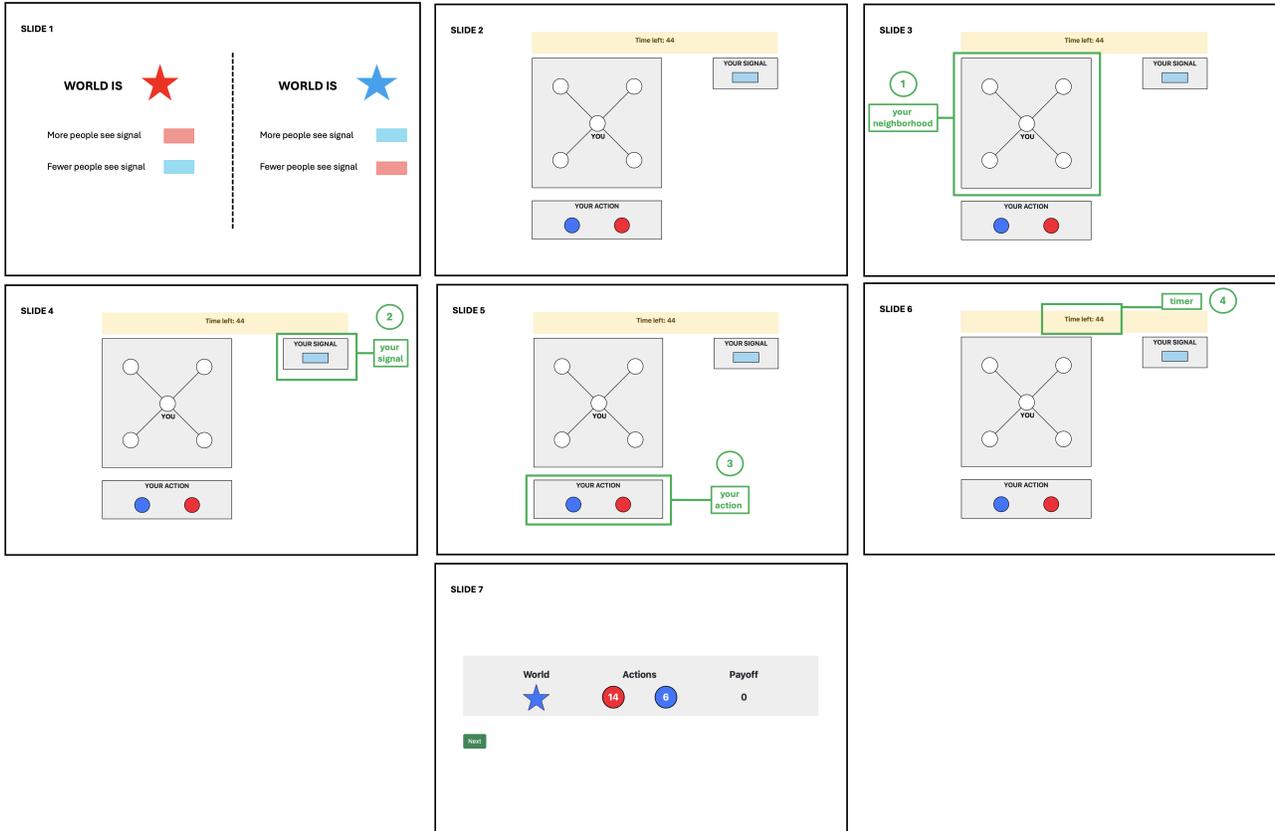


Figure SI1: Slides projected on screen for instructions

SI2. Description of all the network variants

We conducted 4 rounds of each variant (HS , LS , HA , LA), where H is high number of neighbors, L is low number of neighbors, S is symmetric distribution of signals, and A is asymmetric distribution of signals. Figures SI2, SI3 and SI4 present the four variants for networks of 10 (ODLH and USC), 15 (ODLH and USC) and 20 (only USC). In all these examples, the correct state is red.

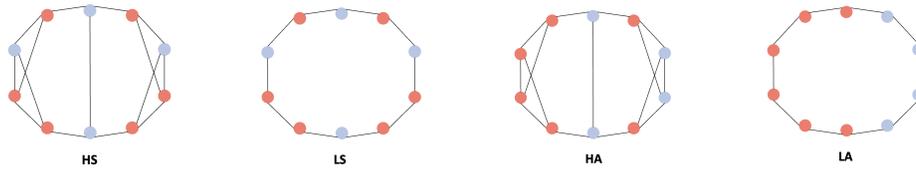


Figure SI2: Treatments HS , LS , HA and LA in networks of 10 (ODLH and USC)

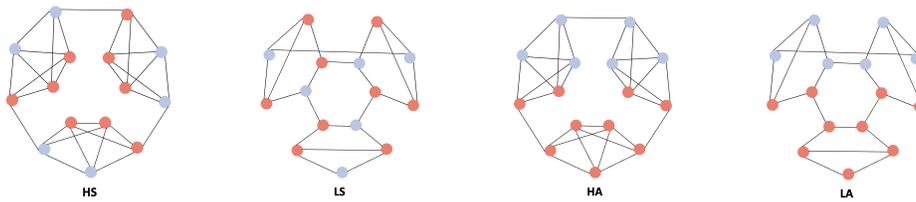


Figure SI3: Treatments HS , LS , HA and LA in networks of 15 (ODLH and USC)

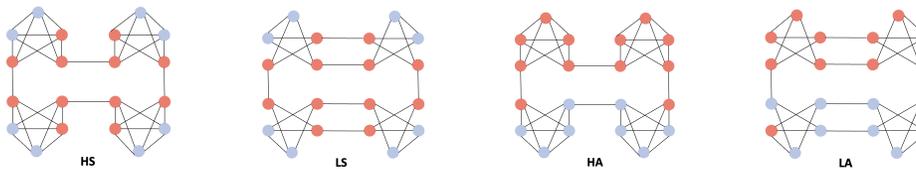


Figure SI4: Treatments HS , LS , HA and LA in networks of 20 (USC)