HHS Public Access
Author manuscript
*Nature*. Author manuscript; available in PMC 2018 April 23.

Published in final edited form as:
*Nature*. 2017 May 17; 545(7654): 370–374. doi:10.1038/nature22332.

# Locally Noisy Autonomous Agents Improve Global Human Coordination in Network Experiments

**Hirokazu Shirado**[1,2] and **Nicholas A. Christakis**[1,2,3,4,*]

[1]Yale Institute for Network Science, Yale University, New Haven, CT 06520 USA

[2]Department of Sociology, Yale University, New Haven, CT 06520 USA

[3]Department of Ecology and Evolutionary Biology, Yale University, New Haven, CT 06520 USA

[4]Department of Biomedical Engineering, Yale University, New Haven, CT 06520 USA

## Abstract

Coordination in groups faces a sub-optimization problem[1–6] and theory suggests that some randomness may help achieve global optima[7–9]. We performed experiments involving a networked color coordination game[10] in which groups of humans interacted with autonomous software agents ("bots"). Subjects (*n*=4,000) were embedded in networks (*n*=230) of 20 nodes to which we sometimes added 3 bots. The bots were programmed with varying levels of behavioral randomness and different geodesic locations. Here, we show that bots acting with small levels of random noise and placed in central locations meaningfully improve the collective performance of human groups, accelerating the median solution time by 55.6%. This is especially the case when the coordination problem is hard. Behavioral randomness worked not only by making the task of humans to whom the bots were connected easier, but also by affecting the game play of the humans among themselves and hence creating further cascades of benefit in global coordination in these heterogeneous systems.

### Keywords

social coordination; noise; social networks; social experiments; AI; complex systems

---

Collective action and large-scale cooperation are important challenges[1–3]. Most work on cooperation has focused on the social dilemma aspect, namely, on getting people to be willing to make sacrifices for the greater good[11,12]. Yet, even when this dilemma can be addressed, there remains another substantial problem: coordination[4–6]. The difficulty of achieving optimal collective action in groups may arise not only from the conflicting

---

[*]**Corresponding Author:** Nicholas A. Christakis, M.D., Ph.D., M.P.H., Yale Institute for Network Science, 17 Hillhouse Ave, New Haven, CT 06520; Tel.: +1-203-436-4749; Fax: +1-203-432-5423. nicholas.christakis@yale.edu.

**Supplementary Information** is submitted in a separate file.

**Competing financial interests:** The authors declare no competing financial interests.

**Data availability:** The data reported in this paper are archived at Yale Institute for Network Science and are available upon request.

**Author Contributions:** HS and NAC designed the project. HS collected the data and performed the statistical calculations. HS and NAC analyzed the results. HS and NAC wrote the manuscript. NAC obtained funding.

interests among individuals, or between individuals and their group, but also as a consequence of the inability of individuals to effectively coordinate their actions globally. Even if all individuals behave properly in their local interactions, this may not result in the optimal outcome for the whole community[1,2].

Prior theoretical work has suggested a surprising, even paradoxical, solution to the coordination problem: adding “noise.”[13–15] Noise is usually defined as meaningless information, and it is often seen as problematic[16]. When it comes to optimization, however, noise can help a system to reach a global optimum. For example, mutation has an essential role in evolution[17]; error can facilitate search for information[18]; random fish schooling may enhance survival[19]; and cooperation may benefit from deviant behavior[7–9,20].

Here, we evaluate the benefits of noise in addressing the coordination problem of human groups[21,22]. And given that human interactions are embedded within social networks, we also consider the impact of network position on the potentially beneficial effect of noise[23]. We first characterize the collective-action dynamics of networks of people interacting in a classic color coordination game[10]. Then, we test the effect of noise on collective performance using autonomous software agents (“bots”), manipulating both the noisiness and geodesic placement of the bots. By adding bots into experimental social networks, we therefore explore the performance of *heterogeneous systems* involving both real humans and autonomous agents, while also demonstrating a possible practical solution to the problem of global coordination itself.

We recruited 4,000 unique subjects online and randomly assigned them to one of eleven conditions in a series of 230 sessions (see SI). Subjects were assigned a location in a network of 20 nodes, generated by a preferential attachment model[24]; the network structure was created *de novo* for each session by attaching new nodes (each with two links) to existing nodes; and subjects were placed into the resulting networks at random. The collective goal is for every node to have a color different than all of its neighbor nodes[10]. This color coordination game successfully captures the problem of systematic failure by sub-optimization in coordination; that is, while each individual attempts to reach a solution that is optimal for that individual, this may not be optimal for the whole group (Fig.1a).

In the sessions, each subject was allowed to choose a color from three choices (green, orange, and purple) at any time. The number of colors made available was the minimum necessary to color the entire network without conflicts, which is known as the “chromatic number”; and all networks in our experiments are, by construction, globally solvable. However, while all the networks allowed the subjects to reach the collective goal, the networks could (by chance) vary in their number of solutions (i.e., the networks ranged from 6 to 13,824 possible “colorings” that would work, known as the “chromatic polynomial” – see SI).

Subjects could see only the colors of neighbors to whom they were connected, in addition to their own color. Thus, although a subject might have solved the problem from his or her own point of view, the game might continue because the network still had conflicts in other regions of the graph. In terms of the optimization problem, the game’s cost function is expressed as the sum of the number of conflicts. As in past work[10], the subjects got paid according to how long it took for all conflicts in the network to be resolved, and they had to complete the task within 5 minutes (see SI for details). All subjects consented, and the Yale Committee of the Use of Human Subjects approved this research.

Within this basic setup, we then introduced 3 bots into the network in exchange for the same number of humans (no bots were placed in the control sessions – see Table S1). Subjects were not informed that there were bots. We manipulated noisiness of the bots as follows: In the "zero noise" condition, the bots behaved with a simple, greedy strategy: when a bot had a chance to minimize color conflicts with its neighbors, it chose that color; otherwise, it maintained its current color. In the other two conditions, the bots behaved with the same greedy strategy most of the time, but they also randomly picked a color from the three permissible options regardless of their local situation – with a probability of 10% ("small noise") or 30% ("large noise"). In all the conditions, the bots made decisions every 1.5 seconds, which was the typical human reaction time (Extended Data Fig.1).

Independent of bot noise, we also manipulated their network location as follows: In the "central" condition, the bots were assigned to the 3 positions that had the largest number of neighbors (the highest network degree). Likewise, in the "peripheral" condition, the bots were assigned to the 3 positions with the lowest degree. In the "random" condition, the bots were randomly assigned to their locations. It was permissible for the bots to be connected to each other, by chance, in all conditions.

As noted, the bots acted using only their local information. To assess the effect of such bot behavior compared to the much more demanding case requiring global knowledge of the entire network structure and its solution space in advance, we also carried out experiments with a "fixed color" condition. In this extra condition, we evaluated all color combinations of each network that resulted in no conflicts, and then assigned the initial colors of three of the nodes based on one of those combinations (chosen at random). That is, during the game, the 3 nodes were not controlled by bots that coordinated with their neighbors, but rather, these nodes simply stayed at their initial colors, which were known to be consistent with a global solution to the problem. We examined this treatment only in the case where the fixed nodes were in the central condition.

In sum, we evaluated eleven conditions: one control condition not involving any bots; nine treatment combinations of noise and location of bots (3 levels of behavioral randomness – 0%, 10%, and 30% – crossed with 3 types of location – random, central, and peripheral), and one final condition with 3 fixed-color nodes. We conducted 30 sessions for the control condition and 20 sessions for each of the treatment conditions (so as to be able *a priori* to evaluate at least a 30% difference in solvability), for a total of 230 sessions and 4,000 subjects.

For the games involving only human subjects, 20 of 30 resulted in an optimal coloring of the network in less than the allotted 5 minutes (median time = 232.4 seconds; IQR 143.7 – 300.0). Although the subjects aimed to eliminate all the conflicts, they often found themselves unable to reach the collective goal only by reducing their local conflicts on an individual basis. For example, as of 105 seconds in Fig.1a (or Video S1), each of the subjects had chosen one of the least common colors among their neighbors; that is, no one person could change their color for the better. A conflict between neighbors, however, still remained. Such states in which players get caught in locally unresolvable conflicts are regarded as local minima of the game's cost function (in contrast to resolvable conflicts which can be addressed by local action). Players would need a moderate level of deviancy from the norm of conflict minimization in order to overcome the local minimum and reach a global solution (e.g., Fig.1a, at 245 seconds).

By analyzing the sessions involving only human subjects, it is possible to discern that games were more likely to be solved when some players occasionally chose a locally inappropriate color, temporarily increasing conflicts[10]; moreover, the effect of such behavioral deviance varied according to the geodesic location of the players, as captured by their network degree (Fig.1b). In addition, and distinctly, some networks could be intrinsically easier to solve (i.e., the chromatic polynomial could be higher) (Fig.1c).

To demonstrate how bots could improve the performance of human groups, Fig.2 shows survival curves of the sessions involving the 9 bot treatments. Before implementing pairwise comparisons of each treated group with the control group, we performed a log-rank test of the null hypothesis that all the survival curves are identical; that hypothesis was rejected ($P$ = 0.024), indicating that at least two of the survival curves differed. The sessions having bots with 10% noise and central locations were the most likely to be solved within the allotted 5 minutes (17 of 20 sessions, or 85%, compared with 20 of the 30 control sessions, or 67%, with humans alone); moreover, the solution was achieved more than 129.3 seconds faster (i.e., 55.6% faster) than sessions involving just humans (median time = 103.1 seconds [IQR 49.5 – 170.1] versus 232.4 seconds [IQR 143.7 – 300.0]), which was significantly better ($P$ = 0.015, log-rank test).

We then examined the difference in effectiveness of the various bot treatments – while furthermore controlling for the intrinsic solvability of the network – using Cox proportional hazard models. Bot behavioral randomness of 10%, central location, and the logarithm of the chromatic polynomial all have a significantly positive impact on the completion time ($P <$ 0.05; $n$=180 bot-treated sessions; see SI). We also evaluated another metric of the complexity of the solution space (i.e., mean convergence steps with linear probabilities) and got similar results (Extended Data Fig.2 and Table S4). The statistical model with full interactions shows that the bots affect the solution time only when they behave with 10% randomness and are placed in the central location in the network (Fig. 3a); moreover, when the network affords many solutions, the beneficial impact of bots decreases, as shown by the three-way interaction (Fig.3b). In short, the bots are especially helpful when the network is globally hard to solve.

We found that the impact of 10%-noise bots was comparable to the impact of assigning three nodes with fixed (constant) colors in a configuration known *ex ante* to be compatible with a global solution. There was no significant difference between the sessions with 10%-noise bots and the sessions with fixed colors ($P$=0.675, log-rank test). Thus, the bots intervention, based on local decision-making alone, is equally as effective as a pre-calculated solution that (in typical circumstances) impractically would require prior global knowledge of the entire network structure and its solution space.

The bots appear to have improved collective performance in part by changing the color-conflict behaviors of human players in the whole system (Extended Data Fig.3). When placed at high-degree nodes, the bots with 0% behavioral randomness reduced the number of conflicts but they increased the duration of unresolvable conflicts; the bots with 30% randomness decreased the duration of unresolvable conflicts but increased overall conflicts; and only the bots with 10% randomness decreased both the number of conflicts and the duration of unresolvable conflicts, compared with the control sessions. In contrast, when placed at low-degree nodes, the bots were less likely to influence the entire network of humans, regardless of their noisiness.

When the bots were placed in high-degree positions, their behavioral randomness was able not only to facilitate the solution of their own conflicts, but also to nudge neighboring humans to change their behavior in ways that appear to have further facilitated a global solution. The bots with 0% behavioral randomness *reduced* the randomness of other human players (Fig.4a), which made the human players, particularly the middle-degree players, come to be stuck in unresolvable conflicts (Fig.4d). The bots with 30% behavioral randomness destabilized the entire network, including the low-degree players, who evinced more noise in their own actions (Fig.4c); as a result, the sessions with 30%-noise bots showed the same level of unresolvable conflicts as those without bots (Fig.4f). The bots with 10% behavioral randomness increased the randomness of the central players but reduced that of the peripheral players (Fig.4b); hence, through the influence of their behavioral randomness, the 10%-noise bots reduced the unresolvable conflicts not only of themselves but also of the entire network, including links between human subjects *unconnected* to the bots (Fig.4e). These results obtain even though the subjects were, in fact, less and less satisfied with their counterparts the noisier the bots were (Extended Data Fig.4).

In a separate, further experiment involving an additional 340 subjects and a matched set of N=20 graphs, we found that these beneficial effects on group coordination and learning obtained even when players knew they were interacting with bots (see SI). The solution time was statistically indistinguishable (Extended Data Fig.5) and the effect on players throughout the system was also similar (Extended Data Fig.6).

Adding autonomous agents with simple strategies to social systems may make it easier for groups of humans to achieve global optima for complex group-wide tasks. Here, the setting was a global coordination game, but other settings might include cooperation, sharing, or navigation[5,12,25]. The bots, however, might only be helpful if they have certain properties, including noisiness or particular geodesic locations. Like other situations[13,14,17,18,20], some noise may be good from the point of view of the group. Moreover, bots with some noise, with solely local information, improved global outcomes here just as much as bots employing global information acquired in advance.

We find that these slightly noisy bots work not only by making the task of humans to whom they are connected easier, but also by affecting the game play of the humans themselves when they interact with still other humans in the group, thus creating cascades of benefit. And this happens even when people know they are interacting with bots. In this sense, even simple artificial intelligence (AI) agents can serve a teaching function, changing the strategy of their human counterparts and modifying human-human interactions, and not just affecting human-bot interactions. More generally, our work illustrates the performance of combined, heterogeneous groups composed neither solely of humans nor solely of robots attempting to coordinate their actions. Future work can explore even more realistic or complex interactions, such as military or commercial robots working within human groups.

While laboratory experiments afford robust causal inference, they must sacrifice some verisimilitude and breadth. Guided by prior theory, we chose to focus on only two aspects of bot contributions (noise and placement) and their impact on one primary outcome (success of global coordination in a standard game[10]). We also necessarily made other design choices, including using a scale-free network limited to 20 people (which was required if the games were to be tractable). But there are other features of social interactions that might affect the ability of groups to coordinate to solve a problem, including group size, network topology[10], and bot fraction; whether the networks are dynamic or static[26,27]; or whether social institutions (e.g., policing, sanctions, or norms) are present. These elements are important directions for future work.

Adding bots of moderate noisiness to strategic positions within human networks might help address diverse problems, especially when the coordination problem is hard. For example, narrowly focused workers might each labor to enhance their own productivity, but this might actually decrease overall company performance. Crowd-sourcing applications in science (such as solving quantum problems[28] or other sorts of “citizen science” ranging from protein folding[29] to the assessment of archeological or astronomical images) might be facilitated by adding some bots or noise to groups working collaboratively. Moreover, our work reinforces the idea that both simple and sophisticated AI might be useful. For instance, simple bots might help reduce racist remarks online[30]. The simplicity and transparency of decision-making in simple AI might also make it intelligible to humans, thereby eliciting an effective, long-term relationship[11]. Simple autonomous agents, when mixed into complex social systems, might offer substantial advantages, and they could help groups of humans to help themselves.

## Methods

A total of 4,000 unique subjects (plus a further 340 for the secondary experiment regarding bot visibility – see SI) participated in our incentivized economic game experiments. They were recruited using Amazon Mechanical Turk (AMT; see SI), and they interacted anonymously over the Internet using customized software playable in a browser window (available at http://breadboard.yale.edu). While keeping other initial conditions the same, we completed 30 sessions for the only-human condition (control) and 20 sessions for each bot-treated condition (treatment). In each session (after passing various tutorials), the subjects were paid a $2 show-up fee and a declining bonus of up to $3 depending on speed to a global solution in which every player in a group had chosen a different color than their connected neighbors. When they did not reach a global solution within 5 minutes, the game was stopped and the subjects earned no bonus.

Except for the control group sessions, the networks had 3 bots in addition to 17 human subjects. These bots were assigned to three geodesic locations (peripheral, central, or random "locations"). The bots were controlled programmatically with a simple, greedy algorithm incorporating a random element; we drew a random number from a uniform distribution between 0.0 and 1.0; if the random number was less than a preset threshold ("behavioral noise"), the bot picked a color among the three color options at random; otherwise, it behaved based on the colors of its neighbors; if the bot's current color was not the least common among its neighbors, it changed to the least common color; otherwise, it maintained the current color.

To evaluate the difference in effectiveness between the various bot treatments, we analyzed the solution time of the N=180 sessions using Cox proportional hazard models. The sessions that were not solved within 300 seconds were regarded as censored. Each network session had a distinct level of complexity with respect to finding a coloring solution because it is generated *de novo*; thus, we controlled for the number of possible color combinations of the network (the "chromatic polynomial"). We also performed various statistical robustness checks (see SI).

We examined the impact of bots' behavioral noise on the humans' behavior using a generalized linear mixed model (GLMM) involving logistic regression (see SI). The dependent variable is the errant color-change rate evinced by the human players (i.e., choices that deviate from the simple, greedy strategy to minimize local conflicts). The model incorporated fixed effects for the behavioral noise of bots, the number of neighbors, the number of neighboring bots, the session length, and random effects for session.

## Extended Data

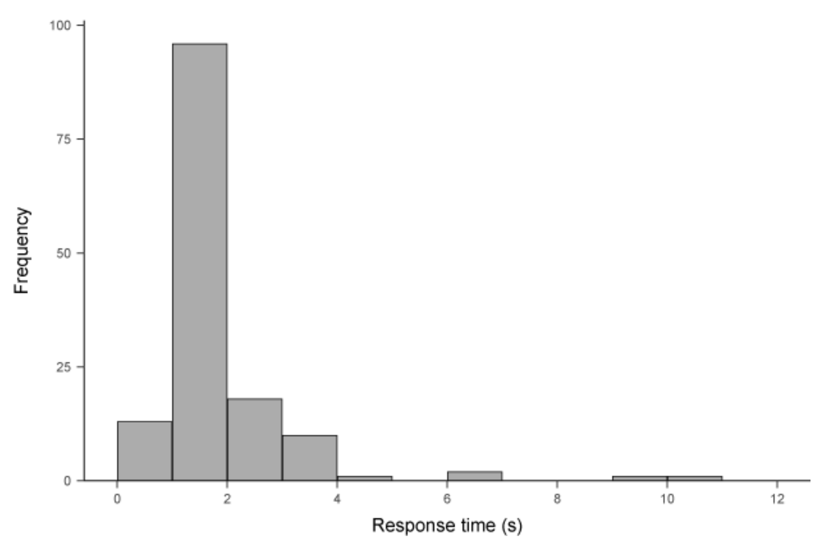


**Extended Data Figure 1. Histogram of the response time of humans in the color-matching test (*n*=142)**

In the color-matching test in our preliminary experiments, subjects were asked three times to click the same color button as a picture on the screen with five options: green, orange, purple, pink, and yellow. This histogram shows the response time (from when a color in question showed up on screen until when a subject clicked a button) for 142 pilot subjects. Most subjects clicked the correct button in 1.0 to 2.0 seconds (median time = 1.59 seconds).

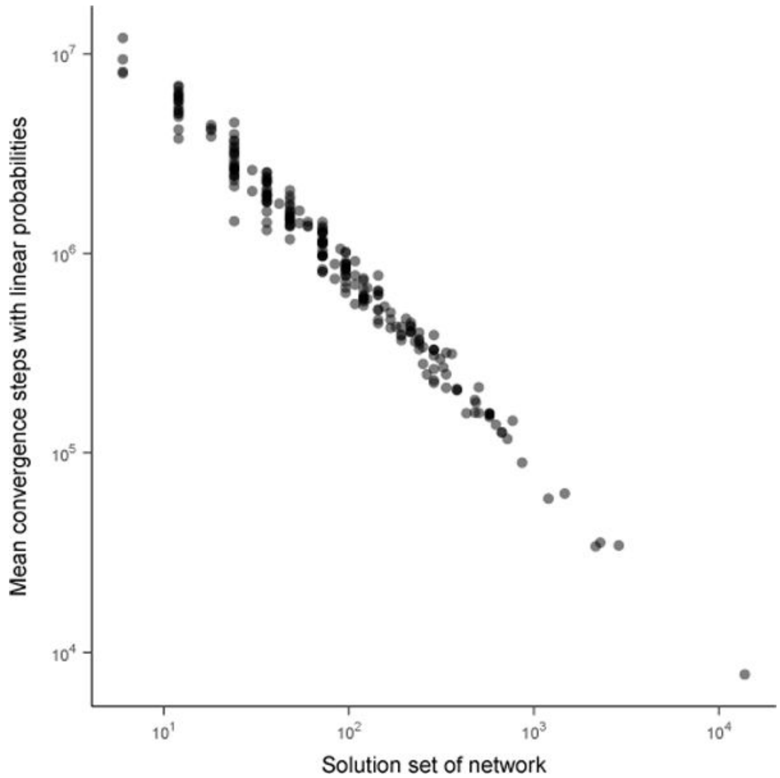


**Extended Data Figure 2. Relationship between different measures of the structure-based complexity of the graph coloring sessions**

The correlated coefficient after logarithmic transformation is −0.990 ($P < 0.001$; $n$=180). The solution set (x axis), known as the chromatic polynomial, is the number of possible color combinations that satisfy the task of coloring the network The linear probability algorithm (y axis) involves computing the following statistics: a node is randomly selected and changes its color to one that is different from its random neighbor until a solution is reached. This algorithm offers the advantage of allowing us to evaluate the landscape of the solution space starting from an arbitrary initial value. The "mean convergence steps" statistic was calculated for 100 iterations of each experimental network given the same initial coloring with the experiment.

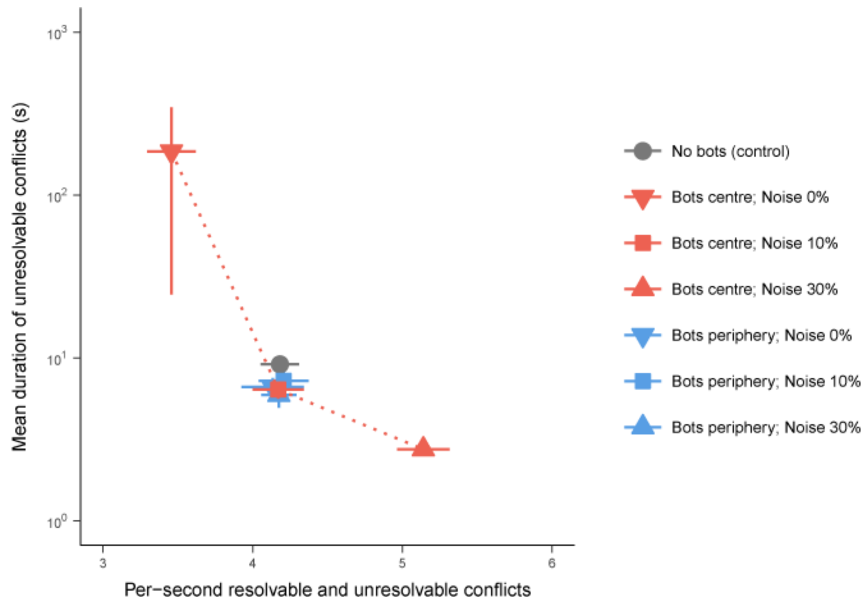


**Extended Data Figure 3. Impact of bots on color conflicts over the entire network**

The error bars are standard errors (*n*=30 for the no-bots sessions; *n*=20 for all the bots-treated sessions). When placed in the center, bots with 0% behavioral noise reduce the number of conflicts but increase the duration of unresolvable conflicts; bots with 30% noise decrease the duration of unresolvable conflicts but increase the overall conflicts; and bots with 10% noise decrease *both* the number of conflicts and the duration of unresolvable conflicts, compared with results of only human players. In contrast to central placement, when bots are placed in the periphery, conflict status does not vary with behavioral noise (data points are overlapping).

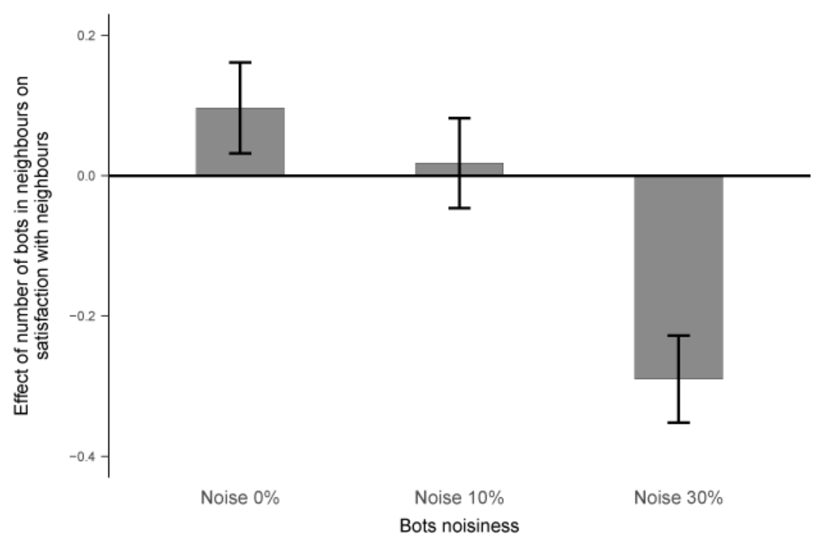


**Extended Data Figure 4. Impact of bots' behavioral noise on players' satisfaction with their neighbors**

After each session was completed, subjects rated their satisfaction with the actions of their neighbors on a five-point scale: very satisfied, satisfied, neither, dissatisfied, and very dissatisfied (the specific question asked was: “How satisfied were you with the actions of your neighbors you were connected with?”). These coefficients show the effect of number of bots among neighbors on their satisfaction with their neighbors, estimated by a proportional odds logistic regression, incorporating number of neighbors and whether the session was solved. The error bars are standard errors (*n*=3,035).

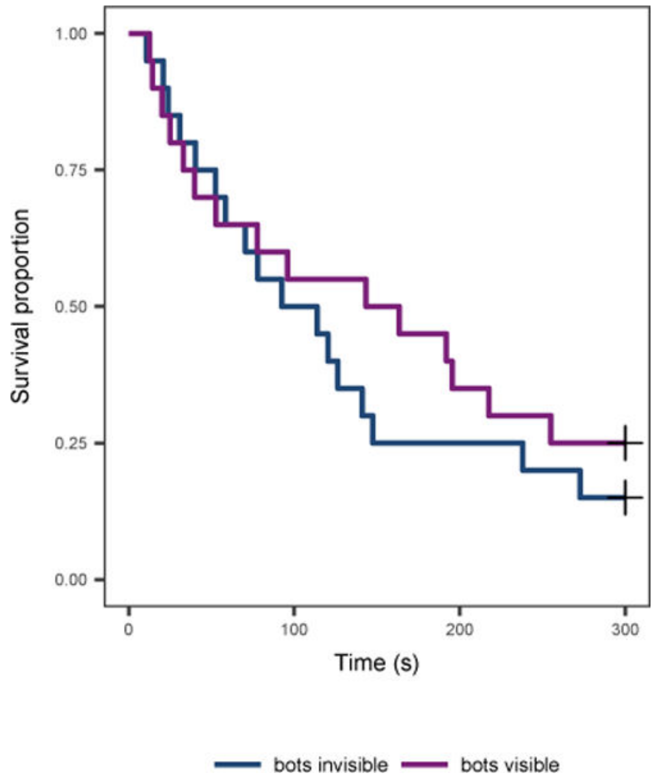


**Extended Data Figure 5. Survival curves for sessions by bot visibility**

The curves show the percentage of sessions unsolved at a given time. Dark blue lines show the *n*=20 sessions (involving *n*=340 additional subjects) where human players were informed of which nodes were played by bots (visible-bots condition; *n*=20), and light blue lines show the sessions where humans were not informed (invisible-bots condition; *n*=20). The difference of the survival curves is not statistically significant ($P = 0.435$, log-rank test).

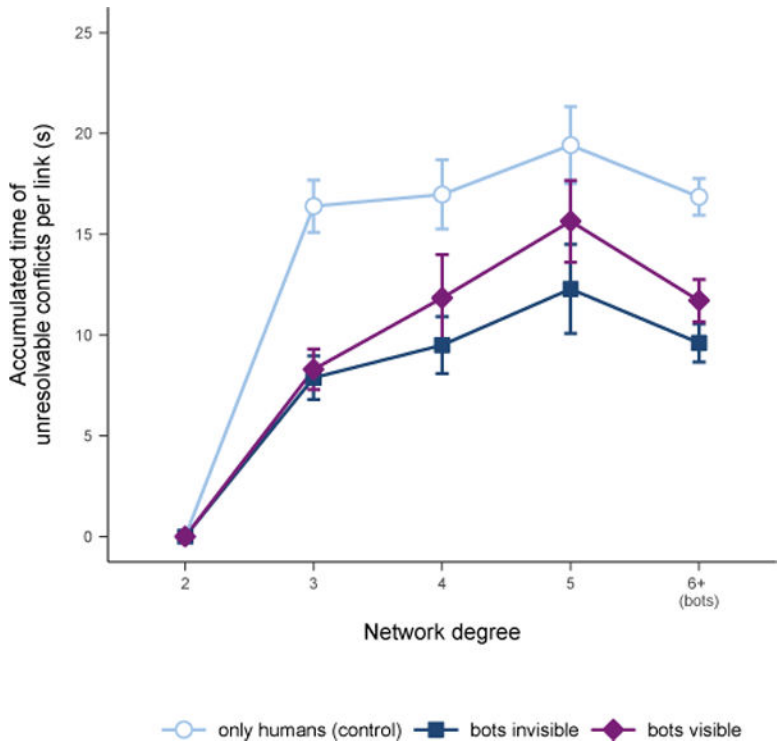


**Extended Data Figure 6. Impact of bot visibility on players' unresolvable conflicts for each geodesic location**

The dark purple line shows results for the sessions where human players were informed of which nodes were played by the bots (visible-bots condition; $n$=20), the dark blue line shows results from the sessions where humans were not informed (invisible-bots condition; $n$=20). In both conditions, the bots were located at high-degree nodes with 10% noise. The light blue line shows results for the sessions with all human players as a control ($n$=30). The error bars are standard errors by session. Except for the addition of the dark purple line (the results of the visible-bots condition), this figure is the same as Fig. 4e. Pertinently, the dark purple and dark blue lines are not statistically distinguishable, suggesting that making the bots visible has a similar effect throughout the network on players' behavior compared to keeping them invisible.

## Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

## Acknowledgments

We thank P. Allison, F. Fu, G. Kraft-Todd, A. Oswald, D. Rand, and D. Spielman for helpful comments. Mark McKnight provided technical support and programming for our Breadboard platform. Support for this research was provided by grants from the Robert Wood Johnson Foundation, the National Institute of Social Sciences, and the National Institutes of Health (P30-AG034420).

## References

1. Hardin G. The tragedy of the commons. Science. 1968; 162:1243–1248.
2. Dawes RM. Social Dilemmas. Annu Rev Psychol. 1980; 31:169–193.

3. Tavoni A, Dannenberg A, Kallis G. Inequality, communication, and the avoidance of disastrous climate change in a public goods game. Proc Natl Acad Sci USA. 2011; 108:11825–11829. [PubMed: 21730154]
4. Calvert R. Leadership and its basis in problems of social coordination. International Political Science Review. 1992; 13:7–24.
5. Dyer JRG, Johansson A, Helbing D, Couzin ID, Krause J. Leadership, consensus decision making and collective behaviour in humans. Philosophical Transactions of the Royal Society B: Biological Sciences. 2009; 364:781–789.
6. Van Huyck JB, Battalio RC, Beil RO. Tacit coordination games, strategic uncertainty, and coordination failure. The American Economic Review. 1990; 80:234–248.
7. Nowak M. Stochastic strategies in the prisoner's dilemma. Theoretical Population Biology. 1990; 38:93–112.
8. Kandori M, Mailath GJ, Rob R. Learning, mutation, and long run equilibria in games. Econometrica. 1993; 61:29–56.
9. Young HP. Learning by trial and error. Games and Economic Behavior. 2009; 65:626–643.
10. Kearns M, Suri S, Montfort N. An Experimental Study of the Coloring Problem on Human Subject Networks. Science. 2006; 313:824–827. [PubMed: 16902134]
11. Axelrod, R. The evolution of cooperation. Basic Books; 1984.
12. Rand DG, Nowak MA. Human cooperation. Trends in Cognitive Sciences. 2013; 17:413–425. [PubMed: 23856025]
13. Jiang JJ, Huang ZG, Huang L, Liu H, Lai YC. Directed dynamical influence is more detectable with noise. Sci Rep. 2016; 6:1–9. [PubMed: 28442746]
14. Eldar A, Elowitz MB. Functional roles for noise in genetic circuits. Nature. 2010; 467:167–173. [PubMed: 20829787]
15. Couzin ID, et al. Uninformed Individuals Promote Democratic Consensus in Animal Groups. Science. 2011; 334:1578–1581. [PubMed: 22174256]
16. Gao J, Barzel B, Barabási AL. Universal resilience patterns in complex networks. Nature. 2016; 530:307–312. [PubMed: 26887493]
17. Sniegowski PD, Gerrish PJ, Lenski RE. Evolution of high mutation rates in experimental populations of E-coli. Nature. 1997; 387:703–705. [PubMed: 9192894]
18. Kirkpatrick S, Vecchi MP. Optimization by simulated annealing. Science. 1983; 220:671–680. [PubMed: 17813860]
19. Kadri U, Brümmer F, Kadri A. Random patterns in fish schooling enhance alertness: a hydrodynamic perspective. EPL (Europhysics Letters). 2016; 116:1–6.
20. Traulsen A, Semmann D, Sommerfeld RD, Krambeck HJ, Milinski M. Human strategy updating in evolutionary games. Proc Natl Acad Sci USA. 2010; 107:2962–2966. [PubMed: 20142470]
21. Helbing D, Yu W. The outbreak of cooperation among success-driven individuals under noisy conditions. Proc Natl Acad Sci USA. 2009; 106:3680–3685. [PubMed: 19237576]
22. Couzin ID, Krause J, Franks NR, Levin SA. Effective leadership and decision-making in animal groups on the move. Nature. 2005; 433:513–517. [PubMed: 15690039]
23. Liu YY, Slotine JJ, Barabási A-Á. Controllability of complex networks. Nature. 2011; 473:167–173. [PubMed: 21562557]
24. Barabasi AL, Albert R. Emergence of scaling in random networks. Science. 1999; 286:509–512. [PubMed: 10521342]
25. Kesting A, Treiber M, Schönhof M, Helbing D. Adaptive cruise control design for active congestion avoidance. Transportation Research Part C: Emerging Technologies. 2008; 16:668–683.
26. Rand DG, Arbesman S, Christakis NA. Dynamic social networks promote cooperation in experiments with humans. Proc Natl Acad Sci USA. 2011; 108:19193–19198. [PubMed: 22084103]
27. Shirado H, Fu F, Fowler JH, Christakis NA. Quality versus quantity of social ties in experimental cooperative networks. Nature Communications. 2013; 4:1–8.

28. Sørensen JJWH, et al. Exploring the quantum speed limit with computer games. Nature. 2016; 532:210–213. [PubMed: 27075097]

29. Cooper S, et al. Predicting protein structures with a multiplayer online game. Nature. 2010; 466:756–760. [PubMed: 20686574]

30. Munger K. Tweetment Effects on the Tweeted: Experimentally Reducing Racist Harassment. Political Behavior. 2016:1–21.

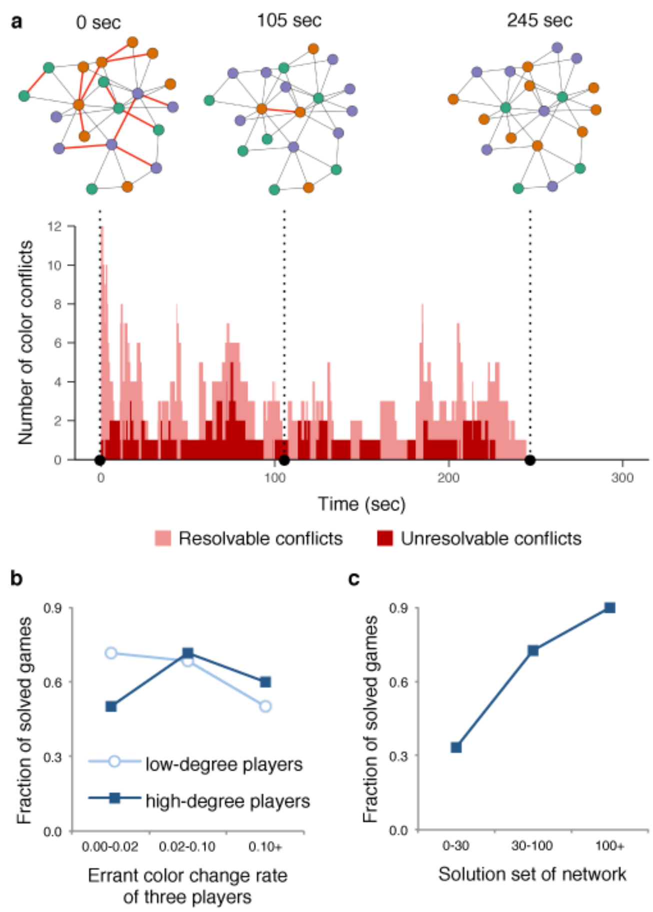


**Figure 1. Results of sessions involving only human players**

(a) An example of the color coordination game. The figures are snapshots with players' node color at 0, 105, and 245 seconds (see Video S1 for full version). Red edges show that the connected players are the same color ("color conflicts"). Some conflicts can be resolved when either player selects the rarest color among his/her neighbors ("resolvable conflicts"); but others cannot ("unresolvable conflicts"). (b) The actual fraction of solved games depending on the behavior of the most central or peripheral three players is shown. The "errant color change rate" is the ratio of color selections (by the subjects) producing *more* color conflicts divided by the opportunities to make such selections (see SI for details). An intermediate level of errant color choice among high-degree human players resulted in the greatest solvability (which comports with the programming strategy for helpful bots). (c) The actual fraction of solved games in relation to the number of possible color combinations (the "chromatic polynomial") is shown; having more possible solutions is associated with a higher solution rate.

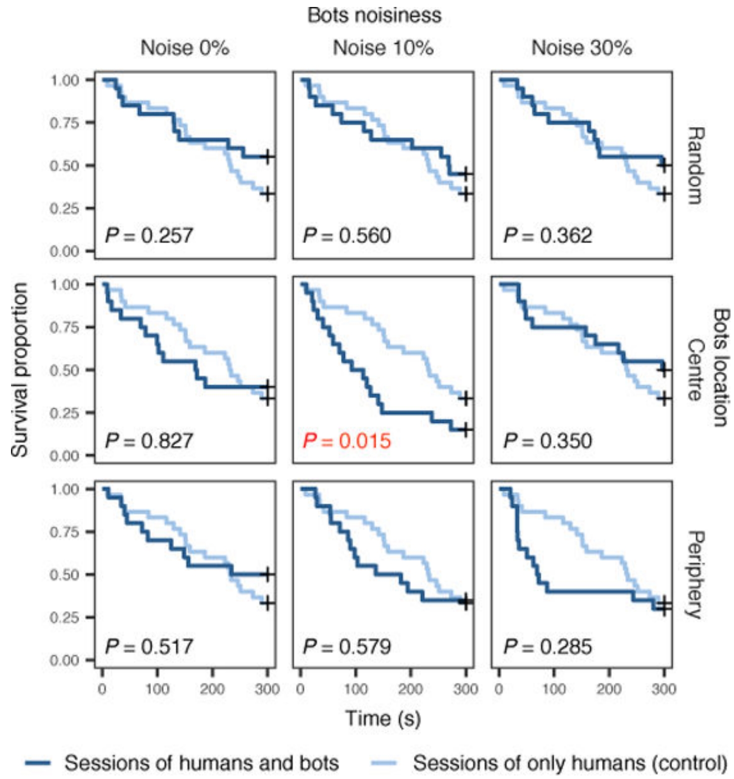


**Figure 2. Survival curves of sessions, by noisiness and location of bots**

The curves show the percentage of sessions unsolved at a given time. Dark blue lines show results for the sessions including bots (*n*=20), by their noise level (horizontal dimension) and geodesic location (vertical dimension). Light blue curves show results for the control sessions involving solely human players (*n*=30). Total *n*=210. Sessions are censored at 300 seconds; *P*-values given by the log-rank test. Bots having 10% behavioral noise and located at the center of the network cause a significant improvement in the solvability of the game (*P*=0.015) and induce 55.6% acceleration in the median time to solution, from 232.4 seconds to 103.1 seconds.

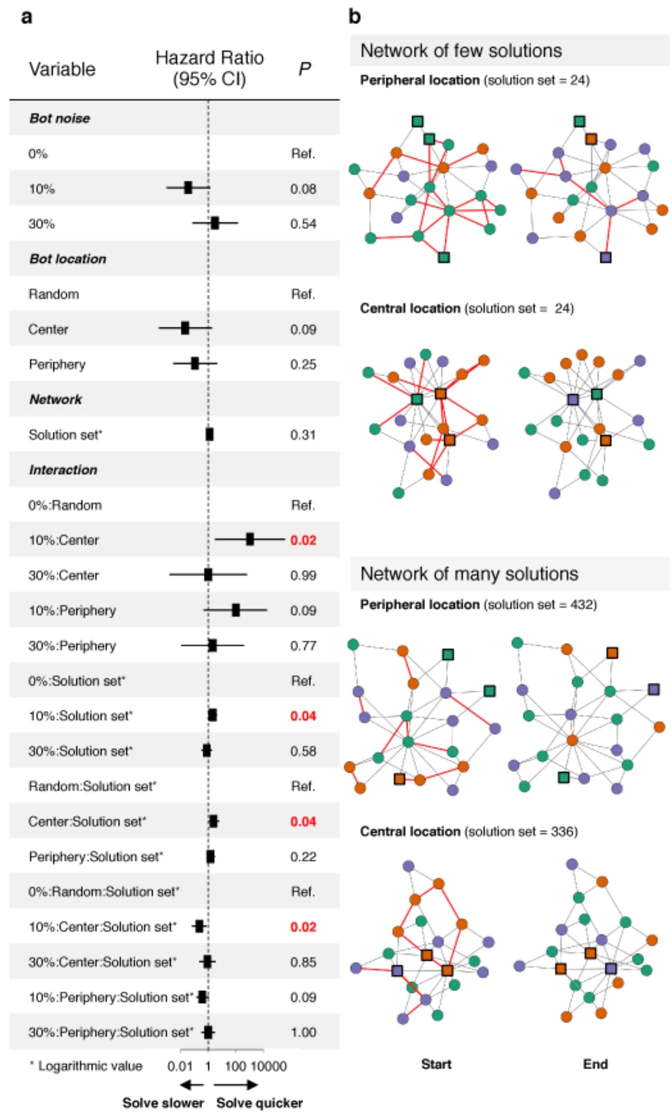


**Figure 3. Results of the survival analysis by bot and network characteristics**

(a) Hazard ratios for game solution time according to bot noise, bot location, number of solutions of the network (chromatic polynomial), and all interactions among these variables ($n$=180; see Table S3 for details). The results show that the benefit of bots varies with the solution space; when a network has few possible color combinations, placing bots in a central location (high-degree nodes) facilitates resolution. (b) These network snapshots show initial and final states of illustrative sessions involving bots with 10% noise. Square nodes show the bots, and round nodes show human players; red edges show color conflicts.

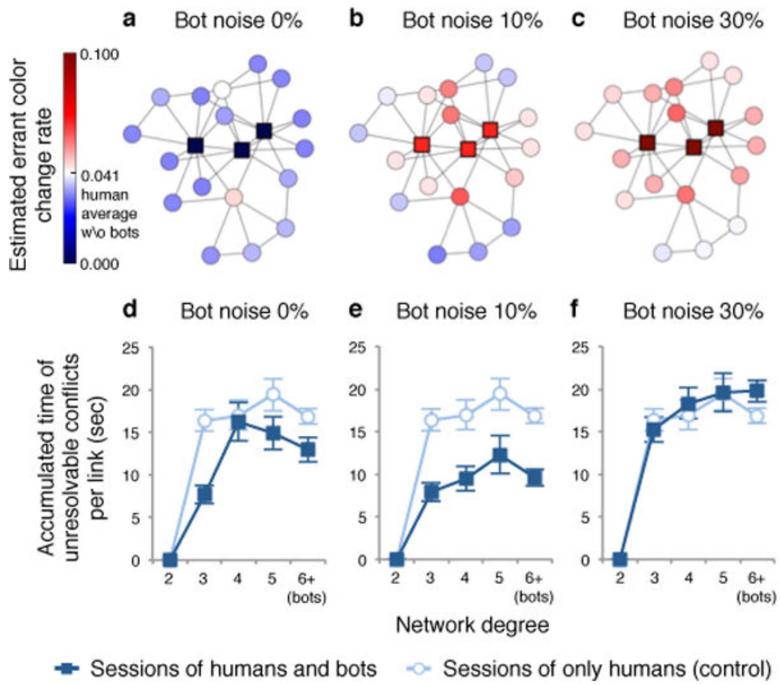


**Figure 4. Impact of bots on the behavior of human players**

(a–c) Snapshots show estimates of the errant color change rate (i.e., humans choosing “wrong” colors) in the same network with central bots, depending on bot noise. Square nodes show bots and round nodes show humans (see SI and Table S5 for regression modeling details). Note that the intermediate white color shows the estimated errant rate of average human players in sessions without bots (0.041); thus, the red color shows that human players behave in a more noisy way as a result of the influence of the bots; the blue color shows the opposite. (d–f) These graphs show the average accumulated time of unresolvable conflicts per link for each geodesic location of players. Dark blue lines show results for sessions with central bots (whose degree was typically ≥ 6) by their noise level, and light blue lines show results for the control sessions with only humans. Bots with 10% noise change the behaviors of the human players in the whole system for the better (e). Error bars are standard errors.