Network Dynamics of Equilibrium Selection in Coordination Decisions

Joshua Becker

Annenberg School for Communication, University of Pennsylvania

joshua.becker@asc.upenn.edu

www.joshua-becker.com

WORKING PAPER

September 15, 2017

Abstract

What social structures promote optimal coordination? Previous work has shown that centralized networks can be advantageous when speed is a priority, because central individuals can coordinate the actions of the group as a whole. However, when multiple strategies compete for adoption, the fastest process may not select the optimal strategy. I present theoretical results showing that centralization increases speed of convergence, but decreases the probability of optimal equilibrium selection in pure coordination games. However, speed is not inherently problematic. In dense networks, coordination is both fast and optimal. The deleterious effects of network centralization are explained instead by the influence of central nodes. Solutions introduced by a central node are more likely to spread regardless of quality, at the expense of less popular, higher payoff solutions.

Keywords

collective intelligence, coordination, social networks, network dynamics, agent based model

Introduction

Social behaviors are shaped by coordination dynamics when decision-making agents have an incentive to make the same choice as their peers. One key example is the adoption of communication technologies, where a given tool may hold no value whatsoever if it is not compatible with the technologies employed by those with whom an agent wishes to communicate. The agent facing a coordination challenge could be an organization, such as a firm choosing a digital communication protocol (Pfieffer, 2012), or an individual, such as a person choosing a mobile messaging platform (Heal & Keunreuther, 2010).

When different strategies offer different payoffs, decision-making agents may have to choose between a popular but low payoff solution - which ensures compatibility with peers - and a less popular, but more preferred solution. The tradeoff between popularity and preference is most salient when many alternative strategies emerge simultaneously. In this scenario, agents may have limited knowledge of the available alternatives, and can thus become invested in an inferior solution before even learning about a less popular but superior option. For example, prior to the establishment of widespread digital communication standards, firms were required to establish conventions directly with industry partners, based on the knowledge of practices within their immediate inter-organizational network (Pfieffer, 2012).

If it were the case that every agent had full knowledge of the solution space from the outset, coordination problems would be trivial: everybody would simply adopt the best solution. However, when agents only become aware of strategies through peer observation, an inferior strategy can quickly become popular and achieve "lock-in" effects. Lock-in effects occur once a behavior with positive externalities achieves sufficiently widespread adoption that alternative strategies will not be adopted even if they offer a clear advantage (Arthur, 1989). Such lock-in effects have been observed historically in the case of technologies such as the QWERTY keyboard (David, 1985), alternating current (David & Bunn, 1987) and VHS tapes (Besen & Farrell, 1994). The QWERTY keyboard stands out for the numerous problems associated with its design, and the numerous failed attempts to introduce superior alternatives (David, 1985).

While coordination models make the formal assumption that alternative behaviors offer a clearly defined payoff, these models nonetheless describe social conventions more generally, where the payoff may not be literally quantifiable (Lewis, 2002 [1969]; Bicchieri, 2005). In addition to characterizing the emergence of social conventions (Centola & Baronchelli, 2015) coordination dynamics also shape group problem solving when there is pressure to reach agreement with peers (Bettenhausen & Murnigan, 1985). In problem solving as well as social conventions, a person may face pressure to adopt some strategy despite having knowledge of a better alternative, due to the need to reach agreement with others.

Centralization

Prominent members of a social network, or "opinion leaders," are a key factor in the formation of collective decisions (Katz & Lazarsfeld, 1955) and play an influential role in the adoption of technological and social innovations (Rogers, 1983 [1962]; Banerjee, Chandrasekhar, Duflo, & Jackson, 2015). When a behavior spreads unchallenged – e.g., a single superior innovation with no alternative – the influence of opinion leaders can be helpful by increasing adoption rates of beneficial technologies, such as health practices (Rogers, 1983). However, when multiple solutions compete for adoption, and agents have a limited knowledge of the solution space, then an opinion leader may inadvertently spread a behavior that is not an optimal strategy. An opinion leader need not be aware of their influence or intentionally promote an agenda in order to lead collective beliefs. Simply by virtue of their greater observability, the individual behavior of opinion leaders can shape collective dynamics due to their influence on a larger share of the population (DeMarzo et al, 2003; Banerjee, Chandrasekhar, Duflo, & Jackson, 2015; Becker, Brackbill, & Centola, 2017). Thus, it is possible that the presence of central nodes is sufficient to reduce the optimality of equilibrium selection, as compared to a network in which influence is equally distributed throughout the population.

Laboratory research on group problem solving has found that centralized networks solve problems faster and with fewer errors, due to the role of central individuals and their positive effect on information exchange (Leavitt, 1951; Bavelas, 1953; Mulder, 1960). Mulder (1960) further argues that centralization is not only a topological feature of information flow, but also can be measured in terms of the decision structure, with regard to who follows whom. In this respect, centralization improves the speed of group decisions due to the fact that “central individuals can integrate the contributions of all individuals” (Mulder, 1960: 12). In other words, central nodes have the potential to act as key coordination agents, a role made possible by their greater knowledge of the information circulating in the network.

These prior studies focus largely on speed of problem solving. However, when some solutions offer a higher payoff than others, speed may not be the only factor of importance, and centralization has the potential to undermine optimal coordination. In order to study the effect of central individuals on coordination dynamics, I use a computational experiment (Macy & Willer, 2002) to compare the probability of optimal coordination in networks characterized by the presence of opinion leaders, which are termed “centralized” networks (Freeman, 1978) against the probability of optimal coordination in networks where everyone is equally connected, or “decentralized” networks.

In contrast with theoretical and empirical research highlighting the beneficial effect of opinion leaders (Katz & Lazarsfeld, 1955; Leavitt, 1951; Mulder, 1960; Rogers, 1983; Banerjee et al, 2013) I find that central individuals can undermine the optimality of collective coordination by reducing the likelihood of optimal equilibrium selection. The deleterious effects of opinion leaders occur when a central individual in a network adopts an inferior strategy, causing that solution to become widely adopted before less popular but higher payoff solutions have a chance to spread.

Prior Analytical Work

Game theoretic analyses of coordination dynamics seek to identify a single equilibrium solution, with the goal of demonstrating that the equilibrium solution will become the dominant strategy in a population. A common approach has been to identify the evolutionary stable equilibrium (Kandori, Mailath, and Rob, 1993) or stochastically stable equilibrium (Young, 1993). A number of analyses using these and related approaches have supported the optimistic conclusion that the payoff-dominant strategy[1] (i.e., the optimal solution) will always emerge as the shared convention in a population (Kandori et al, 1993; Young, 1993; Blume, 1995; Ellison, 2000; Montaneri & Saberi, 2010).

A key characteristic of these analyses is that they require the possibility that a population which has already reached a shared convention can, by chance, switch to another convention, even in the absence of an exogeneous shock (Kandori et al, 1993; Young, 1993). This assumption underlies an important conceptual limitation of these analyses: while they show that the optimal solution is the most likely equilibrium, they do not guarantee that the optimal solution will be the *first* equilibrium reached. As Young (1993) points out, in a population that begins with many different strategies, any potential solution may emerge as the dominant convention.

Even showing that optimal equilibria will emerge asymptotically does not guarantee that the optimal solution will be adopted in reasonable amount of time. Work by Ellison (1993, 2000) highlights two important properties of coordination dynamics: first, suboptimal equilibria can remain stable in networks for non-trivial periods of time; and second, network structure is a key factor shaping coordination dynamics. Ellison demonstrated that in networks characterized by

---

[1]These analyses specifically identify the risk-dominant strategy as the equilibrium of a dynamic system. In pure coordination model examined here (ie, the payoff of miscoordination is constant) the risk-dominant strategy is the payoff dominant strategy.

high local clustering (so that the neighborhoods of adjacent nodes overlap) the optimal solution will spread quickly due to cascading effects, even after a suboptimal solution has been widely adopted. In contrast, it is much more difficult for the optimal solution to replace an established equilibrium in disordered networks with less clustering (random graphs) or networks characterized by global ties (fully connected networks) once an inferior solution has been widely adopted (Ellison 1993; Ellison 2000). This means that an inferior solution can remain the stable equilibrium for long periods of time in all but the most geometrically regular networks.

Once a solution is widely adopted by a population, the self-reinforcing effects of adoption with coordination externalities mean that it is possible for even an inferior convention to remain entrenched indefinitely – nobody has an incentive to deviate from convention. Because any equilibrium may become locked-in indefinitely, it is crucial to identify not only the stochastically stable equilibrium, but also the *first* equilibrium to be reached. A Markov model of coordination with lock-in effects means that any potential solution can be an absorbing state, and may not be analytically tractable (Young, 1993) – the assumption that equilibria can be exited plays a key role in the use of statistical tools for analyzing coordination dynamics (Kandori et al., 1993; Young, 1993). However, by simulating standard models of coordination with absorbing state equilibria, I examine how variation in social structure can make a population more or less likely to select the optimal strategy as the first (and possibly final) equilibrium.

Modeling Coordination in Networks

A coordination problem can be framed in the following way: a population of decision-making agents must frequently perform some behavior, such as getting information from a business partner or sending them a message. They must choose one strategy among many, as for example a particular choice of digital communication technology. Each agent chooses a strategy by combining two piece of information: the payoff of each strategy (ie, their preference) and the number of network neighbors using that particular strategy. Agents choose the strategy that offers the highest expected payoff (the "best response" strategy).

Formally, I study an MxM pure coordination game (ie, with *M* possible strategies) in which a strategy S_i offers payoff V_i upon successful coordination, and a payoff of 0 upon miscoordination. Thus, the expected payoff for a given strategy S_i is simply equal to

(proportion of neighbors using S_i) x (payoff of S_i)

which follows the standard calculation for expected payoff. Number of neighbors can be equivalently substituted for proportion of neighbors, and this expected payoff function offers several possible interpretations in the context of coordination behavior.

In a context where agents play a series of pairwise (two player) coordination games, then the proportion of observable peers employing a given strategy represents the probability of success on a given pairwise interaction. Under one interpretation, an agent interacts only with their network neighbors. For example, a firm choosing a communication protocol must attempt to match as many possible peers in their immediate business network. In another interpretation, agents may have some chance of interacting outside their network neighborhood, but use observations of their immediate network neighborhood as a heuristic to form beliefs about the population at large. This second example describes the following scenario: if a minority of a person's friends use social network app A, and a majority use social network app B, then app B will provide the agent with a higher probability of being able to connect with a random stranger that they might meet at a conference or other social interaction. However, the agent may prefer social network app A, due to its advanced features or low cost, generating a trade-off between popularity and preference.

I follow previous work (Kandori et al, 1993; Blume, 1993; Young, 1993; Ellison, 1993; Montanari & Saberi, 2010) in using a boundedly rational model of coordination behavior with two key assumptions. First, the model assumes that an agent's best response decision depends myopically on their immediate social environment. Second, the model assumes that agents have some chance of error, such that they do not always make a best response decision. The assumption of myopic choice is reflected in the formula for expected payoff, which is based only on the current strategies of an agent's observable peers as determined by the social network.

The second modeling assumption, the presence of statistical noise in agent decisions, is a key assumption of bounded rationality: the expected payoff function identifies an unambiguously preferred strategy (except in the case of ties) but we cannot assume that a decision-making agent will always choose the rationally preferred strategy. The introduction of statistical noise into agent behavior is a critical factor in rendering coordination dynamics analytically tractable, because it allows for the analysis of the ergodic properties of stochastic Markov chains (Young, 1993; Montaneri & Saberi, 2010) but I find that results in simulation are comparable in both noisy and non-noisy conditions. Overall, the model employed here is nearly identical to the several models used in previous coordination research (Kandori et al, 1993; Young, 1993; Blume, 1995; Montanari & Saberi, 2010).

Model Definition

Each run of the simulation is generated as follows:

- N agents are embedded in an undirected, binary network

- Initially, each agent is randomly assigned a strategy from a set of M possible strategies.
- At each time step, an agent is randomly selected from the set of agents whose strategy is not already the best response to their social environment. That agent then adopts the best response strategy with probability (1-ε). With probability ε, the agent randomly selects a strategy from the set of strategies still in use throughout the population.
- The model is run until all agents are employing a best-response strategy – ie, until the system has reached equilibrium.

In previous implementations of this model (Kandori et al, 1993; Blume, 1993; Young, 1993; Ellison, 1993; Montanari & Saberi, 2010) agents who "err" (via the noise term ε) randomly select a strategy from all possible strategies. However, the goal of the present work is to model the emergence of lock-in effects. By modeling agents that select only from strategies currently in use throughout the population, this model allows for unused strategies to "die out" as they are no longer used by any agent. By implementing noise in this way, the model reaches an absorbing state – when every agent uses the same strategy, no agent will ever "err" by adopting an alternative strategy, and the population cannot exit equilibrium.

It is worth noting that a system can reach an equilibrium where multiple strategies are in use throughout a population (Blume, 1995). However, such outcomes are rare, and the networks presented here reach on a shared convention in more than 99% of simulations.

Model Parameters

The main behavior of interest is the case in which a population must select an equilibrium from a large number of possible strategies. To study these dynamics, the main results presented here model the case in which M >> N (i.e., many more solutions [M] than agents [N]), so that each agent starts with a unique strategy. I also show outcomes where M ≤ N, which produces similar results, though the deleterious effects become weaker as the number of solutions approaches 2.

The payoffs for each of the M strategies are independently, identically distributed according to a fixed distribution. In the figures presented here, the payoff for a given strategy is drawn from a log-normal distribution (μ=0, σ=1). The log-normal distribution is chosen due to the fact that innovations often follow a long-tailed payoff distribution, with many mediocre solutions and few high quality solutions (Kauffman, 1993). I show results in the Appendix for a uniform distribution, which produces similar outcomes.

Results are presented with ε=0.1, meaning that 1 in 10 decisions is random, and 9 in 10 decisions follow the best response strategy. As shown in the Appendix, the effect of centralization on coordination does not depend on noise – results are identical when ε=0, such that agents follow a perfect best response strategy. For the sake of generality and consistency with previous research, I present results with noise in the main text.

Network Structure

A network is considered "centralized" when there is a large amount of inequality in the distribution of connectivity. In a highly centralized network, one or a small number of nodes have a large number of connections, while most nodes have only a few connections (Freeman, 1978). To quantify the level of centralization in a network I use the Gini coefficient[2], which is a widely used metric for quantifying the level of inequality in a particular resource distribution. In the case of

[2] Another common metric is the Freeman centralization score (Freeman, 1978), which is less sensitive to variation in network structure (Badham, 2013).

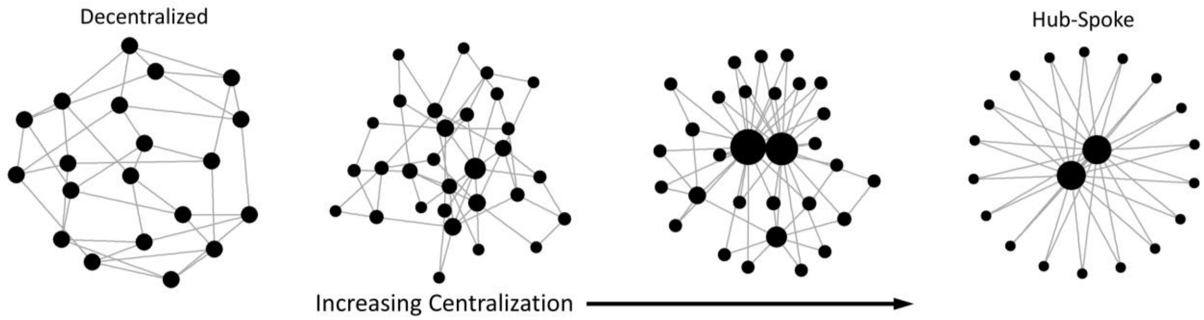

Figure 1. *Schematic showing increasingly centralized networks.* All networks in this figure have the same graph density, with an average of k=4 connections per node. The leftmost graph in this figure is a decentralized random graph, where every node has exactly 4 connections. The rightmost graph is a hub-spoke network, with $\frac{k}{2}$ central nodes and N-2 peripheral nodes. For the centralized networks used to generate the main results, every node has an average of k=20 connections per node, with N=1000 nodes per network.

networks, the Gini coefficient is measured for the degree distribution (Badham, 2013), where each node's "degree" is defined as the number of connections they have to other nodes (Easley & Kleinberg, 2010). Figure 1 shows a schematic representing networks of varying centralization.

To produce centralized networks, I draw on two network generating algorithms. In order to generate highly centralized networks, I employ the well-known preferential attachment mechanism identified by Barabasi and Albert (1999) to explain the emergence of scaling in empirical networks. When the parameters are appropriately tuned, this algorithm can generate networks ranging from moderately centralized random networks (gini=0.25) to highly centralized "hub spoke" networks (gini=0.5), in which a single core group of nodes is connected to every other node, while peripheral nodes are connected only to the core nodes. However, this algorithm cannot be tuned to generate fully decentralized networks (gini=0). Therefore, in order to study networks that vary continuously between fully decentralized and highly centralized networks, I also employ a second network algorithm. This network generator begins with a random decentralized network (gini=0), and employs a preferential rewiring mechanism to generate fully decentralized to moderately centralized networks (gini=0.2) (see Appendix for full definition). Both of these algorithms allow me to hold network density and average degree constant while increasing the level of centralization.

One challenge in modeling network effects is that it is difficult to change one parameter of a network while holding all other parameters constant. In particular, as networks become increasingly centralized, the average path length between two randomly selected nodes decreases. In diffusion models, this results in a more efficient spread of information (Watts & Strogatz, 1998) and therefore it is challenging to distinguish the effects of centralization (presence of more prominent nodes) from the effects of communication efficiency (which may increase the speed of coordination).

In order to disentangle the effects of increasing coordination from the effects of network efficiency, I also compare coordination in fully decentralized networks over a range of density, up to fully connected networks in which every agent can observe every other agent. As density increases, average path length decreases, which has been shown to increase speed of convergence in related models of coordination where every solution offers equal payoff (Dall'Asta et al, 2006).

The Network Dynamics of Equilibrium Selection

In order to compare my results to previous theoretical research on coordination, I begin with a discussion of coordination in clustered lattice graphs, which nearly always select the optimal equilibrium. I then discuss the effect of centralization on the probability of optimal coordination. In order to test whether the effects of centralization can be explained as a result of convergence speed, I also test the effect of network efficiency by varying the density of decentralized networks. I close by discussing the robustness of the main results against variation in modeling assumptions.

In empirical studies of social networks, social structure has been frequently observed to display two key characteristics. First, networks tend to be highly centralized – there are a few well connected individuals, and many poorly connected individuals (Barabasi and Albert, 1999) an observation which has been replicated in a number of empirical studies (Liljeros et al, 2001; Ebel et al, 2002; Barabasi et al, 2002; Kossinetts & Watts, 2006; Eagle & Pentland, 2006). Second, networks are characterized by the presence of long distance ties which disrupt the geometrically regular, highly clustered structure of a lattice graph (Watts & Strogatz, 1998). This second factor – the presence of long distance ties – is particularly relevant in the study of coordination dynamics.

Previous theoretical work (Ellison, 1993; Blume, 1995) has shown that in lattice graphs, the overlapping structure of node neighborhoods makes it easy for the optimal solution to spread quickly, provided that it has sufficient early adoption to take hold. Conditions to ensure this cascading effect are minimal – all it takes is for a single neighborhood (one node and all their contacts) to adopt the strategy, and then it is nearly guaranteed to spread through an entire population (Blume, 1995). In the model studied here, these initial conditions are almost guaranteed to occur. In the case where every node begins with a unique solution, there will of course be one node employing the optimal solution. (I define “optimal solution” as the best solution used by any agent in a given simulation.) The neighbors of this node will quickly adopt that superior strategy, because no other strategy yet offers a popularity advantage, and the optimal solution will spread through the network. This expectation is confirmed by simulation, in which lattice graphs converge on the optimal solution in more than 99% of trials.

However, empirical networks do not display the geometric regularity of lattice graphs (Watts & Strogatz, 1998) and random graphs are more prone to lock-in effects than lattice graphs (Ellison, 2000). Even a small number of long distance ties can disrupt cascading effects (Centola

& Macy, 2007) and coordination dynamics become increasingly less likely to break out of suboptimal equilibria as networks become more disordered (Montanari & Saberi, 2010). These properties indicate that in disordered networks, the first equilibrium to be reached is an important one – but previous results (Kandori et al, 1993; Young, 1993; Montaneri & Saberi, 2010) cannot guarantee that the first equilibrium will be the optimal solution. Instead, I find that it is quite common for random networks to converge on a suboptimal equilibrium. However, not all networks are created equal - the probability of optimal equilibrium selection varies with both network centralization and network density.

The Effect of Network Centralization

Figure 2 (left) shows the probability of optimal coordination in networks with increasing centralization. All the points in this chart show outcomes for networks with an average of 20 degrees per node, and thus network density (0.02) is held constant as centralization increases. In decentralized networks with 20 connections per node, the population selects an optimal strategy in about 75% of simulations. As centralization is increased, however, the probability of optimal equilibrium selection decreases dramatically. In the most centralized networks, which are hub-spoke networks with a core hub of 10 nodes, the optimal solution is selected in only about 15% of simulated trials.

One possibility is that this reduced optimality is due to speed of convergence. Increasing centralization reduces average path length, bringing every node in the network closer together. One effect of this network efficiency is that centralization increases the speed of convergence. Figure 2 (right) shows that the number of updates per agent needed to reach convergence decreases along with the probability of optimal coordination. Thus, it may simply be that the network is

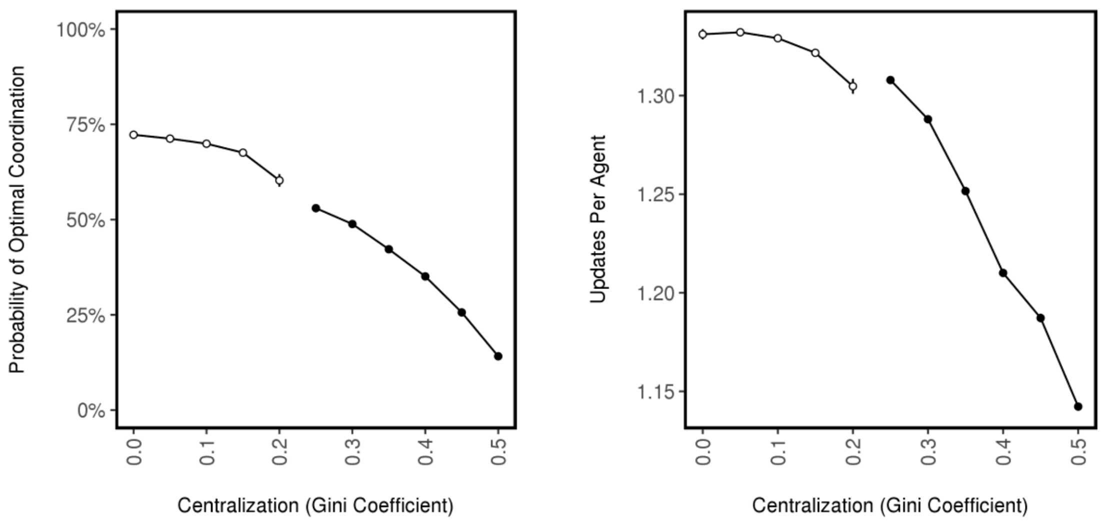

Figure 2. *The effect of network centralization on the probability of optimal coordination.* LEFT: As network centralization increases, the probability of optimal coordination decreases. RIGHT: As network centralization increases, the average number of updates per agent decreases. BOTH PANELS: To obtain the full range of centralization, I employ two network generators, a preferential rewiring algorithm (open circles) and a preferential attachment algorithm (solid circles). I cannot directly control centralization in network generation, and outcomes are binned according to their resulting centralization, rounded to the nearest 0.05. Each point shows at least 3,000 simulations. 95% confidence intervals are drawn, but too small to be visible on some points.

converging more quickly on *any* solution, which decreases optimality as the equilibrium selection becomes noisier.

To further examine the effect of network centralization, I examine the payoff of equilibrium solutions as centralization increases. Figure 3 (solid grey line) shows the average payoff of the equilibrium solution as a function of centralization. As expected, the payoff decreases as centralization increases – this is consistent with the observation that the probability of optimal solutions decreases as centralization increases. However, when outcomes are restricted only to those cases that did achieve optimal equilibrium selection, the average payoff increases as centralization increases (dashed black line, Figure 3). This effect shows that increased centralization does not only increase speed of convergence by making equilibrium selection noisier, but also changes the distribution of likely equilibria.

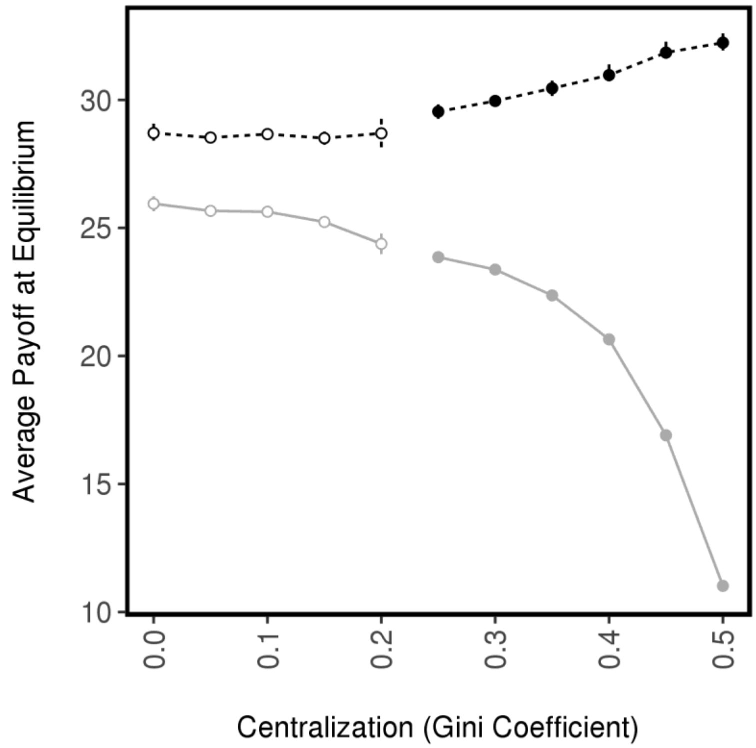

Figure 3. *The effect of network centralization on the payoff of equilibrium solutions.* As network centralization increases, the average payoff of the equilibrium solution decreases (solid grey line). However, when outcomes are restricted to only those simulations that achieved optimal coordination, the payoff of equilibrium solutions increases (dashed black line). This illustrates the influence of central individuals: in centralized networks, only the highest payoff solutions can overcome the influence of central individuals.

If centralization only made equilibrium selection noisier, then the payoff of equilibrium solutions – when restricted to cases in which the optimal solution is selected as equilibrium – would remain constant. In order to explain this effect, it is necessary to examine the role of central nodes, which exacerbate the inherent coordination tradeoff between a strategy's popularity, and that strategy's payoff. The explanation for Figure 3 is that central nodes are more influential than peripheral nodes, and thus their solution is more likely to dominate regardless of payoff. As a result, an optimal solution introduced by a peripheral node will only spread if it offers sufficiently high payoff to overcome the influence of central nodes.

The Influence of Central Individuals

Because every node has an equal probability of introducing the optimal solution, then the equilibrium solution should be (optimally) just as likely to come from a peripheral node as a central node, after controlling for the expected number of each type of node. However, central nodes are more likely to introduce the winning solution than would be expected by their prevalence in the population.

To examine the effect of central nodes, I track the origin of the equilibrium solution in each simulation, measuring the degree D of the agent that introduced the equilibrium solution, and calculate the distribution of D across repeated simulations. This is the degree distribution for only those nodes whose initial solution goes on to become the equilibrium solution, i.e. the "influencer degree distribution." If the equilibrium solution were determined only by payoff, then the influencer degree distribution would be equal to the degree distribution of the network itself. That is, if only 1% of nodes are degree 10 (i.e., have 10 conections), then only 1% of the equilibrium solutions should come from nodes with degree 10.

I then compare the influencer degree distribution with the network degree distribution[3] by calculating the likelihood ratio (L) for nodes of each degree D:

$$L(D) = \frac{proportion\ of\ influencers\ with\ degree\ D}{proportion\ of\ all\ nodes\ with\ degree\ D}$$

where "influencer" refers to a node that introduced the equilibrium solution. If the likelihood ratio L(D) is greater than 1, that means that nodes with degree D are more likely than expected to introduce the equilibrium solution. If the ratio L(D) is less than 1, then nodes with degree D are less likely than expected to introduce the equilibrium solution. Figure 4 shows L(D) for

[3] I use the exact network degree distribution, obtained by counting the total number of nodes with each degree that occur in the simulations used to generate the figure.

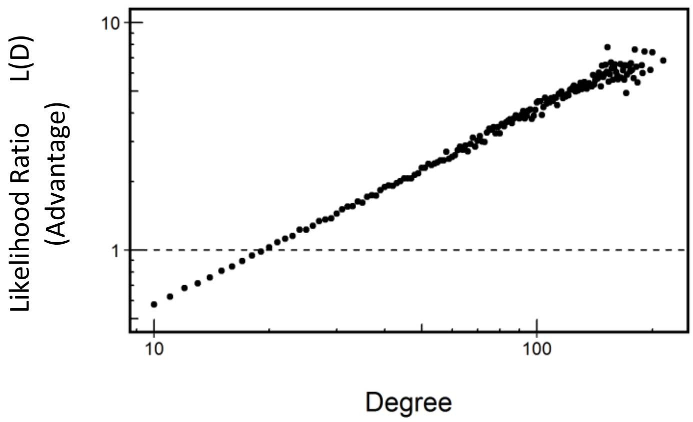

Figure 4. *The influence of central nodes.* Each point shows the relative probability that nodes of a given degree will introduce the equilibrium straetgy. This advantage is measured as the likelihood ratio comparing the probability that the equilibrium solution is selected by a node with degree D with the probability that a randomly selected node has degree D. This figure shows reuslts for BA networks with N=1000 and an average of 20 connectsion per node (m=10, a=1). Data generated from 500k simulations, shown on a log-log scale to ease readability.

centralized networks generated with Barabasi & Albert's preferential attachment algorithm with an average of 20 connections per node.

Figure 4 shows that agents with below average connectivity also have a likelihood ratio less than 1 – and are thus less likely than expected to introduce the equilibrium solution. In contrast, agents with above-average connectivity have a likelihood ratio greater than 1, and this effect increases with degree. Simply put: the more connected an agent is, the more likely it is that the agent's initial solution will become the equilibrium solution, regardless of payoff. This effect does not depend on the particular shape of the network degree distribution. The network used to generate Figure 4 has a "scale-free" degree distribution (Barabasi & Albert, 1999), and the same result is obtained for Erdos-Renyi random graphs, which have a Poisson degree distribution (Figure S2, Appendix).

Central nodes provide a widely observable signal guiding coordination – which increases speed, but also draws attention to the closest solution at hand, before the best solution has a chance to get noticed. However, a central node is just as likely as anyone else to have the best solution and thus, in expectation, offers the same payoff as a randomly selected strategy. One interpretation of the dual effect of centralization is could be that coordinating groups face an unavoidable tradeoff between speed and optimality. However, this tradeoff is not unavoidable, and in very dense networks, coordination is both fast and optimal.

Network Efficiency

Figure 5 (left) shows the effect of network density on the probability of optimal coordination in decentralized networks, where every node has the same number of connections. Networks with 20 connections per agent (density=0.02) are equivalent to the leftmost point on Figure 2 (left). Moderate increases in density have a small but negligible effect on the probability of optimal equilibrium. For networks with 10 connections per node (density=0.01) through networks with 500 connections per node (density=0.5), the optimal solution emerges in approximately 75% of trials. As density increases beyond that point, however, the probability of optimal coordination approaches 1. In fully connected networks, groups nearly always select an optimal equilibrium.

Increasing graph density not only increases the probability of optimal coordination, but also increases the speed of coordination. Figure 5 (right) shows the average number of updates that each node is required to make in order for a population to reach equilibrium, and is comparable to Figure 2 (right). The greatest gains are shown by the sparsest networks, where only a moderate increase in density produces a large decreases in convergence time. In fully connected networks, equilibrium is reached with slightly more than 1 update per person.

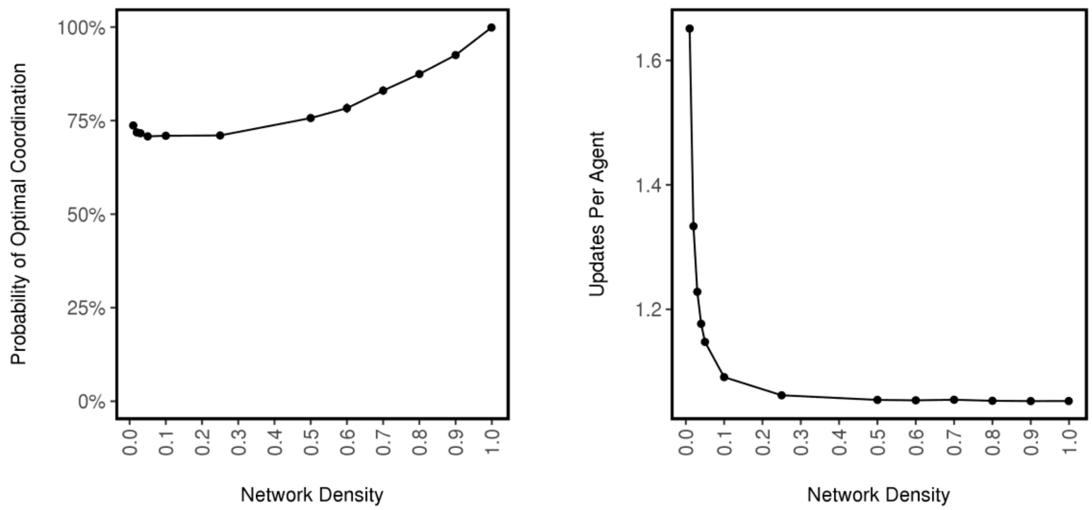

Figure 5. *The effect of network density on the probability of optimal coordination.* LEFT: Network density has a non-monotonic effect on the probability that a population will converge on the optimal strategy in equilibrium. Overall, increasing network density increases the probability of optimal equilibrium selection. In fully connected networks, every agent has full knowledge of the solution space, and a population is nearly guaranteed to converge on the optimal strategy. RIGHT: Increasing network density decreases the average number of updates each agent is required to make before the population reaches a shared convention. In fully connected networks, populations require approximatley 1 update per agent, on average, in order to reach equilibrium. BOTH PANELS: 10k simulation repetitions per point. 95% confidence intervals are drawn, but too small to be visible.

The advantage of network density for both speed and optimality can be illustrated by considering the dynamics of a fully connected network, in which every agent observes every other agent. At the outset, every node has a unique solution. Because every node has full knowledge of the strategy space, and no strategy yet has any advantage of popularity, every node will simply adopt the best strategy. A similar argument can be made for the case where there are fewer strategies than people. In expectation, each strategy will be employed by the same number of people at the outset – thus, again, no strategy starts with a popularity advantage. Because every node has full knowledge of the solution space, the best strategy will be immediately adopted. In simulation, slightly more than 1 update per person is required due to the presence of noise in decision-making.

One unexpected outcome, shown in Figure 5 (left), is that the probability of optimal coordination is non-monotonic with respect to graph density. While simulation is useful for overcoming the tractability limitations of formal analysis, one downside is that it is more difficult to explain "why" a certain result emerges (Macy & Willer, 2002). One explanation for the non-monotonic effect of graph density is that increasing graph density produces two opposing forces. First, graph density helps solutions spread faster, including suboptimal solutions that – by chance – are adopted at higher rates. Thus, a small amount of network density allows an inferior solution to "take advantage" of early adoption due to random chance. Second, increasing graph density allows agents obtain a wider initial knowledge of the solution space. Thus, as density increases further, it becomes less likely that inferior solutions will obtain an initial advantage.

These two opposing forces in decentralized networks indicate a tradeoff between the speed of diffusion enabled by network density, and amount of information available to each agent, which increases with network efficiency. The non-monotonic effect of density in decentralized networks also provides a framework for understanding the problem of centralized networks: as networks become more centralized, adoption speed is increased, but information availability is not increased for most agents. As central nodes become better connected, peripheral nodes become less connected, observing a shrinking proportion of the population as a whole.

Robustness

The results presented so far have demonstrated the properties of only one point in the parameter space with regard to statistical noise, number of solutions, and distribution of solution payoff. In particular, I have assumed a fairly large level of noise, an infinitely large strategy space (so that agents begin with individually unique solutions) and strategy payoffs drawn from a log-normal distribution. The log-normal distribution of strategies reflects a long-tailed innovation process, such that many mediocre solutions and a few excellent solutions are likely to be generated (Kauffman, 1993).

In order to study the factors that shape equilibrium selection in networks, I also examine how robust are the detrimental effects of centralization to variation in assumptions about statistical noise and the strategy space. The main results assumed a high level of noise, and effects are unchanged in the absence of noise (see Figure S1, left, in appendix). Effects are also comparable when individual strategies are drawn from a uniform payoff distribution, rather than a log-normal distribution (Figure S1, right, in appendix). However, one factor impacting the effect of centralization is the size of the solution space.

As the number of solutions is reduced, the effect of centralization becomes weaker, but nonetheless has a substantial impact on the probability of optimal coordination (Figure 6). Even with only 10 unique solutions, the most centralized networks still select an optimal equilibrium less than 80% of the time. However, once the number of solutions reaches its limit (binary coordination, with 2 solutions) the effect of centralization becomes negligible. The decreased impact of centralization can be explained in terms of the initial conditions of the model. At time t=0, every begins with a randomly assigned strategy. When each strategy is unique, the strategies adopted by central nodes have a structural advantage. However, with only two strategies, the structural advantage of each strategy will, in expectation, be approximately the same: each strategy is likely to be adopted by the same number of central individuals.

Although the effect of centralization is small in binary coordination, the most centralized networks nonetheless show a slightly decreased probability of optimal coordination as compared

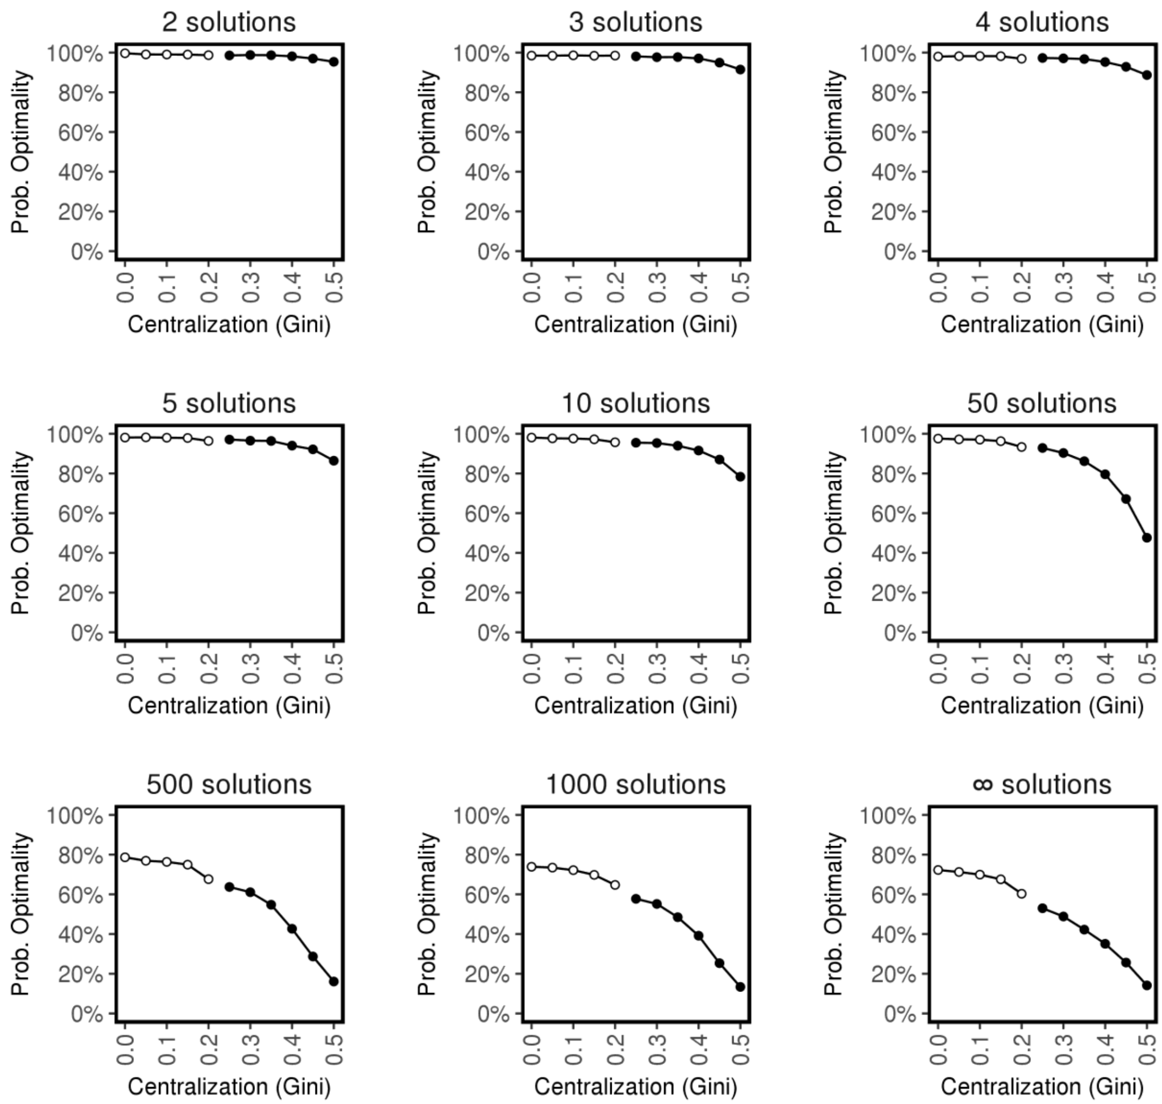

Figure 6. *The detrimental effect of centralization is robust to the size of the strategy space.* This figure replicates Figure 2 (left) across a range of strategy space sizes, where the size of the strategy space determines the number of unique solutions when the model is initialized. In all of the of the figures shown above, the strategies were drawn from an infinite set, so that each agent had a unique solution (lower right panel, this figure). As the number of strategies decreases, the effect of centralization remains the same, but the effect size diminishes. In binary coordination (2 strategies) populations converge on the optimal solution in nearly every simulation.

with decentralized networks. The robustness of this effect shows that even in relatively straightforward decisions, the structural advantage given to solutions adopted by central nodes always has some impact on collective decisions.

Discussion

Central nodes are a bit of a paradox. On the one hand, their centrality enables them to obtain a wide view of the information space. As theorized in previous research, central nodes have the potential to coordinate the behavior of peripheral individuals, by actively integrating their information and activity for the benefit of the group as a whole (Mulder, 1960). However, the social information observed by central nodes is mirror-like, because they are observing individuals who, simultaneously, observe them. Before a central node has the chance to integrate the activities of the peripheral agents in a network, those agents have already been influenced by the central node. Thus, by the time a central node begins to integrate information and make a decision, it will appear as though their initial solution is already quite popular – because their social information reflects their own behavior - and they will be unlikely to use any of the alternative strategies introduced by peripheral nodes.

Although the coordination model presented here captures a highly stylized description of human behavior, as does any model, the effects described here are similar to those observed in related models of collective decision-making. In both a computational model and an experimental test of the wisdom of crowds, Becker et al. (2017) found that decentralized networks perform better than centralized networks on collective estimation tasks. In many ways, estimation decisions are shaped by a fundamentally different form of social influence than coordination decisions. In coordination decisions, individual choices are inherently interdependent, because a solution's payoff depends on the number of peers employing that solution. In estimation decisions, individual decisions are independent, and beliefs are shaped by social influence only to the extent that social information provides a useful signal. Nonetheless, central nodes face a similar paradox in both coordination dynamics and the wisdom of crowds. In forming estimation judgements, a central node is in a position to integrate the information provided by group as a whole. However, the peripheral nodes are simultaneously influenced by the central node, and the network as a whole is drawn toward the belief of central nodes (Becker et al., 2017).

One valuable payoff of computational experiments such as the one conducted here is the ability to identify formal hypotheses that can be tested in laboratory experiments (Macy & Willer, 2002; Centola, Macy & Wheelan, 2005; Smith & Rand, 2016). By simulating a commonly used behavioral model of interpersonal coordination, I have generated hypotheses relating the structure of interaction networks to the probability that a group will reach an optimal equilibrium solution. Centralized networks are expected to coordinate faster, but also are more likely to coordinate on a suboptimal solution.

The theoretical results presented here also draw attention to the danger of opinion leaders. Although the presence of central members in a network can be beneficial when speed is an issue (Mulder, 1960; Rogers, 1983; Banerjee et al, 2013), the influence of central individuals also carries a risk. Most critically, a central node can attract a group towards popular strategies, at the expense of less popular, higher payoff alternatives.

References

Arthur, W. B. (1989). Competing technologies, increasing returns, and lock-in by historical events. *The Economic Journal*, *99*(394), 116–131.

Badham, J. M. (2013). Commentary: measuring the shape of degree distributions. *Network Science*, *1*(2), 213–225.

Banerjee, A., Chandrasekhar, A. G., Duflo, E., & Jackson, M. O. (2013). The diffusion of microfinance. *Science*, *341*(6144), 1236498.

Barabási, A.-L., & Albert, R. (1999). Emergence of scaling in random networks. *Science*, *286*(5439), 509–512.

Barabâsi, A.-L., Jeong, H., Néda, Z., Ravasz, E., Schubert, A., & Vicsek, T. (2002). Evolution of the social network of scientific collaborations. *Physica A: Statistical Mechanics and Its Applications*, *311*(3), 590–614.

Becker, J., Brackbill, D., & Centola, D. (2017). The Network Dynamics of Social Influence in the Wisdom of Crowds. *Manuscript Submitted for Publication.*

Bettenhausen, K., & Murnighan, J. K. (1985). The emergence of norms in competitive decision-making groups. Administrative Science Quarterly, 350–372.

Blume, L. E. (1993). The statistical mechanics of strategic interaction. *Games and Economic Behavior*, *5*(3), 387–424.

Blume, L. E. (1995). The statistical mechanics of best-response strategy revision. *Games and Economic Behavior*, *11*(2), 111–145.

Centola, D., & Baronchelli, A. (2015). The spontaneous emergence of conventions: An experimental study of cultural evolution. *Proceedings of the National Academy of Sciences*, *112*(7), 1989–1994.

Centola, D., & Macy, M. (2007). Complex contagions and the weakness of long ties. *American Journal of Sociology*, *113*(3), 702–734.

Dall'Asta, L., Baronchelli, A., Barrat, A., & Loreto, V. (2006). Nonequilibrium dynamics of language games on complex networks. Physical Review E, 74(3), 036105.

David, P. A. (1985). Clio and the Economics of QWERTY. *The American Economic Review*, *75*(2), 332–337.

David, P. A., & Bunn, J. A. (1988). The economics of gateway technologies and network evolution: Lessons from electricity supply history. *Information Economics and Policy*, *3*(2), 165–202.

DeMarzo, P. M., Vayanos, D., & Zwiebel, J. (2003). Persuasion Bias, Social Influence, and Unidimensional Opinions. *The Quarterly Journal of Economics*, *118*(3), 909–968.

Eagle, N., & Pentland, A. S. (2006). Reality mining: sensing complex social systems. *Personal and Ubiquitous Computing*, *10*(4), 255–268.

Easley, D., & Kleinberg, J. (2010). Networks, crowds, and markets. *Cambridge University Press*, *6*(1), 1–6.

Ebel, H., Mielsch, L.-I., & Bornholdt, S. (2002). Scale-free topology of e-mail networks. *Physical Review E*, *66*(3), 035103.

Ellison, G. (1993). Learning, local interaction, and coordination. *Econometrica: Journal of the Econometric Society*, 1047–1071.

Ellison, G. (2000). Basins of attraction, long-run stochastic stability, and the speed of step-by-step evolution. *The Review of Economic Studies*, *67*(1), 17–45.

Freeman, L. C. (1978). Centrality in social networks conceptual clarification. *Social Networks*, 1(3), 215–239.

Heal, G., & Kunreuther, H. (2010). Social reinforcement: cascades, entrapment, and tipping. *American Economic Journal: Microeconomics, 2*(1), 86–99.

Kandori, M., Mailath, G. J., & Rob, R. (1993). Learning, mutation, and long run equilibria in games. *Econometrica: Journal of the Econometric Society*, 29–56.

Katz, E., & Lazarsfeld, P. (1955). *Personal Influence*, *The part played by people in the flow of mass communications*. Toronto, ON: The Free Press.

Kauffman, S. A. (1993). *The origins of order: Self-organization and selection in evolution*. Oxford University Press, USA.

Kossinets, G., & Watts, D. J. (2006). Empirical analysis of an evolving social network. *Science*, *311*(5757), 88–90.

Leavitt, H. J. (1951). Some effects of certain communication patterns on group performance. *The Journal of Abnormal and Social Psychology*, *46*(1), 38.

Liljeros, F., Edling, C. R., Amaral, L. A. N., Stanley, H. E., & \AAberg, Y. (2001). The web of human sexual contacts. *Nature*, *411*(6840), 907–908.

Macy, M. W., & Willer, R. (2002). From factors to factors: computational sociology and agent-based modeling. *Annual Review of Sociology*, *28*(1), 143–166.

Maslov, S., Sneppen, K., Alon, U., Bornholdt, S., & Schuster, H. G. (2003). *Handbook of graphs and networks*. Wiley-VCH New York:

Montanari, A., & Saberi, A. (2010). The spread of innovations in social networks. *Proceedings of the National Academy of Sciences*, *107*(47), 20196–20201.

Morris, S. (2000). Contagion. *The Review of Economic Studies*, *67*(1), 57–78.

Mulder, M. (1960). Communication structure, decision structure and group performance. *Sociometry*, *23*(1), 1–14.

Rogers, E. M. (1983). *Diffusion of innovations* (3rd ed.). New York, NY: The Free Press.

Schelling, T. C. (1980). *The strategy of conflict*. Harvard university press.

Smith, E. B., & Rand, W. (2016). Simulating macro-level effects from micro-level observations. In *Academy of Management Proceedings* (Vol. 2016, p. 12927). Academy of Management.

Watts, D. J., & Strogatz, S. H. (1998). Collective dynamics of'small-world'networks. *Nature*, *393*(6684), 440.

Young, H. P. (1993). The evolution of conventions. *Econometrica: Journal of the Econometric Society*, 57–84.

Appendix

Network Generators

In the study here, two primary network generators are used. Preferential attachment networks are generated according to the algorithm developed by Barabasi and Albert (1999). To vary the range of centralization produced by this generator, I vary the strength of preferential attachment. As new nodes are connected, they are connected to existing nodes with probability

$$P \sim ck^a$$

where k is a node's degree, a controls the strength of preferential attachment and c is a normalizing constant. This model reduces to the original Barabasi-Albert algorithm when a=1.

The preferential attachment generator produces networks with a Gini coefficient in the range of 0.25 to 0.5. To increase the range available for study, and to allow for the study of a continuous variation from decentralized to centralized networks, I also generate networks with a preferential rewiring algorithm according to the following method:

- A random graph is generated with mean degree D by permuting links of an initial lattice graph (Maslov, Sneppen, & Alon, 2003; see also Centola & Macy, 2008). This generates a random, decentralized network in which every node has the same number of connections.
- A node i is then uniformly selected for rewiring, and one edge is uniformly selected. A second node j≠i is selected according to the probability

  $$P \sim ck^a$$

  where k is a node's degree, a is a preferential attachment parameter, and c is a normalizing constant, as above.

- If node j is not already connected to node i, the new connection is formed and the old connection is broken.
- This process is repeated R times.

This method generates preferentially-rewired random networks where centralization varies with R and a. Combined with the preferential attachment network generator, this allows us to study centralization ranging continuously from fully decentralized (gini=0) to highly centralized (gini=0.5).

Supplementary Figures

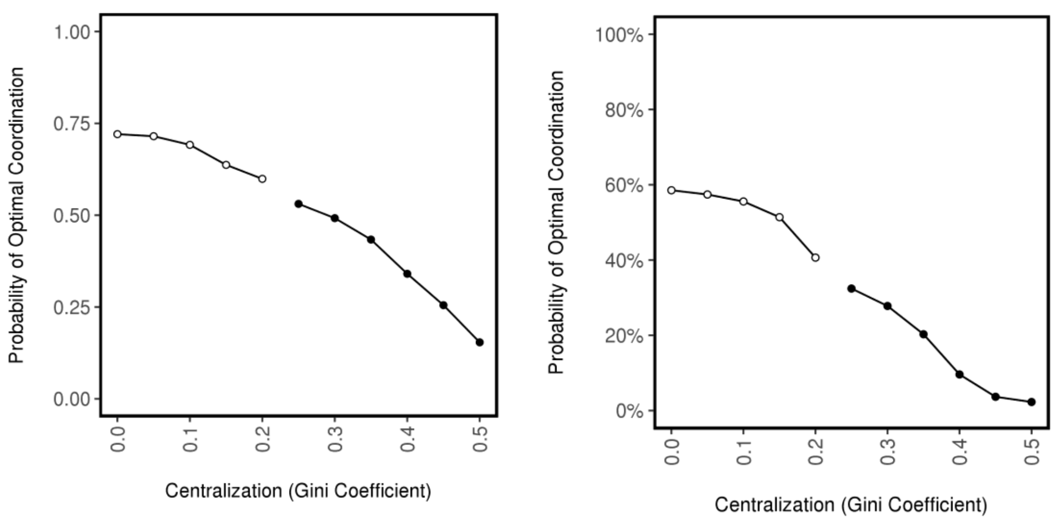

Figure S1. *The detrimental effect of centralization is robust to variation in noise and payoff distribution.* LEFT: A replication of Figure 2 (left) but without noise in agent decisions, producing qualitatively identical results. RIGHT: A replication of Figure 2 (left) but with payoff for each strategy drawn from a uniform (0,1) distribution instead of a log-normal distribution. With a uniform distribution, the population is generally more likely to select a suboptimal equilibrium strategy overall, because the relative payoff between strategies (eg, the difference between best and the second best strategies) is smaller, as compared with a long-tailed log-normal distribution.

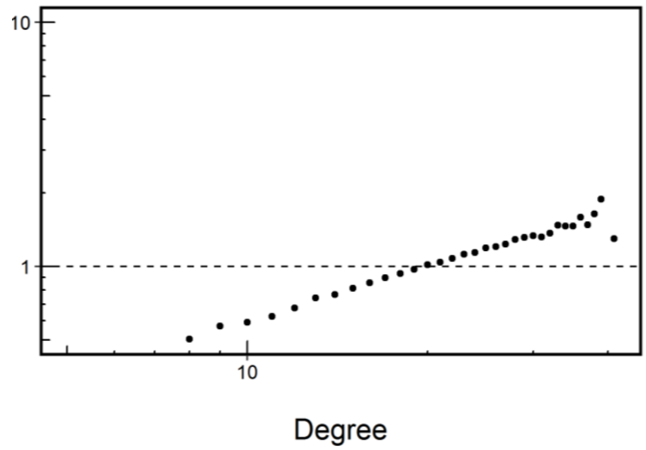

Figure S2. *The influence of central nodes is robust to variation in degree distribution.* This figure replicates Figure 3, using an Erdos-Renyi random graph (mean degree=20) which produces a Poisson degree distribution, instead of a Barabasi-Albert preferential attachment network, which produces a "scale free" degree distribution.