

UC Irvine

UC Irvine Previously Published Works

Title

Self-Assembling Networks

Permalink

<https://escholarship.org/uc/item/887789vc>

Journal

The British Journal for the Philosophy of Science, 70(1)

ISSN

0007-0882

Authors

Barrett, Jeffrey A
Skyrms, Brian
Mohseni, Aydin

Publication Date

2019-03-01

DOI

10.1093/bjps/axx039

Peer reviewed

SELF-ASSEMBLING NETWORKS

JEFFREY A. BARRETT, BRIAN SKYRMS, AYDIN MOHSENI

ABSTRACT. We consider how an epistemic network might self-assemble from the ritualization of the individual decisions of simple heterogeneous agents. In such evolved social networks, inquirers may be significantly more successful than they could be investigating nature on their own. The evolved network may also dramatically lower the epistemic risk faced by even the most talented inquirers. We consider networks that self-assemble in the context of both perfect and imperfect communication and compare the behavior of inquirers in each. This provides a step in bringing together two new and developing research programs, the theory of self-assembling games and the theory of network epistemology.

1. INTRODUCTION

Simple games may self-assemble from individual decision problems. Complex games may self-assemble from simple games. Games self-assemble by means of a learning dynamics acting on the actions of individual agents. The outlines of this theory with some elementary examples have been set out in a series of papers.¹ There is also a growing literature on adaptive networks.²

Social epistemology needs to study how information is transmitted through a social network. Philosophers have been especially interested in the possibility of premature “lock-in” to consensus in a scientific community that learns too fast—that “jumps to conclusions.”³ It has been shown by example that this possibility may be sensitive to the network structure in ways that are not obvious without analysis.⁴ This literature compares fixed network structures. Here we consider an evolving network structure.

Social networks are not just given; they evolve. Here we consider how an epistemic network might self-assemble as heterogeneous agents investigate nature and each other. When successful, even very simple agents may self-assemble a social network that allows them to be more successful than they could have been had they

Date: August 26, 2016.

¹See Barrett and Skyrms [4] for a general description of the theory and examples and Barrett [3], Huttegger and Skyrms [11], Skyrms and Pemantle [20], and Bala and Goyal [1] for further examples. See McGregor [15] for examples of evolved signaling networks in nature.

²See [8] for a recent survey of this work.

³See, for example, Kitcher [12][13] and Strevens [21][22].

⁴See Zollman [24][23] and Rosenstock, O’Connor, and Bruner [17].

investigated nature on their own. We focus on reinforcement learning, a low rationality way of “feeling ones way toward a better network.” This is not because we think that this is the “right” dynamics. Various learning dynamics are relevant to this large field, and all deserve to be investigated. We choose reinforcement learning because (1) it has a long track record applied to interactive decisions, (2) success of reinforcement learning would suggest that more sophisticated learning might also be successful, and (3) its degree of success would provide a baseline against which to measure the effectiveness of other forms of learning.

We consider self-assembly of epistemic networks under two different conditions. In the first case, information is transmitted along a link with fidelity; in the second there are transmission errors. In both cases, networks form by means of polymerization, but the different conditions result in different kinds of network structures.⁵ The resulting networks help the epistemic community in ways that we will describe in detail. In short, agents in the self-assembled networks are nearly always significantly more successful together than they would have been had they investigated nature on their own. The self-evolved networks promote the epistemic goals of the community and lower the agents’ individual epistemic risk.

2. THE BASIC MODEL

Consider a group of inquirers who set out to investigate nature and use what they learn for the purpose of successful action. When an inquirer investigates, we will suppose that she does one of two things: she may observe nature or consult an agent by asking her what she has learned so far. For the present model, we will say that an inquirer knows the truth, and is hence prepared for successful action or to share the truth with others, if she knows how to solve the particular empirical problem at hand. And she knows this if her most recent observation was successful or if the most recent inquirer she consulted successfully communicated the truth.

If every inquirer were a perfect observer, each might learn the truth from nature directly. In this case there would be no point in consulting other agents. But if inquirers differ in the reliability of their observations, less reliable inquirers may learn from their more reliable colleagues. We will suppose that while some inquirers typically learn the truth when they make an observation, others are less reliable. Given this, the community would do best if those Inquirers who are most reliable observe nature directly and those who are less reliable learn from their more reliable colleagues. But which inquirers are epistemically reliable is itself a matter for inquiry. Inquirers must learn this even as they investigate nature. If they are

⁵See Barrett and Skyrms [4] for a characterization of the evolution of complex games polymerization, modular composition, and template transfer. In the present models, networks self-assemble primarily by means of polymerization.

successful in determining who can be trusted, they may evolve a social network where they are more more successful together than they could be apart.

We will consider five-agent and twenty-agent models where each inquirer is initially assigned a random, unbiased *epistemic reliability* in the interval $[0, 1]$. An inquirer’s epistemic reliability represents the probability that she will successfully determine the truth on a particular observation. Each inquirer is also assigned a *social reliability*. An inquirer’s social reliability represents the probability that she will reliably determine the truth when she consults an agent who knows the truth. For the sake of simplicity, we will suppose that an inquirer’s epistemic and social reliabilities do not change over time.

We will suppose that the inquirers determine who or what to consult by simple reinforcement learning.⁶ It is the learning dynamics that provides the feedback that may ritualize the individual actions of the agents over time to evolve increasingly complex games. A central question here is the extent to which simple reinforcement learning might allow the inquirers to self-assemble a game where the epistemic community successfully exploits the most reliable observers.

Simple reinforcement learning can be modeled by adding and drawing balls from an urn. To this end, we will suppose that each agent is equipped with an urn that initially contains one ball corresponding to nature and one ball corresponding to each agent. An agent decides whether to observe nature or consult a particular agent by randomly drawing a ball from her urn. She then consults the source indicated on the ball. The ball is subsequently returned to the urn and a second ball of the same type may be added to the urn if the consultation was successful.

The inquirers investigate the world in rounds of play, repeated over multiple generations. Each generation is characterized by a particular empirical problem that must be solved for successful action. An agent who knows how to solve the problem knows the truth. Each round, the inquirers take turns attempting to solve the current problem by observation or consultation with an agent. We will consider simulations consisting of 200 generations where each generation has 100 full rounds of play.

At the start of each generation no inquirer can solve the problem characterizing that generation. That is, no inquirer knows the truth. In each round, the inquirers are randomly ordered, and they play once each in turn. A play begins with the inquirer whose turn it is drawing a ball from her urn.

⁶See Herrnstein [9] for an early description of simple reinforcement learning. See Roth and Erev [18] [6] for discussions of more subtle types of reinforcement learning and how human agents exhibit reinforcement learning in practice. See Barrett and Zollman [5], Skyrms [19], Huttegger, Skyrms, Tarrès, and Wagner [10], and Barrett, Cochran, Fujiwara, and Huttegger [2] for discussions of these and other learning dynamics in the context of signaling games.

If the ball represents *nature*, the inquirer makes an observation and learns the truth with a probability equal to her *epistemic reliability*. If she learns the truth, then she is successful in action and prepared to share the truth if consulted. In this case, she returns the *nature* ball to her urn and adds another ball of the same type. Otherwise, she ends up believing something that is false, fails in action, and is unprepared to share the truth if consulted. In this case, she simply returns the *nature* ball to her urn.

If the ball represents *agent n*, the inquirer consults agent *n*. If *n* knows the truth, then the inquirer learns the truth with probability equal to her *social reliability*. If the inquirer learns the truth, then she is successful in action and prepared to share the truth with another agent if consulted. In this case, she returns the agent-*n* ball to her urn and adds another ball of the same type. Otherwise, she ends up believing something that is false, fails in action, and is unprepared to share the truth if consulted. In this case, she simply returns the *agent n* ball to her urn.⁷

After each inquirer has played once in random order, the round ends and the inquirers begin a new round with the same empirical problem and what they have learned regarding its solution. At the end of a generation, the inquirers start a new investigation with a fresh empirical problem and begin in a state of complete ignorance.


3. EQUILIBRIUM ANALYSIS

We will begin by considering the conditions for a deterministic equilibrium network. We will distinguish between *perfect communication* where each inquirer has an unbiased random epistemic reliability in the interval $[0, 1]$ but a social reliability of 1 and *imperfect communication* where an inquirer's social reliability is also randomly determined. In the first case, inquirers always learn the truth when they consult an agent who knows the truth; in the second, an inquirer might consult an agent who knows the truth but not learn it. After considering the equilibrium conditions for perfect and imperfect communication on a single, simultaneous play of the game, we will consider what actually happens when one runs the evolutionary game under the sequential dynamics described in the last section, then compare this medium-term behavior to the deterministic equilibria.⁸

On a single play of the game, each agent either consults nature or precisely one other agent. The combination of the pure strategies corresponding to these actions defines a network. A *Nash network* is one such that no agent could gain

⁷Note that an inquirer might consult herself. If so, she will learn the truth if and only if her current belief is true.

⁸We consider pure strategies and deterministic networks in this section. As we will see later, the random networks that evolve on the learning dynamics are well-approximated by the deterministic Nash equilibrium networks.



(A) Perfect Communication Equilibrium (B) Imperfect Communication Equilibrium

FIGURE 1. Representative equilibria for perfect and imperfect communication. In (A), the most reliable agent ϕ consults nature, while all other agents have a path to ϕ . In (B), the most reliable agent consults nature, independent agents ($e_i > e_\phi s_i$) consult nature, and dependent agents ($e_i < e_\phi s_i$) directly consult the most reliable agent ϕ .

in payoff by unilaterally changing her strategy. It is a *strict Nash network* if any agent unilaterally changing her strategy would decrease her payoffs. We will take an agent’s payoff to be her expected return on the single play where the value of getting the truth is 1 and not getting it is 0. For perfect communication we will set aside the measure-zero cases where agent epistemic reliabilities are exactly 0 or 1 and for imperfect communication the cases where agent epistemic or social reliabilities are exactly 0 or 1. Finally, we will suppose that there is always a single most epistemically reliable agent in the population.

Claim 1. *The Nash networks for perfect communication are exactly those where the most reliable agent consults nature and every other agent has a path to the most reliable agent.*

Proof. If the most reliable agent consults nature and every other agent has a path to the most reliable agent, then the network is Nash. Consider such a network and suppose that the most epistemically reliable agent unilaterally deviates consulting another agent rather than directly consulting nature. In that case, no one consults nature, so no one learns the truth and the most reliable agent does strictly worse. Suppose that another agent consults nature directly instead of consulting the most

reliable agent. In that case, the deviating agent does strictly worse since she has a lower epistemically reliability than the most reliable agent. If there are more than two agents, then the deviating agent might consult a third agent instead of consulting the most reliable agent. If there is still a chain from the third agent to the most reliable agent, then the deviating agent gets the same payoff as she did initially and hence does not gain by the deviation. Otherwise, the deviation breaks the path to the most reliable agent and the deviating agent does strictly worse. Note that the only strict Nash networks in the case of the perfect communication are those with fewer than three agents.

Only networks where the most reliable agent consults nature and every other agent has a path to the most reliable agent are Nash. There are two cases to consider. First, consider a network where the most reliable agent does not consult nature directly. In this case, she would do at best as well as an agent who is less reliable and consults nature. So such a network is not Nash since the most reliable agent could do strictly better by consulting nature directly. Second, consider a network where there is an agent without a path to the most reliable agent. That agent would do strictly better by consulting the most reliable agent. \square

Now consider imperfect communication. Let $1, 2, \dots, n$ index the n agents in the population, and e_1, e_2, \dots, e_n denote their corresponding epistemic reliabilities. Let ϕ denote the most reliable agent, and $e_\phi = \max(\{e_i : i = 1, 2, \dots, n\})$ her epistemic reliability. Let s_1, s_2, \dots, s_n denote the agents' social reliabilities.

Setting aside the case, for a moment, where there are agents such that $e_i = e_\phi s_i$, there are two cases to consider: (I) there are no agents i such that $e_i > e_\phi s_i$ and (II) there are agents i such that $e_i > e_\phi s_i$.

Claim 2. *The Nash networks for imperfect communication are exactly those where in case (I) the most reliable agent ϕ consults nature and all other agents consult ϕ and in case (II) ϕ consults nature, agents i where $e_i > e_\phi s_i$ consult nature, and all other agents consult ϕ . Further, such networks are strict Nash.*

Proof. Consider case (I). If there are no agents i such that $e_i > e_\phi s_i$, networks where the most reliable agent ϕ consults nature and all other agents consult ϕ are strict Nash. Start with such a network. If ϕ unilaterally deviates and consults another agent, then no one consults nature and ϕ does strictly worse. If another agent k deviates from consulting ϕ , then she either (1) consults nature directly and does strictly worse since $e_k < e_\phi s_k$ or (2) consults another agent and does strictly worse since her chain to ϕ is lengthened and $s_k < 1$.⁹

If there are no agents i such that $e_i > e_\phi s_i$, only networks where the most reliable agent ϕ consults nature and all other agents consult ϕ are Nash. Consider

⁹Again, we will discuss the case where there are agents such that $e_i = e_\phi s_i$ later.

a network where the most reliable agent ϕ does not consult nature. Whatever she is doing, since ϕ has the highest epistemic reliability, she would do strictly better by consulting nature directly. Consider a network where ϕ consults nature, but someone else k does not consult ϕ . Either k is consulting nature, in which case she does strictly worse than she would consulting ϕ since $e_k < e_\phi s_k$, or she is consulting a third agent who is not the most epistemically reliable which is strictly worse than consulting ϕ , so in either case, she would do better by consulting ϕ .

Consider case (II). We will call those agents i such that $e_i > e_\phi s_i$ *independent* and all other agents, except the most epistemically reliable agent ϕ , *dependent*. If there are independent agents, then networks where the most reliable agent ϕ consults nature, independent agents consult nature, and dependent agents consult ϕ are strict Nash. In this case, if ϕ unilaterally deviates, she either consults a dependent agent and gets no payoff since she then has no path to nature or consults an independent agent k and does strictly worse than she would consulting nature directly since $s_\phi e_k < e_\phi$. If a dependent agent deviates from consulting ϕ , then she does strictly worse consulting nature since $e_i < e_\phi s_i$ for dependent agents and strictly worse consulting an independent agent since these agents are not the most epistemically reliable. If an independent agent deviates from consulting nature, then she does strictly worse consulting the most reliable agent since $e_i > e_\phi s_i$ for independent agents and does yet worse consulting anyone else.

If there are independent agents, then only networks where the most reliable agent ϕ consults nature, independent agents consult nature, and dependent agents consult ϕ are Nash. Consider a network where the most reliable agent ϕ does not consult nature directly. Whatever she is doing, as the most epistemically reliable agent ϕ would do better consulting nature directly. Consider a network where a dependent agent consults nature. Since $e_i < e_\phi s_i$ for dependent agents, she would do better consulting ϕ . Consider a network where she consults an independent agent k . She would do better consulting ϕ since $e_\phi > e_k$. Consider a network where an independent agent does not consult nature. If she consults ϕ , then she would do better consulting nature since $e_i > e_\phi s_i$ for independent agents. And in a network where she consults any other agent agent, she would do yet better consulting nature directly. \square

Now consider the case we initially set aside where there are agents such that $e_i = e_\phi s_i$. These agents would do equally well consulting nature directly or consulting the most epistemically reliable agent ϕ . Hence, agents satisfying this condition would allow for equilibria that are not strict Nash. Importantly, except for this measure-zero case, the equilibrium networks for imperfect communication are strict Nash.

In the case of perfect communication, then, there are many Nash equilibria like figure 1(A).¹⁰ While all agents are ultimately linked to the most reliable inquirer ϕ , it may be by way of a relatively long path. Here the success of the inquirers does not depend on any particular agent identifying the most reliable inquirer. When one introduces imperfect communication, however, there is now a single strict Nash equilibria with a star-shaped structure around the most reliable inquirer as in figure 1(B) and with independent inquirers consulting nature directly, for a fixed distribution of epistemic reliabilities, the worse the social reliability of the community, the more independent inquirers. Here the dependent inquirers all directly identify the most reliable agent.

The next step is to see what actually happens in the evolutionary sequential game. Given the complexity of this game under even simple reinforcement learning, this requires simulation.

4. PERFECT COMMUNICATION

In the case of perfect communication we suppose that each inquirer has an epistemic reliability randomly selected from the interval $[0, 1]$ with uniform probabilities but a social reliability of 1. Each inquirer, hence, will learn the truth about nature by consulting any agent who knows the truth.

On simulation, perfectly communicating inquirers are always found to self-assemble a network that allows them to be more successful together than they could be apart. In brief, epistemically reliable inquirers come to be consulted more often by others, while they themselves are more likely to observe nature directly. Epistemically unreliable inquirers are somewhat less likely to be consulted by other inquirers and are found to rarely observe nature directly.

Figure 2 illustrates an epistemic network that self-assembled on a typical run of the five-agent model. The nodes represent nature and the five agents, and the strength of the vertices represent the evolved likelihood of each agent consulting another source. Figure 2(A) gives the full graph while figure 2(B) includes just the most significant edges (with weights above 0.2). On this run of the model, agent 5 evolved dispositions to observe nature directly, and each of the other agents evolved dispositions to consult agent 5 with high probability. A network like this is typical when one of the inquirers is significantly more reliable than the others in investigating nature. An optimally reliable network would have just the structure of figure 2(B). Here the self-assembled network approximates the optimally network precisely insofar as figure 2(A) approximates the structure of figure 2(B).

¹⁰If one considers indeterministic networks, there are yet more Nash equilibria as a agent can do just as well mixing over paths that surely lead to the most reliable agent.




FIGURE 2. Five-agent model with perfect communication

The differentiation in the roles played by epistemically reliable and unreliable agents in the evolved network can be seen by considering both who is consulted most often and how the consulting dispositions of each type of agent evolves. In general, epistemically reliable agents are consulted more frequently over multiple runs of the model. In the five-agent model, the top quartile of reliable agents are consulted by the population on average 0.217 of the time, in contrast to the general population average of 0.164, which is close to the frequency of consultation of 0.167 expected if the population were consulting at random with uniform probabilities.¹¹ This effect is independent of scale. When one moves to the twenty-agent model, the top quartile of reliable agents are consulted by the population on average 0.058 of the time, in contrast to the general population average of 0.047, which again is close to the frequency of consultation of 0.048, expected if the population were consulting at random.

Correspondingly, epistemically unreliable agents are consulted less frequently on the evolved networks, with the bottom quartile being consulted on average 0.123 and 0.042 of the time in the five-agent and twenty-agent models, respectively. While one might be surprised that epistemically unreliable agents are consulted at all, note that, since we are assuming perfect communication, an agent may be unreliable at consulting nature directly yet become a consistent source of reliable information by learning to consult with reliable agents. Indeed, in the context of perfect communication, epistemically unreliable agents *typically* become conduits

¹¹The cumulative statistics here and below are means for 100 runs of the model.

of reliable information and may hence evolve to be routinely consulted by other members of the network.

Figures 3 and 4 illustrate the frequency of being consulted and the frequency of consulting nature as a function of an agent's epistemic reliability, each on one hundred runs of the twenty-agent model. Each data point in these figures represents a single agent at the end of a run of the model.

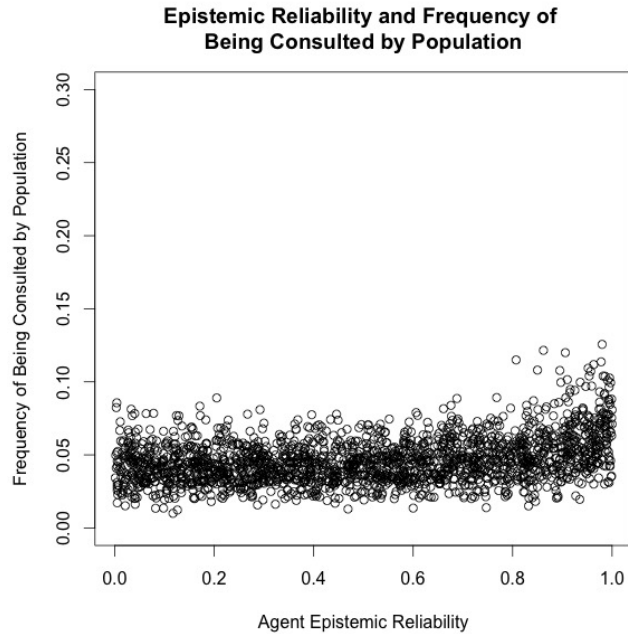


FIGURE 3. Epistemic Reliability and Frequency of Being Consulted with Perfect Communication

Regarding the dispositions of the inquirers by type, epistemically unreliable inquirers tend to consult reliable inquirers, and reliable inquirers tend to observe nature. The bottom quartile of unreliable agents rarely observe nature, with frequencies of 0.004 and less than 0.001 for the five-agent and twenty-agent models respectively. Whereas the top quartile of reliable agents come to observe nature 0.446 and 0.229 of the time in the two models. The simulations suggest that there is a threshold effect at work here. In Figure 4 note that an agent almost never evolves to observe nature directly if his epistemic reliability is less than 0.5.

The upshot is that, even on simple reinforcement learning, the epistemic community typically learns who the most reliable inquirers are. These inquirers observe nature more often and are consulted more often by their colleagues. And the results of their observations are disseminated throughout the evolved network.




FIGURE 4. Epistemic Reliability and Frequency of Consulting Nature with Perfect Communication

The network that self assembles promotes overall reliability. Indeed, most inquirers are much more successful in the evolved network than had they made their own observations. If they were to act asocially and each simply observe nature directly, their expected success rate would be equal to the mean epistemic reliability of the community, which, since we are supposing an even distribution of epistemic reliability of the agents on the interval $[0, 1]$, would be 50%. In contrast, when the inquirers are allowed to self-assemble by simple reinforcement, the mean success rate for the community is 0.760 for the five-agent model and 0.786 for the twenty-agent model. Indeed, even the most unreliable quartile of agents significantly outperform the baseline, with success rates for the two models of 0.732 and 0.780 respectively, which is nearly as good as the mean success rate. Those inquirers who are not themselves reliable at directly observing nature learn from reliable inquirers, or from inquirers who learn from reliable inquirers, etc.

The reliability of the most reliable quartile of inquirers in the evolved network is comparable to what it would have been had they consulted nature directly with success rates of 0.859 and 0.818 for the top quartile of agents in the five-agent and twenty-agent models. This is just less than the 0.875 success rate the top quartile of inquirers would have had they investigated nature directly without social interactions. The lower success rate for the most reliable inquirers in the

twenty-agent model appears to be the result of there being more potential sources of information. On reinforcement learning, these options serve to distract the most reliable inquirers from directly observing nature more frequently. The most reliable inquirers, then, do better in smaller communities.

But there is a significant tradeoff. The larger the community, the broader the dissemination of knowledge. Further, other things being equal, more agents means a better chance of there being exceptionally gifted inquirers. This matters since the community's average success rate is largely a function of the epistemic reliability of the most reliable inquirers. The most reliable inquirers evolve to observe nature directly, and are most often consulted by their colleagues in the evolved network. It is their success that drives the success of the community. This is seen in the fact that the correlation between the epistemic reliability of the most reliable investigator in the community and a typical agent's success rate is 0.704 while the correlation between a typical agent's reliability and that agent's success rate is just 0.200 in the twenty-agent model.

The reliability of each agent is randomly drawn with uniform probabilities over the unit interval. Let the reliability of an agent i be X_i , and let X_{max} be the reliability of the most reliable agent in the community. For a community of N agents, $X_{max} = \sup\{X_i\}$ for $i = 1, 2, \dots, N$. So $Pr(X_{max} \leq x) = Pr(X_{max} \leq x, i = 1, 2, \dots, N)$, since, if the most reliable agent has a reliability less than x , so do the rest of the agents. The reliability values are independent and identically distributed, so the cumulative distribution function is

$$Pr(X_{max} \leq x) = \prod_{i=1}^N Pr(X_i \leq x) = x^N.$$

And the probability density function is Nx^{N-1} , by differentiation. Integrating over the interval, the expected value of X_{max} is

$$E[X] = \int_0^1 x(Nx^{N-1})dx = \int_0^1 Nx^N dx = \frac{N}{N+1}.$$

The expected maximum epistemic reliability for an agent in the five- and twenty-agent models, then, is $\frac{5}{6} \approx 0.833$ and $\frac{20}{21} \approx 0.952$ respectively. The mean success rates of the two models after 200 generations 0.760 and 0.786 are both well below these bounds, but the difference suggests that the expected epistemic reliability of the most reliable agents in the larger community ultimately helps more than the distractions of the larger community hurt.¹²

¹²Note that it is the actual reliability of the most reliable agents in practice that matters, not their innate epistemic reliability. If the increased distractions faced by the most reliable inquirers undermine their de facto reliability, this, in turn, undermines the reliability of the entire community.

Figure 4 shows the evolution of the population success rate for the twenty-agent model and the five-agent model as the networks self-assemble over 200 generations.

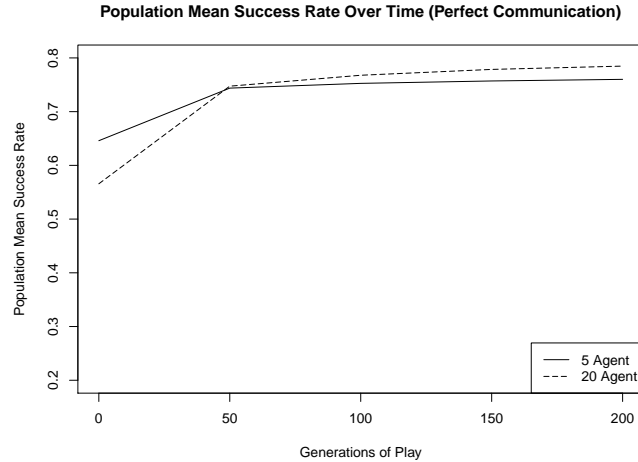


FIGURE 5. Mean Population Success Rate Over Time

That said, the increased distractions of a larger community undermines the reliability of the most reliable inquirers. The top quartile of reliable inquirers do better in the five-agent model than in the twenty, but they do not do as well as they would on their own since they still face distractions from less reliable colleagues in the smaller community. While one might be tempted to think of the difference between their success rate in the evolved network and the success rate they would enjoy on their own as a sacrifice for the sake of the epistemic good of the community, this would be to miss a significant point. Before inquiry, no one knows who is reliable and who is not. Indeed, given human nature as we find it, unreliable inquirers may well think themselves to be among the most reliable.¹³ Given this, the difference between the success rate of the most reliable inquirers in the evolved network and the success rate they would have if they investigated nature on their own is better thought of as the cost of finding out that they were in fact the most reliable inquirers and of being assured along the way that they will do nearly as well as their most reliable colleagues even if it turns out that they are not.

If the most reliable inquirers learn that they are in fact the most reliable, perhaps by noting the role they come to play in inquiry, they will benefit less from the

¹³See Kruger and Dunning [14] for evidence that suggests that both competent and incompetent agents are incompetent at assessing their own relative competence. The incompetent tend to believe they are better than they are, and the competent tend to believe that they are worse than they are.

insurance provided by the self-assembled social network. The most reliable agents may continue serving the epistemic goals of community if they gain something from their role in the evolved network that compensates for the distractions that comes with their interactions. Given that the expected cost of participation is small, if there were a small payoff for being an oft-consulted agent, that might keep the most talented inquirers engaged in the network. This suggests a positive role for *fame* in maintaining the stability of a successful epistemic network.¹⁴

It is also worth noting that even the most talented inquirers might find good reason to continue to serve as active members of an epistemic community under the plausible assumption that no single type of reliability is sufficient to address every problem posed by nature. The basic model might be extended to investigate this by providing the agents with various types of epistemic reliability, each appropriate to a corresponding type of problem that nature may present.¹⁵

5. IMPERFECT COMMUNICATION

The evolution of the social network is more subtle when one allows for social reliabilities less than 1. We will suppose here that each agent is also randomly assigned a social reliability in the interval $[0.5, 1]$ with uniform probability at the beginning of each run of the model. The resulting imperfect communication might represent poor eavesdropping abilities, a tendency toward misunderstanding, or just recurrent failures to convince agents to report their results truthfully.¹⁶

While the average population success rate declines with imperfect communication, the success rate of the most epistemically reliable agents is barely diminished. The most epistemically reliable agents are still consulting nature and doing as well as they would in the context of perfect communication, or, for that matter, on their own.

¹⁴There are, of course, other factors one might consider. An epistemic community might thrive even after losing its most talented inquirers if there were a regular source of new inquirers and if the community could learn to identify which new inquirers were reliable. Or nature may pose problems that require different types of expertise to resolve such that no agent might be entirely successful alone. Such considerations suggest extensions of the present model.

¹⁵One would expect a more flexible learning dynamics like win-stay/lose-randomize to outperform simple reinforcement learning on such a model inasmuch as the inquirers must learn which of their colleagues are most reliable for each type of problem and learn to consult those colleagues whenever that type of problem is presented by nature. A hybrid learning dynamics like win-stay/lose-randomize with reinforcement might be particularly well-suited to this task. The thought is that the reinforcement part of the dynamics might evolve to track the overall epistemic reliability of the agents while the flexibility of win-stay/lose-randomize would allow the inquirers to find quickly from among the most reliable agents those agents who have the right type of reliability for the current problem posed by nature. See [2] for a discussion of low-rationality hybrid learning.

¹⁶Note that the agents' expected epistemic reliability is higher than their social reliability since the former is selected over the interval $[0, 1]$ and the latter over the interval $[0.5, 1]$. On simulation, lower expected social reliabilities lead to the evolution of more independent inquirers since there are more agents who do better investigating nature on their own than consulting the most epistemically reliable agents.

For the range of social reliabilities we are considering, the average population success rates for the five-agent and twenty-agent models decrease from 0.760 to 0.637 and from 0.786 to 0.647 respectively. That said, the success rates for the top quartile of epistemically reliable agents stay relatively stable at 0.846 and 0.824 compared to 0.859 and 0.818 in the case of perfect communication. The bottom quartile of epistemically unreliable inquirers, on the other hand, suffer a significant drop in reliability in the evolved network from 0.732 and 0.780, to 0.584 and 0.556. Of course, this is nevertheless much better than the 0.125 success rate these agents would expect observing nature on their own.

The net effect of imperfect communication is to shift the evolved network toward a greater reliance on directly consulting nature and on consulting those agents directly who reliably observe nature. In the case of imperfect communication, the average relative frequency for a typical agent being consulted drops from 0.164 to 0.104 and from 0.047 to 0.033 in the five-agent and twenty-population models respectively. Epistemically reliable agents, however, experience an increase in the relative frequency of being consulted. Specifically, the relative frequency of being consulted for the top quartile of epistemically reliable agents increases from 0.217 to 0.226 and from 0.058 to 0.093 in the two models, while the relative frequencies for the bottom quartile drop significantly from 0.123 to 0.025 and 0.042 to 0.011. This effect can be seen in the comparison of Figures 3 and 6.

With imperfect communication, longer signaling chains are less reliable as each link compounds the possibility of error, hence consulting directly with the most epistemically reliable agents becomes comparatively more attractive and such agents attract the attention of their colleagues. Indeed, in the case of imperfect communication, finding someone who is good at consulting nature directly and learning from them directly is the best way to learn the truth for inquirers who are not themselves good at observing. In contrast, the best inquirers might remain relative anonymous in the context of perfect communication since the reliable results of their observations are widely available from even unreliable agents if they are well-connected to reliable agents.

While imperfect communication leads to a modest increase in the relative consultation rate of the most reliable agents, there is an order-of-magnitude, absolute increase in direct investigations of nature. This is clearly seen by comparing Figures 4 and 7. Here the average frequency of consulting nature increases from 0.178 to 0.479 and 0.069 to 0.346 in the five-agent and twenty-agent models. The frequency of consulting nature for the top quartile of epistemically reliable inquirers increases from 0.446 to 0.879 and 0.229 to 0.818, nearly doubling and quadrupling for the two models. In short, when communication is imperfect, the most epistemically reliable inquirers almost always observe nature directly. More generally, any




FIGURE 6. Epistemic Reliability and Frequency of Being Consulted with Imperfect Communication

inquirer whose epistemic reliability is higher than their social reliability does better observing nature directly.

In the context of imperfect communication, the inquirers in a signaling chain will only be successful if the initial inquirer is epistemically reliable and either everyone else is socially reliable or the chain is very short. Consequently, inquirers tend either to observe nature directly or to learn directly from reliable inquirers who themselves investigate nature directly. The effect of imperfect communication, then, is to shorten signaling chains and to increase the breadth of direct investigations of nature. As a result, a network that self-assembles in the context of unreliable communication, and is hence less successful overall, might nevertheless be more resilient inasmuch as the degree of success the community achieves depends on the direct investigations of a broader collection of inquirers.

Given the range of social reliabilities considered here, the resulting network typically exhibits a star structure with a few central reliable agents and everyone else linked to those central agents.¹⁷ Figure 8 illustrates a typical structure for the five-agent model. Figure 8(A) gives the full graph while figure 8(B) includes just the most significant edges (with weights above 0.2). The networks that self-assemble

¹⁷If one allows for a broader range of social reliabilities, one sees an increasing number of inquirers with low social reliabilities evolve to consult nature directly.

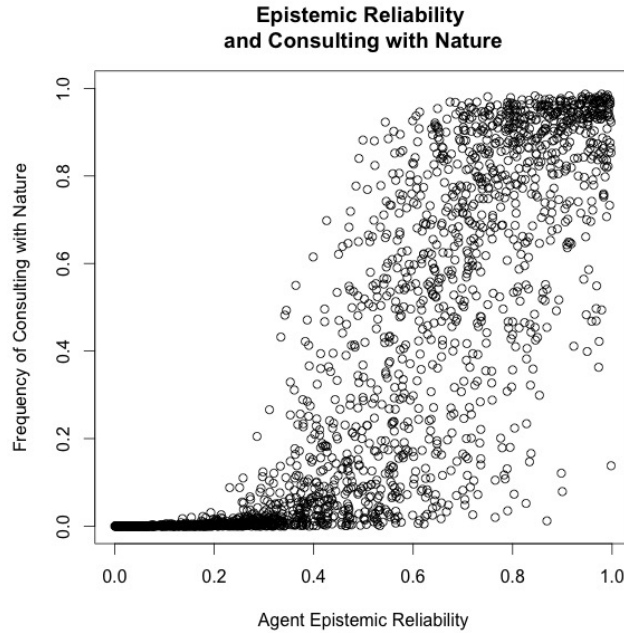


FIGURE 7. Epistemic Reliability and Frequency of Consulting Nature with Imperfect Communication

with imperfect communication contrast with those that evolve in the case of perfect communication, as illustrated in figure 2, by exhibiting a compact star structure with shorter signaling chains. On the simulation illustrated by figure 8, agent 4 was significantly more epistemically reliable than the other agents. When the reliability of the most reliable agents is close, the star structure may form around two or more agents. Such composite star structures, however, are corresponding less likely.¹⁸ Again, the network that self-assembles here is optimal insofar as figure 8(A) approximates the structure of figure 8(B).

With imperfect communication, the community must directly identify the most epistemically reliable inquirers since agents cannot count on the reliable dissemination of information. Here the most reliable agents are directly consulted the most often. They have evolved a reputation for reliability and play a central role in inquiry.

¹⁸See [7] for an extension of the Bala-Goyal model [1] where heterogeneous agents typically evolve networks with star structures. On that model, as here, the networks typically form around a “high-value” agent who helps to ensure the stability and efficiency of the evolving network. In the present model, the type of star structure that results is due to imperfect communication making it likely that agents will link to the most epistemically reliable agents directly.




FIGURE 8. Five-agent model with imperfect communication

Perhaps unsurprisingly, socially reliable inquirers are on average consulted more frequently than socially unreliable inquirers. However, the impact of social reliability on the frequency of being consulted is somewhat less significant than that of epistemic reliability. The top quartile of socially reliable inquirers are consulted by the population on average 0.140 and 0.047 of the time for the five-agent and twenty-agent models. In contrast, the bottom quartile of socially reliable inquirers were consulted by the population on average 0.083 and 0.023 of the time in the two models, for a factor of about two difference.

Considering the dispositions of the inquirers by type, socially reliable inquirers are on average more likely than other agents to consult their peers, while socially unreliable inquirers are more likely to observe nature directly. This effect is significant. The top quartile of socially reliable inquirers observe nature on average only 0.371 and 0.237 of the time and their peers the rest of the time, in the five-agent and twenty-agent models. Such inquirers do best by finding the most epistemically reliable inquirers, then using their social skills to glean reliable information from them. The bottom quartile of socially unreliable agents observe nature more frequently, at an average of 0.619 and 0.476 of the time, and their peers correspondingly less in the two models. Socially unreliable agents are sometimes most successful observing nature for themselves since they cannot reliably learn the results of even the most talented observers.

6. DISCUSSION

The modeled inquirers start by consulting each possible source of information indiscriminately. As they learn from their experience, however, they self-assemble a network by polymerization where they play mixed strategies that strongly favor consulting the most reliable inquirers or inquirers who are reliably connected to the most reliable inquirers. Typical inquirers in the self-assembled network do much better together than they could possibly do apart.

In the case of perfect communication, any agent, no matter how epistemically unreliable, may come to be a reliable source of information by finding a reliable inquirer to consult. This allows for arbitrarily long signaling chains. In the case of imperfect communication, however, the length of the signaling chains goes down and inquirers tend to either investigate nature directly or consult the most epistemically reliable inquirers directly.

In each case, the most reliable inquirers tend to observe nature directly, the most unreliable agents are rarely found to observe nature directly, and the most epistemically and socially reliable agents form the core of the resulting networks. While socially reliable inquirers are on average consulted more frequently than socially unreliable inquirers, the impact of social reliability on the frequency of being consulted is much less significant than that of epistemic reliability.

Imperfect communication leads the epistemic community to identify the most epistemically reliable inquirers, and these agents are directly consulted the most often. They are central to inquiry and famous for their reliability within the community.

There are epistemic trade-offs in the number of inquirers in the epistemic community. The more inquirers there are, the broader the dissemination of knowledge. A larger community also means a greater chance that there will be exceptionally gifted inquirers who might drive higher levels of success for everyone. But more inquirers also means more potential distractions for the most reliable inquirers, and, inasmuch as they drive the epistemic success of the community, what is bad for them is bad for everyone.

While the most talented observers may end up consulting less reliable inquirers from time to time in a self-assembled network, they nevertheless do nearly as well as they would had they been identified at the beginning of inquiry and allowed to observe nature without the distractions of their less competent colleagues. They might have done slightly better on their own, however, does not mean that they do not benefit from being a part of the general epistemic community.

In a community of heterogeneous agents, one must learn who is best at what in order to take advantage of differences in epistemic competence. And, insofar as

they are risk averse, learning this benefits even the most reliable agents. In the case of perfect communication, the most reliable inquirers pay a small cost in reliability for being a part of the evolving network, but the option is a situation where nearly everyone in the community does much worse. Consequently, insofar as the inquirers themselves do not know up front how reliable they are, the possibility of evolving a reliable network significantly limits each inquirer's personal epistemic risk. And the self-assembled network is only reliable when the inquirers learn what roles they are well-suited for and learn to play those roles. The present model shows how such a network may evolve by means of the simple reinforcement learning of individual agents.

The simulation results for both perfect and imperfect communication agree well with the equilibrium analysis. In particular, the random networks that evolve on simulation are well approximated by the deterministic Nash networks characterized in claims 1 and 2. This can be seen graphically by comparing the Nash network for perfect communication in figure 1(A) to the network that evolves with perfect communication in figure 2(B) and the strict Nash network for imperfect communication in figure 1(B) to the network that evolves for imperfect communication in figure 8(B).

In the case of perfect communication, one finds the agents evolving long chains connecting to the most reliable inquirer who, in turn, usually consults nature. There is little advantage for a given agent to identify and consult with the most epistemically reliable agent directly. In the case of imperfect communication, however, one sees the evolution of more compact star networks where a community of social reliable agents identifies the most epistemically reliable agents and consults those agents directly. Simple reinforcement learning does not always readily identify the single most epistemically reliable agent, especially if there other agents with nearly the same epistemic reliability. It does, however, typically allow those agents with the required social competence to identify the most epistemically reliable agents and learn from them. Agents with low social reliability typically evolve to consult nature directly.

There are two natural extensions of the present models. The first is to introduce *costs* for inquiry and allow agents to *charge* for sharing information. Just as there might be costs to checking nature or consulting with another agent, a particularly reliable agent might come to charge for sharing information or simply be rewarded for being often consulted. To implement this, one might allow for analog payoffs. This degree of freedom would also allow one to associate a random distribution of possible payoffs with each agent for checking nature or consulting another agent. One might expect such models to exhibit quite different dynamical properties and equilibria. A second extension would be to allow nature to present the agents with

different *types* of problems with heterogeneous agents who exhibit different degrees of competence for each type of problem. To be successful, the agents would need to coevolve the ability to identify the type of the problem presented and to consult agents who are good at solving that particular type of problem. In this case, one might expect a conditional network to evolve where agents identify different experts for each type of problem and learn to consult experts accordingly. Since each of these extensions requires a learning dynamics that is more sophisticated than simple reinforcement, these models will be discussed elsewhere.

Here we have shown how an epistemic network may spontaneously self-assemble from the evolving dispositions of even very simple, low-rationality agents. The resulting epistemic network is well customized to the strengths and weaknesses of the heterogeneous agents. The most reliable inquirers do nearly as well as they would simply consulting nature directly, other agents identify the most epistemically reliable agents and hence do much better than they could possibly do on their own, and those agents whose social reliability is very low follow their best strategy and investigate nature directly.

REFERENCES

- [1] Bala, V. and Goyal, S. [2000] ‘A Noncooperative Model of Network Formation,’ *Econometrica* 68:71181–1229. doi:710.1111/1468-0262.00155.
- [2] Barrett, J. A., C. T. Cochran, N. Fujiwara, S. Huttegger [2016] ‘Hybrid Learning in Signaling Games,’ draft paper.
- [3] Barrett, J. A. (2014) “Rule-Following and the Evolution of Basic Concepts,” *Philosophy of Science* 81(5): 829–39.
- [4] Barrett, J. A. and B. Skyrms [2015] ‘Self-Assembling Games,’ *British Journal for the Philosophy of Science*. First published online 13 September 2015. doi: 10.1093/bjps/axv043.
- [5] Barrett, J. A. and K. Zollman [2009] ‘The Role of Forgetting in the Evolution and Learning of Language,’ *Journal of Experimental and Theoretical Artificial Intelligence* 21(4): 293–309.
- [6] Erev, I. and A. E. Roth [1998] ‘Predicting How People Play Games: Reinforcement Learning in Experimental Games with Unique, Mixed Strategy Equilibria’ *American Economic Review* 88: 848–81.
- [7] Goeree, J. K. and A. Riedl and A. Ule [2009] ‘In Search of Stars: Network Formation among Heterogeneous Agents’ *Games and Economic Behavior* 67(2):445–466.
- [8] Gross, T. and H. Sayama (eds.) [2009] *Adaptive Networks: Theory, Models and Applications*. Dordrecht: Springer.
- [9] Herrnstein, R. J. [1970] ‘On the Law of Effect,’ *Journal of the Experimental Analysis of Behavior* 13: 243–266.
- [10] Huttegger, S., B. Skyrms, P. Tarrès, and E. Wagner [2014] ‘Some Dynamics of Signaling Games,’ *Proceedings of the National Academy of Sciences* 111(S3): 10873–10880.
- [11] Huttegger, S. and B. Skyrms [2013] ‘Emergence of a Signaling Network with Probe and Adjust,’ in *Cooperation and its Evolution*. B. Calcott, R. Joyce and K. Sterelny (eds.), Cambridge, MA: MIT Press, 265–274.
- [12] Kitcher, P. [1990] ‘The Division of Cognitive Labor,’ *The Journal of Philosophy* 87(1), 5–22.
- [13] Kitcher, P. [1993] *The Advancement of Science*. New York: Oxford University Press.
- [14] Kruger, J. and D. Dunning [1999] ‘Unskilled and Unaware of It: How Difficulties in Recognizing One’s Own Incompetence Lead to Inflated Self-Assessments,’ *Journal of Personality and Social Psychology* 6:1121–34.
- [15] McGregor, P. [2005] *Animal Communication Networks*. Cambridge: Cambridge University Press.
- [16] Pemantle, R. and B. Skyrms [2003] ‘Network Formation by Reinforcement Learning: The Long and Medium Run,’ *Mathematical Social Sciences* 48(3): 315–327.
- [17] Rosenstock, S., C. O’Connor, and J. Bruner [2016] ‘In Epistemic Networks, Is Less Really More?’ draft paper.
- [18] Roth, A. E. and I. Erev [1995] ‘Learning in Extensive Form Games: Experimental Data and Simple Dynamical Models in the Immediate Term,’ *Games and Economic Behavior* 8:164–212.
- [19] Skyrms, B. [2010] *Signals: Evolution, Learning, & Information*. New York: Oxford University Press.
- [20] Skyrms, B. and R. Pemantle [2000] ‘A Dynamic Model of Social Network Formation,’ *Proceedings of the National Academy of Science* 97: 9340–9346.
- [21] Strevens, M. [2003] ‘The Role of the Priority Rule in Science,’ *Journal of Philosophy* 100(2): 55–79.
- [22] Strevens, M. [2013] ‘Herding and the Quest for Credit,’ *Journal of Economic Methodology* 20:19–34.

- [23] Zollman, K. [2012] 'Network Epistemology: Communication in Epistemic Communities,' *Philosophy Compass* 8: 15–27
- [24] Zollman, K. [2007] 'The Communication Structure of Epistemic Communities' *Philosophy of Science* 74(5): 574–587

UC IRVINE, IRVINE, CA 92697, USA

E-mail address: `j.barrett@uci.edu`

E-mail address: `bskyrms@uci.edu`

E-mail address: `aydin.mohseni@gmail.com`