Beliefs and Social Networks

Kevin J.S. Zollman

Abstract: Network epistemology is a growing field which studies the relationship between social network structure and belief. The field draws on work from many disciplines, including computer science, economics, philosophy, physics, political science, psychology, and sociology. While the conclusions of the field suggest that the relationship between network structure and belief is quite complex, there are a few general lessons that can be drawn. This chapter discusses what is known, and where fruitful new interdisciplinary work could help expand our understanding.

Brief Biography: Kevin J.S. Zollman is the Herbert A. Simon Professor of Philosophy and the Director of the Institute for Complex Social Dynamics at Carnegie Mellon University. His research focuses on mathematical models of social behavior including the study of information flow in social networks.

Keywords: Network epistemology, social networks, belief, learning

Introduction

Social epistemic pathologies like polarization, group-think, economic bubbles, and misinformation cause significant real-world harms beyond false beliefs. Studying these pathologies involves two main approaches: one examines individual-level causes, such as social psychological effects, and how social interaction alters beliefs. The other investigates larger structural causes, like communication patterns, to understand belief formation within the broader epistemic environment.

Network Epistemology is of this second type. It is the study of how social network structure influences belief acquisition, transmission, and abandonment. It covers both pathological and non-pathological societies and gives us a lens through which we can understand many of the social processes that spread beliefs.

Social networks vary significantly in structure (see for example, figure 1). Some are highly connected while others are sparse. Some feature massive inequality in social connections, while others are more uniform and homogeneous. Some feature clustering, where one’s friends are friends with one another, while others have very little clustering. Understanding how these differences affect the spread and acquisition of belief is the central question of network epistemology.

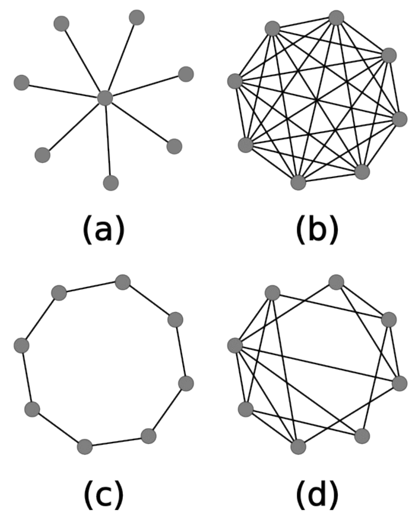

Figure 1: An illustration of four networks (a) the star, (b) the complete network, (c) the ring, and (d) a small world network.

A broad conclusion of this research program is that there is not a simple relationship between network structure and belief. Rather, depending on the type of learning problem and other features of the community, radically different networks might cause or prevent epistemic pathologies. As we will see, even very simple questions, like “are more connections good or bad for the uptake of reliable beliefs?”, do not have unequivocal answers.

All is not lost, however. Once some particulars are specified, both models and experiments appear to provide robust conclusions about how social network structure affects beliefs. Of course, much remains unstudied. There is ample opportunity for novel theoretical and empirical work.

Network epistemology is an interdisciplinary field which studies the relationship between social networks and the beliefs of its members. In keeping with this wide interdisciplinarity is diversity in methods.

Because of the complexity of the studied system, theoretical methods are common. Both mathematical and computational models are ubiquitous. Some models, especially from economics, are based on expected utility maximization, while others use “boundedly rational” agents. Agent-based models simulate communities with agents following simple rules and interacting via social networks. Scholars analyze the resulting social dynamics to identify patterns linking agents and their networks to aggregate behavior.

These models are critical because they allow us to investigate changes in large-scale social structure which would be difficult or impossible to study via traditional experimental methods. While it might be possible to design a laboratory experiment that altered the social network among five or ten people, doing so for a thousand or ten thousand would be difficult. Doing so at the scale of online social media is simply impossible. Furthermore, these methods allow us to understand the relationship between individual behavior and social outcomes in ways not possible with experiments. One cannot ask subjects to cease being influenced by confirmation bias (for example) in order to understand the social effects of confirmation bias. With models this is possible.

While the starting point of the literature is often formal models, scholars connect these models with data. Lab experimentation is common, where human subjects are placed into an artificial setting that mimics a model (this mimicry can come in varying degrees of fidelity).

Models are tested by experiment in two ways. First, researchers compare the particular design assumptions of the model to experiments. For example, suppose an agent-based model features agents who average their opinion with the opinions of others around them. An experiment might be conducted to determine the degree to which “averaging” is a good model for how those subjects behave in analogous laboratory settings.

Second, experiments compare macroscopic predictions of a model to macroscopic outcomes of an experiment. For example, if a model predicts that one network structure leads to more accurate beliefs than another, the experiment might seek to confirm the superiority of the better network using human subjects. The macroscopic results might match even if the microscopic processes do not.

Beyond laboratory experiments, a model may also be compared to observational data in less contrived settings. So, for example, a model might predict the rate of adoption of a new belief or technology. Data from sales or from Twitter might be used to validate the model.

In the sections that follow, I will present several classes of models that differ in terms of the underlying learning problem being modeled. Changing the learning problem focuses on a different aspect of social influence on belief. The first two sections consider models of evidence amalgamation, where there are different ways a community might pool information they already have. The third considers a model of how new information flows through a network. In all these models, information merely arrives. The fourth section considers models of information search, where people must actively seek out new information. Information exchange can sometimes be quite complicated, with people offering reasons or arguments for their beliefs. The fifth section describes models of this form. Many of the models of the first five sections treat the social network structure as fixed. In the penultimate section, we consider models where the social network’s structure is malleable. Of course, this is not a complete review of the literature, but rather an illustration of the breadth of problems, methods, and conclusions.

Information cascades

Models of *information cascades* in economics represent an early example of network epistemology (Bikhchandani, Hirshleifer, and Welch, 1992, Banerjee 1992). Consider a simple example: There is a bag filled with red or blue marbles. You and a collection of other people will each privately draw a small number of marbles. Every person will make an announcement about whether they think the bag contains more red marbles or more blue marbles in total. You will act in a predetermined order, and each person will hear the announcement of those who act before them. How will (or should) someone act in this setting? How should their announcement depend on those before them? And what will be the aggregate reliability of the group in announcing the correct answer?

While highly stylized, this model is meant to capture a social dynamic where the spread of belief is mediated by action. An investor might only learn about the beliefs of others through observing stock purchases. I might only learn my friend's food preferences by seeing her order food. Or a politician might only come to know what voters want by observing who they vote for.

Information cascades: Theory

Each agent has a prior probability estimate of marbles in the bag. They update both on the information they received and on the announcement of others using Bayes' theorem (see BAYESIAN AND RATIONAL CHOICE APPROACHES TO MEASURING CONFIDENCE and A PREDICTIVE ARCHITECTURE FOR THE ATTITUDES in this handbook). Each agent will be rewarded if their personal answer is correct. When it comes to their turn, they choose an announcement which will maximize their expected utility relative to their beliefs.

The striking result is that in finite groups, it is possible for the group to converge to the wrong answer despite the fact that the information, in aggregate, was not misleading. Even an infinitely large group can converge on the wrong answer. Many find it surprising that groups of fully rational individuals, arranged in this way, do not optimally integrate the available information.

To illustrate, consider a simple example. Everyone will have the opportunity to draw three marbles. You are fourth in line, and you draw two blue and one red marble. You form the belief that there are more blue than red marbles. Before you announce, you are allowed to see the announcement of three other people, and all three announce "red."

You know that the first person must have seen more red than blue marbles. So they must have either seen 2-1 or 3-0 for red. That means that combined with your information, it's either 3-3 or 4-2 for red in total. Already after seeing that first announcement, you would announce "red."

The second person also announces red. You have less information about what they saw, because it's possible that they got the same draw you did and reasoned in the same way you did about the first person. You can (maybe) eliminate the possibility that they got a draw of 3-0 for blue. So, this will slightly increase the chance that red is the correct answer in your mind.

The third person also announces red. Their announcement gives you very little information about their draw, since they would say "red" regardless. When it's your turn, you announce red. So does everyone else after you as well because they reason similarly to you.[1]

Information cascade models show how rational individuals could be party to economic bubbles, fads, and other groupthink-like behaviors. These models illustrate that individual-level cognitive biases are not *necessary* for social pathologies.

Information cascades: Experiments

The information cascade framework is also an experimental methodology. Individuals are placed into this setting and their behavior is observed. Subjects are paid according to whether their private statements match the secret underlying state.

---

[1]This does depend on the assumption of “bounded signals.” If there is some chance that somebody can get a signal so strong that it overwhelms any amount of social information to the contrary, then groups will eventually converge on the correct answer (Smith & Sørensen, 2000).

At the group level, the predictions of the model are largely confirmed. Groups do not take full advantage of the information that was supplied to it (Anderson & Holt, 1997, Hung & Plott, 2001, Kübler & Weizsäcker, 2005, Ziegelmeyer et al., 2010). Although very long cascades may behave differently than the theory predicts (Kübler & Weizsäcker, 2005, Goeree et al., 2007).

At the individual level, results are mixed. Agents do not respond exactly as predicted by the Bayesian model. Sometimes this is beneficial to the group, when an individual ignores the social information and acts on their private information (Anderson & Holt, 1997, Hung & Plott 2001). This generates information that benefits those behind the person who deviates.[2] In contrast, subjects do not realize that they should ignore the statements by people later in the cascades (Ziegelmeyer et al., 2010). This worsens the cascade effect.

Several experimental variations have been conducted. One might change the way subjects are rewarded: by paying them if the majority of their group guesses correctly or if their guess conforms with the guesses of others. Hung and Plott (2001) found that people behave differently depending on their incentives. Rather

[2] I won't call this "irrational" because they may be acting altruistically, although experiments have not confirmed that they are.

than receiving private information for free, subjects might be asked to pay to receive their private information. Kübler and Weizsäcker (2004) find that people pay too frequently, showing that they overvalue private information relative to the Bayesian ideal. Ziegelmeyer and colleagues (2010) give subjects a “stronger” signal that should be sufficient to break the cascade but find that it does not always do so.

Summary of information cascades models

The information cascade models and experiments highlight a couple of important themes. First, social pathologies—in this case something like group-think or economic bubbles—might be produced by ideally rational agents. This illustrates that if we want to explain these phenomena, we should not necessarily rely on biases in individual belief formation.

Second, the model illustrates that in some cases, sharing of information might be harmful. Each agent learning what others are doing in front of them makes them *individually* better, but makes the group as a whole worse. This has the same basic structure as a public goods problem, where individual selfish behavior harms everyone.

Models of social influence

The information cascades model is a good model for interactions that take place over time (like the purchasing of stock in a company). Of course, much of belief change takes place in a less structured setting. I might tell you my opinion about how the Pittsburgh Steelers will perform in the coming NFL season today. You might learn others' opinions tomorrow. Maybe in a few days, you might tell me your opinion on the matter. These interactions feature more symmetric social influence, where information can go both ways.

A different model of belief sharing is needed. In this section, all the models and experiments will have a similar structure: agents form an initial belief about some question, either entirely based on their own reasoning or from some information they are provided. For example, everyone might form an opinion about the prospects of the Pittsburgh Steelers by independently evaluating their past performance.

No additional information will arrive from the outside, the remaining part of the process involves aggregating their beliefs. Agents are arranged in a social network that represent how declarations are shared. Each agent declares a belief to their neighbors, listens to the declaration of their neighbors, and potentially changes their belief. The process is repeated (perhaps indefinitely) until the beliefs of all the agents stabilize.

Pooling models

Modeling symmetric social influence with full Bayesian individuals is possible, but extraordinarily difficult (Acemoglu & Ozdaglar, 2011, Sadler, 2014, Hązła et al., 2021). Consider what would be involved. Agents must have subjective probability distributions over all possible distributions of information, and all possible declarations of every person in the network. If they are unaware of the network structure, they must have a distribution over that as well.

Instead of trying to model all this uncertainty, many scholars have turned to “boundedly rational” models of belief change. Two simple models have been reinvented often.

In the first model, belief is a binary state (you either believe something or you do not) and each agent changes their belief to match the majority of their neighbors.

This model has many names; I will call it the repeated majority model.[3] The second model represents belief as a continuous variable (often in [0,1]). Each agent updates their beliefs by taking a linear average of the beliefs of their neighbors. This is called the linear pooling model.[4]

Both models entail that agents will "double count" information they receive. Since they respond to their neighbors in every round, this will lead to them to count opinions of others multiple times. More highly connected people will be over-counted relative to the less connected.

Although the models are quite similar, they come to strikingly different conclusions about the long run. In the continuous model, all agents converge on a single consensus opinion under very weak assumptions about the social network and the averaging procedure (Demarzo, Vayanos, & Zwiebel, 2003, Golub & Jackson, 2010, Golub & Sadler 2016). The discrete belief model will often stabilize

---

[3] This model has been reinvented too many times to count. For a list of references see (Zollman 2010a).

[4] Like the first, this model has been reinvented several times. The most prominent first inventors are (French 1956; DeGroot 1974; Lehrer and Wagner 1981).

with heterogeneous beliefs, especially when the number of social network connections is relatively low (Zollman, 2010a, Zehmakan, 2020).

When groups are judged by the reliability of the majority, however, both models predict that *regular* networks will be closer to the truth. If all agents are equally reliable in their initial opinions, regular social networks are those with no inequality in social connections. The ring and complete graph (figure 1b and 1c) are regular, while many other networks, like the star and small world graph (figure 1a and 1d), are not. When agents are not equally reliable, however, the distribution of network connections should correspond to the difference in reliability. The most reliable networks involve a correlation between individual reliability and the number of network connections (Demarzo, Vayanos, & Zwiebel, 2003, Golub & Jackson, 2010, Zollman, 2010a, 2011).

A well-studied variation on the linear pooling model requires agents to ignore opinions that are too divergent from their own. This model, called the bounded confidence model, represents a relatively minimal change to micro-process that leads to significant macroscopic changes. No longer is consensus guaranteed. The model produces certain types of polarization and echo chambers (Hegselmann & Krause, 2002, Hegselmann, 2023).

Pooling models: experiments

Two experiments tested whether subjects arranged in social networks are well modeled by the repeated majority model (Grimm & Mengel, 2020; Chandrasekhar, Larreguy, & Xandri, 2020). Many subjects followed the repeated majority model rather than the full Bayesian one, especially when subjects were less informed about the structure of their social networks.[5]

Brandts, Giritligil, and Weber (2015) and Becker, Brackbill, and Centola (2017) tested a macroscopic prediction of the pooling model. Both experiments confirmed that the group estimates were biased toward the initial guesses of people with more connections. In contrast, while Corazzini and colleagues (2012) found that people do not appropriately discount repeated information, the macroscopic results from two different networks cannot be explained by the linear pooling model.

---

[5] Choi, Gale, and Kariv (2005) conducted a similar experiment, but with only three-person directed graphs. They found behavior was largely (but not entirely) consistent with full Bayesian rationality. It is possible that people are capable of more sophistication when the size of the groups is relatively small.

Granovskiy and colleagues (2015) perform an experiment that mimics the pooling model but where there was no network structure: everyone interacts with everyone. They develop a model of their data which shares some, but not all, features with the bounded confidence model, where people ignore some opinions.

Jönsson, Hahn, and Olsson (2015) perform a similar experiment with network structure. In one of their experiments, they found that the small world network (see figure 1d) outperformed the complete network in terms of both individual and collective error. This result contradicts the pooling model where the complete network (figure 1b) is predicted to be among the best.

Trust and ongoing information acquisition

The models discussed so far leave out two important parts of some real-world deliberations. First, information might arrive during the process of deliberation. While we are discussing the prospects of the Pittsburgh Steelers, for example, we might learn some news about the players. Second, we might also regard some people as less reliable sources of information than others. My neighbor is always pessimistic about the Steelers, so I might discount his opinion more than someone else.

There are three models which incorporate this possibility in very different ways. Interestingly, they come to broadly similar conclusions. Adding these two features are sufficient to produce polarization.

Axelrod (1997) was the first to model polarization.[6] His agents live on a checkerboard. An agent's beliefs are represented as a vector of numbers. Agents imitate the beliefs of others who are near them in the social network, but they are more likely to imitate agents who have similar beliefs across the whole vector. This simple process instantiates homophily—where agents tend to interact with others who are like them. Homophily in this form creates the conditions for a structured and polarized society. A related model comes to similar conclusions (Weatherall & O'Connor, 2020).

---

[6] While a common term, "polarization" is actually the name for a large cluster of different concepts. Detailed discussion is beyond the scope of this paper, but interested readers should consult (Bramson et al., 2017).

The Laputa framework is one that allows for new information and dynamic trust (Olsson, 2011). Agents receive information both from each other and from the world. Agents have heterogeneous reliability, and they each have a conjecture about others' reliability. They update these in a broadly Bayesian manner. The Laputa framework has been employed to analyze several themes including polarization (Olsson, 2013, 2020, Hahn, Hansen, & Olsson, 2020), trust and misinformation (Morreau & Olsson, 2022, Vallinder & Olsson, 2014), and network structure (Hahn, Hansen, & Olsson, 2020). Aspects of the model have been confirmed empirically (Collins et al., 2018).

Abstracting away from the micro-process, one might also model agents as "pulled" toward the truth. The bounded confidence model has been adapted in this way (Hegselmann & Krause, 2006). Like with the Laputa model, the bounded confidence model can capture lack of trust, although the micro-process is modeled quite differently. It can also explain polarization (Hegselmann & Krause, 2002, Hegselmann, 2023) and be used to study misinformation (Douven & Hegselmann, 2021).

These three models are a study in contrasts. All three model similar things, but the micro details diverge. The Laputa model is Bayesian (albeit bounded) and has many parameters (Olsson, 2011). The bounded confidence model has only a small handful of parameters. It does not aim to model individual behavior exactly, but instead abstracts away from that to generate a more macro-scale model which can capture relevant phenomena (Hegselmann & Krause, 2006). The Axelrod model is like the bounded confidence model, but models several simultaneous beliefs.

Despite their structural differences, all three models predict polarization via a similar underlying dynamics: agents form separate communities that have a high degree of internal trust but with low trust in other communities. In the Laputa model this happens because agents are trying to learn one another's reliability. In the bounded confidence model and the Axelrod model, this is imposed by assumption.

Summary of social influence models

Social influence models illustrate the dangers of highly unequal social connections. Groups that feature highly varied networks (like online social networks), will often fail to aggregate information in an optimal way. Rather, more "flat" social structures are superior.

When we add the arrival of new information with the ability to gauge the reliability of another, we also find a clear explanation for some types of polarization. That relatively simple models can lead to polarization provides some evidence that larger scale social dynamics may be more to blame for polarization than individual-level psychological effects.

Contagious ideas

The pooling and voting models of the previous section are meant to capture deliberation about something where everyone has an opinion. Another type of "belief spread" occurs when new information is transmitted through a society where no one has thought about the matter before. In the morning, you may not yet have an idea of what a controversial politician said the night before. But, as you talk to friends, look at social media, or read the paper, you might come to know the ridiculous things he said.

This type of belief spread mimics the spread of a disease, and epidemiological models have been adapted to this task (for more discussion see BELIEF PROPAGATION ACROSS SOCIETY in this handbook). Adopting a new belief is like catching a disease. Once you believe it, there may be a period of "infectiousness" where you tell other people what you now believe. Forgetting the belief, adopting a contradictory belief, or finding it no longer newsworthy, is the equivalent of recovering from a disease.

Contagions: theory

The classic epidemiological model, the SIR model, uses differential equations and ignores social network structure. It categorizes everyone into three “bins” as susceptible to a disease, currently infectious, or recovered. The size of each group changes deterministically: the infectious make susceptibles infectious and recover after a certain period.

This model is also common in the study of *innovation diffusion*.[7] Since our focus is on network structure, we will set this model to the side to focus on network models.[8] Many models have been developed with different diffusion micro-

---

[7] Bass (1969) was the first to employ an epidemiological model to study innovation diffusion.

[8] See (Meade & Islam, 2006) for a survey of non-network-based models and their connection with data. The first to incorporate networks in epidemiology was (Hethcote, 1978). The first in innovation diffusion was (Abrahamson & Rosenkopf,

processes (Young, 2009, Dunn & Gallego, 2010, Abrahamson & Rosenkopf, 1997, Kuandykov & Sokolov, 2010, Goldenberg & Efroni, 2001, H. Choi, Kim, & Lee, 2010, Delre et al., 2010, Delre, Jager, & Janssen, 2007, Hohnisch, Pittnauer, & Stauffer, 2008, Bartal, Pliskin, & Tsur, 2020).

These models can be interpreted in two different ways. If the underlying belief is an innovation—i.e., a genuinely good piece of information—the ideal network will be one that maximizes the spread in minimal time. Alternatively, the belief might be misinformation, in which case the best network is one that contains or slows the spread.

This contrast illustrates a difficulty of regulating misinformation in online social networks. Without knowing if the information is good or bad, we simply cannot design anything like optimal social networks.

If our goal is to spread information as quickly as possible, the optimal network is clearly the complete network. More contact means more opportunity for information to be transmitted. If we want the best network with the fewest

---

1997).

connections, the star network (figure 1a) can also be good because all nodes are close to one another.[9]

LaCroix, Geil, and O'Connor (2021) use a double-diffusion model to understand why false information spreads more widely than a correction. For example, why does a new scientific result spread further than a retraction or failed replication of that result? They model the initial sharing as an innovation diffusion model, but assume that the retraction is only taken up and shared by someone who has been initially "infected" by the false idea. This model elegantly shows why combating misinformation is so hard: corrections are less contagious than the initial misinformation.

The traditional diffusion model inherits the assumption from epidemiology that one can take up an idea after only a single contact. (Or to say it more precisely, the probability that one believes the idea is independent given the number of times one has been exposed.) However, uptake of ideas may not follow this pattern. Instead, someone may wait until they've heard an idea multiple times before

[9] Here I focused on the speed of the transmission of the idea. But there are also other things that one might consider, like how many people will eventually adopt the idea in an arbitrary random graph (Newman 2002, 2018).

believing or transmitting it to others. Models of “complex contagions” allow for this possibility (Centola & Macy,2007).

Complex contagions are encouraged by very different network structures. If someone only adopts a new belief if they hear it from two or more people, then the star is completely inefficient. No information will spread because only the center person could possibly be told the same fact from two different people. Instead, networks with high clustering are needed (Centola & Macy, 2007). In simple contagions, clustering slows down the spread of information. In complex contagions it facilitates spread.

Contagions: connections with data

Diffusion of new information is much more difficult to study in a laboratory setting. More common are natural experiments.

Substantial work has matched non-network models to data, especially in the context of diffusion of new products (Meade & Islam, 2006, Rogers, 2003). Some have attempted to explicitly match network diffusion models to data, most often data from online social networks (Hohnisch, Pittnauer, & Stauffer, 2008, Goel et al., 2016, Takayasu et al., 2015, Myers, Zhu, & Leskovec, 2012).

Work has been done testing some assumptions of those models against data (Lerman & Ghosh, 2010). In particular, many studies have focused specifically on whether ideas follow the “simple” or “complex” contagion model (Centola & Macy, 2007, Centola 2018). There is some evidence that online social media functions more like a complex than a simple contagion (Romero, Meeder, & Kleinberg, 2011, Mønsted et al., 2017, Fink et al., 2016, 2015), although there is disagreement (Hodas & Lerman, 2014).

Summary of contagion models

The study of contagion models illustrates how small changes in the micro-level process can generate big changes on the macro-level. For example, the shift from “simple” contagions—where a single exposure can make you believe something—to “complex” contagions—where you must hear something several times before you believe it—can radically alter the relationship between network structure and the speed of belief spread. In simple contagions, clustering leads to slower spread, and in complex contagions, clustering leads to faster spread.

Supplementing these models with a model of correcting misinformation also shows how complex interactive processes can be. This model provides some micro-structural explanation for the common observation that lies spread faster than the truth (LaCroix, Geil, & O'Connor, 2021).

Modeling search

Common to all the models discussed above is that information arrives independently of the beliefs of the agents themselves. Search problems create a complex relationship between beliefs, actions, and what information arrives in the future. They do this by creating some version of explore—exploit tradeoffs.

There are two main modeling paradigms that emphasize different aspects of the search process. The first is the *bandit problem* model. Named after slang for slot machines (a "one armed bandit"), bandit problems give agents a choice among a small number of options. Each option is characterized by a probability distribution over potential payoffs. Agents want to play the option with higher average payoff as much as possible (called "exploitation"). They don't know, however, which option is optimal. So they must explore the various options to determine which is best.

The second paradigm, known as a *landscape search problem*, models a search problem involving many (often hundreds or thousands) of options arranged in a mathematical space. Each option provides a payoff for anyone who tries it. The agents must explore the space in some intelligent way in order to find the option with the highest payoff.

In both these models, we can devise a notion of problem difficulty. In bandit problems, if distributions have large variances or the payoffs of the two options are very close, it is difficult to identify the superior option. If the payoff difference is large and variance small, the problem will be quite easy.

In landscapes, the difficulty is determined by how “rugged” the landscape is. If the landscape is smooth (as in figure 2a), any reasonable search procedure will find the optimal solution. If the landscape features many local maxima separated by valleys (as in figure 2b), then it might be quite difficult.

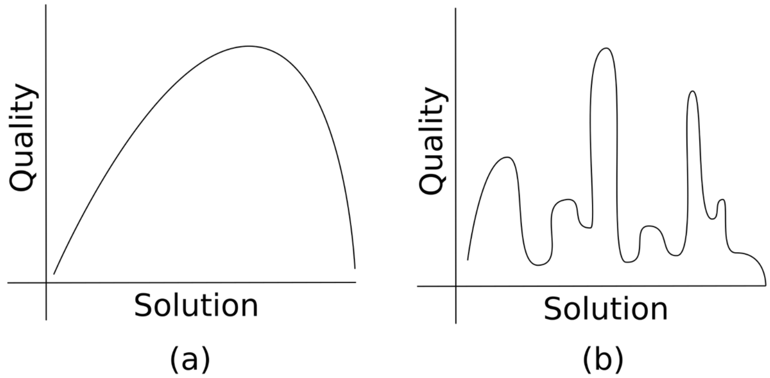

Figure 2: A heuristic illustration of two types of landscapes. The x-axis represents the space of solutions as a one-dimensional space. The y-axis represents the quality of an option. (a) illustrates a relatively easy "smooth" landscape, where a hill climbing algorithm will easily find the global optimum. (b) illustrates a "rugged" landscape, where hill climbing algorithms risk getting stuck in local optima that are not globally optimal.

Modeling search: Theory

One central finding from both paradigms is that diversity of search is beneficial (Zollman, 2010b, Wu, 2023a, Wu & O'Connor, 2023, Hong & Page, 2004, Thoma, 2015, Grim et al., 2019, Barkoczi & Galesic, 2016, Lazer & Friedman, 2007, Campbell, Izquierdo, & Goldstone, 2022, Pöyhönen, 2017). In both cases, distributing members of a group among many different options will help the group find the best solution. For bandit problems, diverse behavior limits the premature abandonment of one of the options. In landscape search, diversity ensures a larger amount of the terrain is explored.

The value of diversity of behaviors in these models depends critically on the difficulty of the problem (Rosenstock, Bruner, & O'Connor, 2017, Grim et al., 2019). Diversity of behavior or opinion is not important in easy problems that feature a smooth landscape or bandits with radically different payoffs.[10] Individuals facing easy problems are likely to solve them on their own, and so maintaining diversity has little epistemic benefit.

In difficult problems, diversity is beneficial. This diversity can be maintained at the individual and the collective levels. At the individual level, people can adopt strategies that encourage exploration. However, there is a complication. Exploration in these models is a type of public good: each individual takes on a cost by exploring potentially inferior options, while the group benefits from it. As a result, individuals alone may generate insufficient diversity for optimal exploration (Bolton & Harris, 1999, Kummerfeld & Zollman, 2016).

---

[10] Here we are focused on diversity of opinion or behavior, not on demographic diversity. There is a complicated relationship between the two, which is beyond the scope of this chapter.

Social strategies for maintaining diversity are more counter intuitive. One can encourage experimentation by limiting group information through a sparser social network (Zollman, 2007, 2010b, Grim et al., 2013, Wu, 2023a), by giving individuals “unreasonably” strong initial opinions (Zollman, 2007), making them ignore certain subcommunities (Wu, 2023b, Fazelpour & Steel 2022), and making them subject to confirmation bias (Gabriel & O’Connor, 2024).

Many of these social mechanisms come with serious costs. In order to preserve diversity you must also slow down achieving consensus. Whether that cost is worth paying depends on particulars of the learning problem the group is confronting. In addition to slowing down learning, some of these social processes are particularly likely to lead to polarized communities, which might have its own negative effects (Wu & O’Connor, 2023).

Modeling search: Experiments

Many experiments have been conducted on how individuals choose in the context of explore-exploit decisions when acting alone (for examples, see Steyvers, Lee, & Wagenmakers, 2009, Lee et al., 2011, Zhang & Yu, 2013a, Daw et al., 2006, Speekenbrink & Konstantinidis, 2014, Zhang & Yu 2013b). Because they focus on individuals acting alone, they cannot be uncritically applied to social settings. They can, however, inform how models of this type are built.

Wisdom and Goldstone (2007) look at social learning strategies in landscape search problems. They find that individual strategies for learning from others are extremely complex, including changes in imitation behavior over time and across social settings.

However, large scale predictions of these models are consistent with experiments. Several experiments have placed subjects in social networks and asked them to solve landscape search style problems. They find that limiting information can sometimes improve the performance of the group by encouraging more diverse behavior (Mason, Jones, & Goldstone, 2008; Mason & Watts, 2012, Derex & Boyd, 2016, Charness, Cooper, & Grossman, 2020, Wisdom, Song, & Goldstone, 2004, Goldstone et al., 2013).

Summary of search models

Both landscape search problems and bandit problems provide a useful testbed for social learning when information must be actively acquired (rather than arriving to passive observers). The results from models and experiments have consistently found a benefit for behavioral or opinion diversity when problems are difficult. There are many potential mechanisms for maintaining this diversity, and research about their various costs and benefits is ongoing.

Argumentation models

Almost all the models discussed so far focus on a single belief. That belief might be about the optimal bandit, the best point in a landscape, a piece of gossip, or the state of the world. Of course, real social interaction often involves the communication of different interconnected beliefs. You might try to convince me to believe something by presenting me with an argument involving several other beliefs that I already hold. Or, I might come to agree with you about one proposition because I agree with you about several others. A few models have been developed to account for this (for a more comprehensive overview, see De Tarlé, et al., 2025).

Dung (1995) developed a simple model of argumentation, which represents an argument as a series of claims that “attack” other claims. Borg and colleagues (2017, 2018) apply this model to a kind of landscape search problem. In contrast to some of the landscape search problems discussed in the previous section, Borg and colleagues find that whether communication is helpful is mediated by how agents share information and whether they are willing to share arguments critical of their current “location” in the landscape.

Mäs and Flache (2013) develop a less structured model of argumentation. For Mäs and Flache, agents have “reasons” that speak either for or against a target belief. They don’t model a fixed social network, but allow one to evolve over time. They endeavor to determine how homophily might be responsible for polarization. As with Axelrod’s model (see section 2.3), they find homophily is sufficient to produce polarization.[11]

Network formation

[11] (Singer et al. 2019) develop a similar model, but there is no social network structure in their model.

In most of the models considered so far, the network structure is taken as exogenous and fixed. In real human societies, the network structure is dynamic and influenced by the choices of individuals. Models that look at network change and network formation account for this.

Economics has produced one modeling paradigm: each individual receives some payoff from the structure of the network. They create or destroy connections to others and are striving to achieve the highest possible payoff. An equilibrium is reached when no person can improve their situation by adding or deleting a connection (Jackson & Wolinsky, 1996, Jackson, 2008, Bala & Goyal, 2000, Goyal 2007).

Suppose each agent has some piece of information that can be transferred along the social network, perhaps with error as it passes through intermediaries. The information is valuable, but maintaining connections has a cost. These benefits and costs are made precise via a utility function. In this model, the ideal network is either empty, complete, or the star because it balances benefits against costs for the agents collectively (Jackson & Wolinsky, 1996). Sometimes networks of this form are inconsistent with individual incentives, leading to a public goods problem (Jackson, 2008). Even when they are consistent with incentives it can be quite difficult for communities to find the optimal social arrangement because of the large search space (Watts 2001, Huttegger, Skyrms, & Zollman, 2014, Zollman, 2017).

Many models of dynamic network formation explain polarization. They create an underlying process whereby both individuals' beliefs and their social network evolve over time (Kivinen 2017, Fu & Wang, 2008, Macy et al., 2003, Del Vicario et al., 2017, Santos, Lelkes, & Levin, 2021, Mäs & Flache, 2013).[12] A common theme is that polarization can emerge from processes that lead to belief homophily—the tendency of people to interact more with people who have similar beliefs.

Conclusion

Network epistemology is a field rife for valuable interdisciplinary collaboration. Many of the social processes are too difficult to study entirely in the lab or naturalized settings. Instead, models are necessary. The models themselves require substantial empirical input, either testing the mechanisms that they use or

[12] The evolving distrust relationship present in some of the models can be seen as a type of evolving network. O'Connor and Weatherall (2018) show that distrust can lead to polarization in bandit problem models.

their macroscopic predictions. While some of this testing has occurred, there remain many processes where the modeler must use her best guess. Many macroscopic predictions of these models remain untested.

The questions tackled by network epistemology are nonetheless critical. Social epistemic pathologies like polarization, misinformation, groupthink, economics bubbles, and the like are detrimental to our global society. Despite their importance, we remain ignorant about much of the individual-level and social-level contributing causes. Understanding the dynamics of networks is one component we must tackle in order to improve our society's ability form and propagate true beliefs.

Acknowledgements

I am grateful to Kenny Easwaran, Tania Lombrozo, and Neil Van Leeuwen for helpful comments.

References

Abrahamson, E. & Rosenkopf, L. (1997). Social Network Effects on the Extent of Innovation Diffusion: A Computer Simulation. *Organization Science* 8 (3), 289–

309. https://doi.org/10.1287/orsc.8.3.289.

Acemoglu, D. & Ozdaglar, A. (2011). Opinion Dynamics and Learning in Social Networks. *Dynamic Games and Applications* 1 (1), 3–49. https://doi.org/10.1007/s13235-010-0004-1.

Anderson, L.R. & Holt, C.A. (1997). Information Cascades in the Laboratory. *The American Economic Review* 87 (5), 847–62.

Axelrod, R. (1997). The Dissemination of Culture: A Model with Local Convergence and Global Polarization. *Journal of Conflict Resolution* 41 (2), 203–26.

Bala, V. & Goyal, S. (2000). A Noncooperative Model of Network Formation. *Econometrica* 68 (5), 1129–81.

Banerjee, A.V. (1992). A Simple Model of Herd Behavior. *The Quarterly Journal of Economics* 107 (3), 797–817.

Barkoczi, D. & Galesic, M. (2016). Social Learning Strategies Modify the Effect of Network Structure on Group Performance. *Nature Communications* 7 (1), 13109. https://doi.org/10.1038/ncomms13109.

Bartal, A., Pliskin, N., & Tsur, O. (2020). Local/Global Contagion of Viral/Non-Viral Information: Analysis of Contagion Spread in Online Social Networks. *PLOS ONE* 15 (4), e0230811. https://doi.org/10.1371/journal.pone.0230811.

Bass, F.M. (1969). A New Product Growth for Model Consumer Durables. *Management Science* 15 (5), 215–27.

Becker, J., Brackbill, D., & Centola, D. (2017). Network Dynamics of Social Influence in the Wisdom of Crowds. *Proceedings of the National Academy of Sciences of the United States of America* 114 (26), E5070–76. https://doi.org/10.1073/pnas.1615978114.

Bikhchandani, S, Hirshleifer, D., & Welch, I. (1992). A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades. *Journal of Political Economy* 100 (5), 992–1026.

Bolton, P. & Harris, C. (1999). Strategic Experimentation. *Econometrica* 67 (2), 349–74.

Borg, A.M., Frey, D., Šešelja, D, & Straßer, C. (2017). Examining Network Effects in an Argumentative Agent-Based Model of Scientific Inquiry. In *Logic, Rationality, and Interaction*, edited by Alexandru Baltag, Jeremy Seligman, and Tomoyuki Yamada, 391–406. Berlin, Heidelberg: Springer. https://doi.org/10.1007/978-3-662-55665-8_27.

Borg, A.M., Frey, D., Šešelja, D., & Straßer, C. (2018). Epistemic Effects of Scientific Interaction : Approaching the Question with an Argumentative Agent-Based Model. *Historical Social Research* 43 (1), 285–307. https://doi.org/10.12759/hsr.43.2018.1.285-307.

Bramson, A., Grim, P., Singer, D.J., Berger, W.J., Sack, G., Fisher, S., Flocken, C., & Holman, B. (2017). Understanding Polarization: Meanings, Measures, and

Model Evaluation. *Philosophy of Science* 84 (1),115–59. https://doi.org/10.1086/688938.

Brandts, J., Giritligil, A.E., & Weber, R.A.. (2015). An Experimental Study of Persuasion Bias and Social Influence in Networks. *European Economic Review* 80, 214–29. https://doi.org/10.1016/j.euroecorev.2015.07.007.

Campbell, C.M., Izquierdo, E.J. & Goldstone, R.L. (2022). Partial Copying and the Role of Diversity in Social Learning Performance. *Collective Intelligence* 1 (1), 263391372210818. https://doi.org/10.1177/26339137221081849.

Centola, D. (2018). *How Behavior Spreads: The Science of Complex Contagions*. Princeton: Princeton University Press.

Centola, D. & Macy, M.. (2007). Complex Contagions and the Weakness of Long Ties. *American Journal of Sociology* 113 (3), 702–34.

Chandrasekhar, A.G., Larreguy,H. & Xandri, J.P. (2020). Testing Models of Social Learning on Networks: Evidence from Two Experiments. *Econometrica* 88 (1), 1–32.

Charness, G. Cooper, D.J. & Grossman, Z. (2020). Silence Is Golden: Team Problem Solving and Communication Costs. *Experimental Economics* 23 (3), 668–93. https://doi.org/10.1007/s10683-019-09627-w.

Choi, H., Kim, S.H., and Lee, J. (2010). Role of Network Structure and Network Effects in Diffusion of Innovations. *Industrial Marketing Management*, Case

Study Research in Industrial Marketing, 39 (1), 170–77. https://doi.org/10.1016/j.indmarman.2008.08.006.

Choi, S., Gale, D., & Kariv, S. (2005). Behavioral Aspects of Learning in Social Networks: An Experimental Study. In *Experimental and Behavioral Economics*, edited by John Morgan, 25–61. Amsterdam: Elsevier.

Collins, P.J., Hahn, U., von Gerber, Y. & Olsson, E.J. (2018). The Bi-Directional Relationship between Source Characteristics and Message Content. *Frontiers in Psychology* 9. https://www.frontiersin.org/articles/10.3389/fpsyg.2018.00018.

Corazzini, L., Pavesi, F., Petrovich, B. & Stanca, L. (2012). Influential Listeners: An Experiment on Persuasion Bias in Social Networks. *European Economic Review* 56 (6), 1276–88. https://doi.org/10.1016/j.euroecorev.2012.05.005.

Daw, N.D., O'Doherty, J.P. Dayan, P., Seymour, B & Dolan, R.J. (2006). Cortical Substrates for Exploratory Decisions in Humans. *Nature* 441 (7095), 876–79. https://doi.org/10.1038/nature04766.

De Tarlé, L. D, Michelini, M., Šešelja, D., & Staßer, C. S. (2025). Argumentative Agent-Based Models. *Journal of Applied Logics - IfCoLog Journal of Logics and Their Applications*, *12*(3), 489–547.

DeGroot, M.H. (1974). Reaching a Consensus. *Journal of the American Statistical Association* 69 (345), 118–21.

Del Vicario, M., Scala, A., Caldarelli,G., Stanley, H.E., & Quattrociocchi, W. (2017). Modeling Confirmation Bias and Polarization. *Scientific Reports* 7 (January), 1–9. https://doi.org/10.1038/srep40391.

Delre, S.A., Jager, W., Bijmolt, T.H.A., & Janssen. M.A. (2010). Will It Spread or Not? The Effects of Social Influences and Network Topology on Innovation Diffusion. *Journal of Product Innovation Management* 27 (2), 267–82. https://doi.org/10.1111/j.1540-5885.2010.00714.x.

Delre, S.A., Jager, W. & Janssen. M.A. (2007). Diffusion Dynamics in Small-World Networks with Heterogeneous Consumers. *Computational and Mathematical Organization Theory* 13 (2), 185–202. https://doi.org/10.1007/s10588-006-9007-2.

Demarzo, P.M., Vayanos,D. & Jeffrey Zwiebel. (2003). "Persuasion Bias, Social Influence, and Uni-Dimensional Opinions." *Quarterly Journal of Economics* 118 (3), 909–68.

Derex, M., & Boyd, R. (2016). Partial Connectivity Increases Cultural Accumulation within Groups. *Proceedings of the National Academy of Sciences* 113 (11), 2982–87. https://doi.org/10.1073/pnas.1518798113.

Douven, I., & Hegselmann, R. (2021). Mis- and Disinformation in a Bounded Confidence Model. *Artificial Intelligence* 291 (February),103415. https://doi.org/10.1016/j.artint.2020.103415.

Dung, P.M. (1995). On the Acceptability of Arguments and Its Fundamental Role in Nonmonotonic Reasoning, Logic Programming and n-Person Games. *Artificial Intelligence* 77 (2), 321–57. https://doi.org/10.1016/0004-3702(94)00041-X.

Dunn, A.G., & Gallego. B. (2010). Diffusion of Competing Innovations: The Effects of Network Structure on the Provision of Healthcare. *Journal of Artificial Societies and Social Simulation* 13 (4), 8.

Fazelpour, S. & Steel, D. (2022). Diversity, Trust and Conformity. *Philosophy of Science* 89 (2), 209–31.

Fink, C., Schmidt, A., Barash, V. Cameron, C. & Macy, M. (2015). Complex Contagions and the Diffusion of Popular Twitter Hashtags in Nigeria. *Social Network Analysis and Mining* 6 (1), 1. https://doi.org/10.1007/s13278-015-0311-z.

Fink, Clay, Schmidt, A., Barash, V., Kelly, J., Cameron, C., & Macy, M. (2016). Investigating the Observability of Complex Contagion in Empirical Social Networks. *Proceedings of the International AAAI Conference on Web and Social Media* 10 (1), 121–30. https://doi.org/10.1609/icwsm.v10i1.14751.

French, J.R. 1956. A Formal Theory of Social Power. *Psychological Review* 63 (3), 181–94.

Fu, F. & Wang, L. (2008). Coevolutionary Dynamics of Opinions and Networks: From

Diversity to Uniformity. *Physical Review E* 78 (1), 016104. https://doi.org/10.1103/PhysRevE.78.016104.

Gabriel, N. & O'Connor, C. (2024). Can Confirmation Bias Improve Group Learning? *Philosophy of Science* 91 (2), 329–50. https://doi.org/10.1017/psa.2023.176.

Goel, S., Anderson, A., Hofman, J & Watts, D.J. (2016). The Structural Virality of Online Diffusion. *Management Science* 62 (1), 180–96. https://doi.org/10.1287/mnsc.2015.2158.

Goeree, J.K., Palfrey, T.P., Rogers,B.W. & McKelvey R.D. (2007). Self-Correcting Information Cascades. *Review of Economic Studies* 74 (3), 733–62. https://doi.org/10.1111/j.1467-937X.2007.00438.x.

Goldenberg, J. & Efroni. S. (2001). Using Cellular Automata Modeling of the Emergence of Innovations. *Technological Forecasting and Social Change* 68 (3), 293–308. https://doi.org/10.1016/S0040-1625(00)00095-0.

Goldstone, R.L., Wisdom, T.N. Roberts, M.E. & Frey, S. (2013). Learning Along With Others. In *Psychology of Learning and Motivation*, 58:1–45. Elsevier. https://doi.org/10.1016/B978-0-12-407237-4.00001-3.

Golub, B. & Jackson. M.O. (2010). Naïve Learning in Social Networks and the Wisdom of Crowds. *American Economic Journal: Microeconomics* 2 (1), 112–49.

Golub, B., & Sadler. E. (2016). Learning in Social Networks. In *The Oxford Handbook of The Economics of Networks*, edited by Yann Bramoullé, Andrea Galeotti, and Brian W. Rogers. Oxford: Oxford University Press.

Goyal, S. (2007). *Connections: An Introduction to the Economics of Networks*. Princeton: Princeton University Press.

Granovskiy, B., Gold, J.M., Sumpter, D.J.T. & Goldstone, R. (2015). Integration of Social Information by Human Groups. *Topics in Cognitive Science* 7 (3), 469–93. https://doi.org/10.1111/tops.12150.

Grim, P., Singer, D.J., Bramson, A., Holman, B., McGeehan, S. & Berger, W.J. (2019). Diversity, Ability, and Expertise in Epistemic Communities. *Philosophy of Science* 86 (1), 98–123. https://doi.org/10.1086/701070.

Grim, P., Singer, D.J., Fisher, S.,Bramson, A., Berger, W.J., Reade, C., Flocken, C. & Sales, A. (2013). Scientific Networks on Data Landscapes: Question Difficulty, Epistemic Success, and Convergence. *Episteme* 10 (4), 441–64. https://doi.org/10.1017/epi.2013.36.Scientific.

Grimm, V. & Mengel, F. (2020). Experiments on Belief Formation in Networks. *Journal of the European Economic Association* 18 (1), 49–82. https://doi.org/10.1093/jeea/jvy038.

Hahn, U., Hansen, J.U. & Olsson, E.J. (2020). *Truth Tracking Performance of Social Networks: How Connectivity and Clustering Can Make Groups Less Competent*. *Synthese*. Vol. 197. Springer Netherlands. https://doi.org/10.1007/s11229-018-01936-6.

Hązła, J., Jadbabaie, A. Mossel, E. & Rahimian, M.A. (2021). Bayesian Decision Making in Groups Is Hard. *Operations Research* 69 (2), 632–54. https://doi.org/10.1287/opre.2020.2000.

Hegselmann, R. (2023). Bounded Confidence Revisited: What We Overlooked, Underestimated, and Got Wrong. *Journal of Artificial Societies and Social Simulation* 26 (4), 11.

Hegselmann, R. & Krause, U. (2002). Opinion Dynamics and Bounded Confidence. *Journal of Artificial Societies and Social Simulation* 5 (3). http://jass.soc.survey.ac.uk/5/3/2.html.

———. (2006). Truth and Cognitive Division of Labor: First Steps Toward a Computer Aided Social Epistemology. *Journal of Artificial Societies and Social Simulation* 9 (3), 1–29.

Hethcote, H.W. 1978. An Immunization Model for a Heterogeneous Population. *Theoretical Population Biology* 14 (3), 338–49. https://doi.org/10.1016/0040-5809(78)90011-4.

Hodas, N.O., & Lerman, K. (2014). The Simple Rules of Social Contagion. *Scientific Reports* 4 (1), 4343. https://doi.org/10.1038/srep04343.

Hohnisch, M., Pittnauer, S. & Stauffer, D. (2008). A Percolation-Based Model Explaining Delayed Takeoff in New-Product Diffusion. *Industrial and Corporate Change* 17 (5), 1001–17. https://doi.org/10.1093/icc/dtn031.

Hong, L. & Page, S.E. (2004). Groups of Diverse Problem Solvers Can Outperform Groups of High-Ability Problem Solvers. *Proceedings of the National Academy of Sciences* 101 (46), 16385–89. https://doi.org/10.1073/pnas.0403723101.

Hung, A.A. & Plott, C.R.. (2001). Information Cascades: Replication and an Extension to Majority Rule and Conformity-Rewarding Institutions. *American Economic Review* 91 (5), 1508–20. https://doi.org/10.1257/aer.91.5.1508.

Huttegger, S.M., Skyrms, B. & Zollman, K.J.S.. (2014). Probe and Adjust in Information Transfer Games. *Erkenntnis* 79:835–53.

Jackson, M.O. (2008). *Social and Economic Networks*. Princeton: Princeton University Press.

Jackson, M.O. & Wolinsky. A. (1996). A Strategic Model of Social and Economic Networks. *Journal of Economic Theory* 74 (0108), 44–74.

Jönsson, M.L., Hahn, U. & Olsson. E.J. (2015). The Kind of Group You Want to Belong to: Effects of Group Structure on Group Accuracy. *Cognition* 142:191–204. https://doi.org/10.1016/j.cognition.2015.04.013.

Kivinen, St. (2017). Polarization in Strategic Networks. *Economics Letters* 154:81–83. https://doi.org/10.1016/j.econlet.2017.02.032.

Kuandykov, L. & Sokolov, M. (2010). Impact of Social Neighborhood on Diffusion of Innovation S-Curve. *Decision Support Systems* 48 (4), 531–35. https://doi.org/10.1016/j.dss.2009.11.003.

Kübler, D. & Weizsäcker, G. (2004). Limited Depth of Reasoning and Failure of Cascade Formation in the Laboratory. *Review of Economic Studies* 71 (2), 425–41. https://doi.org/10.1111/0034-6527.00290.

———. (2005). Are Longer Cascades More Stable? *Journal of the European Economic Association* 3 (2/3), 330–39.

Kummerfeld, E. & Zollman, K.J.S. (2016). Conservatism and the Scientific State of Nature. *British Journal for the Philosophy of Science* 67 (4), 1057–76.

LaCroix, T., Geil, A. & O'Connor, C. (2021). The Dynamics of Retraction in Epistemic Networks. *Philosophy of Science* 88 (3), 415–38. https://doi.org/10.1086/712817.

Lazer, D. & Friedman, A. (2007). The Network Structure of Exploration and Exploitation. *Administrative Science Quarterly* 52 (4), 667–94. https://doi.org/10.2189/asqu.52.4.667.

Lee, M.D., Zhang, S., Munro, M. & Steyvers, M. (2011). Psychological Models of Human and Optimal Performance in Bandit Problems. *Cognitive Systems Research*, The 9th International Conference on Cognitive Modeling. Manchester, UK, July 2009, 12 (2), 164–74. https://doi.org/10.1016/j.cogsys.2010.07.007.

Lehrer, K. & Wagner, C. (1981). *Rational Consensus in Science and Society: A Philosophical and Mathematical Study*. Dordrecht: D. Reidel.

Lerman, K. & Ghosh, R. (2010). Information Contagion: An Empirical Study of the Spread of News on Digg and Twitter Social Networks. *Proceedings of the International AAAI Conference on Web and Social Media* 4 (1), 90–97. https://doi.org/10.1609/icwsm.v4i1.14021.

Macy, M.W., Kitts, J.A., Flache, A. & Benard, S. (2003). Polarization in Dynamic Networks: A Hopfield Model of Emergent Structure. In *Dynamic Social Network Modeling and Analysis*, edited by R Breiger, K Carley, and P Pattison, 162–73. Washington, DC: The National Academies Press.

Mäs, M., & Flache, A. (2013). Differentiation without Distancing. Explaining Bi-Polarization of Opinions without Negative Influence. *PloS One* 8 (11), e74516. https://doi.org/10.1371/journal.pone.0074516.

Mason, W.A. Jones, A. & Goldstone, R.L. (2008). Propagation of Innovations in Networked Groups. *Journal of Experimental Psychology: General* 137 (3), 422–33. https://doi.org/10.1037/a0012798.

Mason, W.A. & Watts, D.J. (2012). Collaborative Learning in Networks. *Proceedings of the National Academy of Sciences of the United States of America* 109 (3), 764–69. https://doi.org/10.1073/pnas.1110069108.

Meade, N. & Islam, T. (2006). Modelling and Forecasting the Diffusion of Innovation – A 25-Year Review. *International Journal of Forecasting* 22 (3), 519–45. https://doi.org/10.1016/j.ijforecast.2006.01.005.

Mønsted, B., Sapieżyński, P., Ferrara., E., & Lehmann, S. (2017). Evidence of Complex Contagion of Information in Social Media: An Experiment Using Twitter Bots. *PLOS ONE* 12 (9), e0184148. https://doi.org/10.1371/journal.pone.0184148.

Morreau, M. & Olsson, E.J. (2022). Learning from Ranters: The Effect of Information Resistance on the Epistemic Quality of Social Network Deliberation. In *Social Virtue Epistemology*, edited by Mark Alfano, Colin Klein, & Jeroen de Ridder. New York: Routledge.

Myers, S.A., Zhu, C. & Leskovec. J. (2012). Information Diffusion and External Influence in Networks. In *Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 33–41. KDD ’12. New York, NY, USA: Association for Computing Machinery. https://doi.org/10.1145/2339530.2339540.

Newman, M.E.J. (2002). Spread of Epidemic Disease on Networks. *Physical Review E* 66 (1), 016128. https://doi.org/10.1103/PhysRevE.66.016128.

———. (2018). *Networks*. Second Edition. Oxford: Oxford University Press.

O’Connor, C. & Weatherall, J.O.. (2018). Scientific Polarization. *European Journal for Philosophy of Science* 8 (3), 855–75. https://doi.org/10.1007/s13194-018-0213-9.

Olsson, E.J. (2011). A Simulation Approach to Veritistic Social Epistemology.

Episteme 8 (2), 127–43. https://doi.org/10.3366/epi.2011.0012.

———. (2013). A Bayesian Simulation Model of Group Deliberation and Polarization. *Bayesian Argumentation: The Practical Side of Probability*, 113–33. https://doi.org/10.1007/978-94-007-5357-0_6.

———. (2020). Why Bayesian Agents Polarize. In *The Epistemology of Group Disagreement*, edited by F. Croncano-Berrocal & A. Carter. Routledge.

Pöyhönen, S. (2017). Value of Cognitive Diversity in Science. *Synthese* 194 (11), 4519–40. https://doi.org/10.1007/s11229-016-1147-4.

Rogers, E.M. (2003). *Diffusion of Innovations*. Fifth Edition. New York: Free Press.

Romero, D.M., Meeder, B., & Kleinberg, J. (2011). Differences in the Mechanics of Information Diffusion across Topics: Idioms, Political Hashtags, and Complex Contagion on Twitter. In *Proceedings of the 20th International Conference on World Wide Web*, 695–704. WWW ’11. New York, NY, USA: Association for Computing Machinery. https://doi.org/10.1145/1963405.1963503.

Rosenstock, S., Bruner, J.P., & O'Connor, C. (2017). In Epistemic Networks, Is Less Really More? *Philosophy of Science* 84 (2), 234–52. https://doi.org/10.1086/690717.

Sadler, E. (2014). Bounded Rationality and the Network Economist’s Dilemma. SSRN Scholarly Paper. Rochester, NY. https://doi.org/10.2139/ssrn.2501690.

Santos, F., Lelkes, Y. & Levin, S.A.. (2021). Link Recommendation Algorithms and Dynamics of Polarization in Online Social Networks. *Proceedings of the National Academy of Sciences* 118 (50), e2102141118. https://doi.org/10.1073/pnas.2102141118.

Singer, D.J., Bramson, A., Grim, P., Holman, B., Jung, J., Kovaka, K., Ranginani, A., & Berger, W.J.. (2019). Rational Social and Political Polarization. *Philosophical Studies* 176 (9), 2243–67. https://doi.org/10.1007/s11098-018-1124-5.

Smith, L. & Sørensen, P. (2000). Pathological Outcomes of Observational Learning. *Econometrica* 68 (2), 371–98. https://doi.org/10.1111/1468-0262.00113.

Speekenbrink, M. & Konstantinidis, E. (2014). Uncertainty and Exploration in a Restless Bandit Task. *Proceedings of the Annual Meeting of the Cognitive Science Society* 36:7.

Steyvers, M., Lee, M.D., & Wagenmakers. E-J. (2009). A Bayesian Analysis of Human Decision-Making on Bandit Problems. *Journal of Mathematical Psychology*, Special Issue: Dynamic Decision Making, 53 (3), 168–79. https://doi.org/10.1016/j.jmp.2008.11.002.

Takayasu, M., Sato, K., Sano, Y., Yamada, K., Miura, W., & Takayasu, H. (2015). Rumor Diffusion and Convergence during the 3.11 Earthquake: A Twitter Case Study. *PLOS ONE* 10 (4), e0121443. https://doi.org/10.1371/journal.pone.0121443.

Thoma, J. (2015). The Epistemic Division of Labor Revisited. *Philosophy of Science* 82 (3), 454–72. https://doi.org/10.1086/681768.

Vallinder, A. & Olsson, E.J.. (2014). Trust and the Value of Overconfidence: A Bayesian Perspective on Social Network Communication. *Synthese* 191 (9), 1991–2007. https://doi.org/10.1007/s11229-013-0375-0.

Watts, A. (2001). A Dynamic Model of Network Formation. *Games and Economic Behavior* 34 (2), 331–41. https://doi.org/10.1006/game.2000.0803.

Weatherall, J.O. & O'Connor, C. (2020). Endogenous Epistemic Factionalization. *Synthese*, 1–23. https://doi.org/10.1007/s11229-020-02675-3.

Wisdom, T.N. & Goldstone, R.L. (2007). Group Problem Solving Behavior in a Networked Puzzle Game. In *Proceedings of CogSci 2007*, 1665–70.

Wisdom, T,N., Song, X., & Goldstone, R.L. (2004). The Effects of Peer Information on Problem-Solving in a Networked Group.

Wu, J. (2023a). Better than Best: Epistemic Landscapes and Diversity of Practice in Science.

———. (2023b). Epistemic Advantage on the Margin: A Network Standpoint Epistemology. *Philosophy and Phenomenological Research* 106 (3), 755–77. https://doi.org/10.1111/phpr.12895.

Wu, J. & O'Connor, C. (2023). How Should We Promote Transient Diversity in Science? *Synthese* 201 (2), 37. https://doi.org/10.1007/s11229-023-04037-1.

Young, H.P. (2009). Innovation Diffusion in Heterogeneous Populations: Contagion, Social Influence, and Social Learning. *American Economic Review* 99 (5), 1899–1924. https://doi.org/10.1257/aer.99.5.1899.

Zehmakan, A.N. (2020). Opinion Forming in Erdős–Rényi Random Graph and Expanders. *Discrete Applied Mathematics* 277:280–90. https://doi.org/10.1016/j.dam.2019.10.001.

Zhang, S. & Yu, A.J. (2013a). Cheap but Clever: Human Active Learning in a Bandit Setting. *Proceedings of the Annual Meeting of the Cognitive Science Society* 35:7.

———. (2013b). Forgetful Bayes and Myopic Planning: Human Learning and Decision-Making in a Bandit Setting. *NIPS*, 9.

Ziegelmeyer, A., Koessler, F., Bracht, J., & Winter, E. (2010). Fragility of Information Cascades: An Experimental Study Using Elicited Beliefs. *Experimental Economics* 13 (2), 121–45. https://doi.org/10.1007/s10683-009-9232-x.

Zollman, K.J.S. (2007). The Communication Structure of Epistemic Communities. *Philosophy of Science* 74 (5), 574–87. https://doi.org/10.1086/525605.

———. (2010a). Social Structure and the Effects of Conformity. *Synthese* 172 (3), 317–40. https://doi.org/10.1007/s11229-008-9393-8.

———. (2010b). The Epistemic Benefit of Transient Diversity. *Erkenntnis* 72 (1), 17–35.

———. (2011). Social Network Structure and the Achievement of Consensus.

Politics, Philosophy & Economics 11 (1), 26–44. https://doi.org/10.1177/1470594X11416766.

———. (2017). Learning to Collaborate. In *Scientific Collaboration and Collective Knowledge*, edited by Thomas Boyer-Kassem, Conor Mayo-Wilson, and Michael Weisberg. New York: Oxford University Press. https://doi.org/10.1093/oso/9780190680534.003.0003.