False Beliefs and the Social Structure of Science: Some Models and Case Studies

Cailin O'Connor and James Owen Weatherall

March 2019

We use models and historical cases to try and understand some of the ways scientific beliefs can go wrong. In particular, we consider questions like: how do conformity and social trust influence the spread of beliefs? What is the ideal network structure for a scientific community? And how do industrial propagandists influence the progress of science, as well as public belief?

1 Introduction

Humans are inherently social animals. It should come as no surprise that our systems of knowledge formation are social as well. We pass information and evidence from person to person via testimony. As a result, for each of us, the majority of things we believe were learned from other people. This is an enormously powerful and successful ability, in general. Human culture and technological innovation would not be possible without the social spread of knowledge. But the social spread of knowledge opens us up to a potential problem—when we trust the testimony of others, we will sometimes trust false things as well as true ones.[1]

In this paper, we consider how the social spread of knowledge happens, especially in scientific communities and regarding scientific beliefs. In particular, we will use a set of agent-based models, drawn from what is called the *network epistemology* framework, to illustrate how the social aspects of science can influence theory adoption and lead to potential problems. Section 2 introduces this framework, and makes clear how it can be used to represent the spread of scientific beliefs. Section 3 is the heart of the paper, where we discuss a number of historical cases in science where 1) social factors influenced outcomes and 2) network epistemology models can help us understand how. As will become clear, these models can be applied both to cases where aspects of social psychology influence belief spread, and to cases where pernicious influencers or propagandists attempt to shape scientific belief. Section 4 concludes by briefly discussing

[1] For a nice discussion of this tradeoff, see (Mayo-Wilson, 2014).

some social and political issues that arise as a result of the persistence of false scientific beliefs.

2 Network Epistemology Models of Science

Mathematical models are commonly used in the social sciences as a tool to understand human interaction. These models are especially useful in cases where it is difficult to fully investigate a group of people empirically. In the case of information spread and theory change, scientific communities often adopt a new theory slowly over the course of many years, and as a result of thousands of interactions between scientists and dozens of experiments. Although we cannot directly observe each of these interactions, we can use a combination of empirical observation and modeling work to develop an understanding of how these processes occur.

The network epistemology framework was first introduced by economists Bala and Goyal (1998) to model the spread of social knowledge.[2] *Epistemology* refers to the fact that these models are about knowledge and belief formation. More recently, it has been used to directly model scientific communities, starting with the work of the philosopher of science Kevin Zollman (Zollman, 2007, 2010b). Since then, this framework has been used widely to model scientific consensus and the spread of scientific belief.[3]

How does this sort of model work? Here we will give a relatively non-technical description, as is appropriate for an interdisciplinary book. There are two features of these models—a *decision problem* and a *network*. Let us first discuss the decision problem. In all of the models we consider, agents face a *two-armed bandit problem*. A 'bandit' is another name for a slot machine. The idea is that they face a problem analogous to choosing one of two arms on a slot machine, where one arm pays out more often than the other. We use this problem to represent situations in science where agents are considering two possibilities—say, two theories, or two medical treatments—that may be more or less successful, and attempting to choose the best one. In what follows, we will refer to the two possible actions/theories as 'Arm A' and 'Arm B', in keeping with the two-armed bandit metaphor; in general, Arm B will refer to the more successful of the two actions (B is for 'better').

How do agents choose? These models assume that agents have a belief about which arm is best, that they can gather evidence by actually testing the world, and that their belief is sensitive to the evidence gathered. Belief is represented

[2]There is another, highly influential, framework for modeling the spread of belief that treats beliefs like viruses in a contagion (Rogers, 2010). We focus on network epistemology as a better modeling framework for science because agents can gather and share evidence from the world to support belief.

[3]For more on the use of this framework in philosophy of science see Zollman (2010b); Mayo-Wilson et al. (2011); Zollman (2013); Kummerfeld and Zollman (2015); Holman and Bruner (2015); Rosenstock et al. (2017); Borg et al. (2017); Frey and Šešelja (2017a,b); Holman and Bruner (2017); Weatherall and O'Connor (2017); Weatherall et al. (2018); O'Connor and Weatherall (2018); Weatherall and O'Connor (2018); O'Connor and Weatherall (2019).

by a *credence*—this is a number between 0 and 1 reflecting a degree of belief in something. For example, suppose that Karen thinks there is a 60% chance that arm B is better than arm A. We can say she has credence .6 that arm B is best. As mentioned, credences are sensitive to the evidence that agents gather from the world. Typically these models start with agents who have random credences about whether A or B is better. In each round, agents decide to make a test. They can either pull arm A or arm B. A common assumption in these models is that their beliefs guide their choice—if a scientist thinks B is the better theory, she is more likely to test it. Then, based on the evidence they gather, they change their credence. Suppose an agent pulls arm B, and it is very successful. That agent's credence that B is better should go up. If arm B fails, this credence goes down. Many of these models assume agents change their beliefs using some version of *Bayes rule*, which probability theory tells us is the rationally best way to change beliefs in light of evidence. From model to model there are variations on exactly how beliefs are represented, evidence is gathered, and beliefs are updated.

Now let us turn to the network aspect of the model. Agents are part of a network where each node is an individual, and each edge is a social connection. Agents in this sort of model share the evidence they gather with their neighbors in the network. In this way, beliefs can spread throughout a network. How does this work? Suppose a community is settled on some theory, represented by arm A. This could be 'cigarette smoking is safe'. Now one scientist develops a different theory, represented by arm B, 'cigarette smoking causes cancer'. This new belief leads the scientist to change behaviors. Now that they think cigarettes might be dangerous, they begin to gather data about whether this is indeed the case. If their evidence is persuasive, as they share it with their peers, these peers will also become convinced and begin to change their own behaviors to reflect the new belief. Over time the entire network may switch to the new, successful theory. For this reason, a common endpoint of these models is that all agents have settled on a correct consensus.

This is not what always happens, though. Social influence can also lead the entire group to settle on incorrect beliefs. For instance, a few pieces of spurious evidence in favor of a false belief, if shared widely, may be enough to convince an entire community of something incorrect. In the next section, we will discuss this possibility at more length. And we will discuss what sorts of conditions make false consensus more or less likely in a scientific community.

3 Cases and Results

The models described in the last section have been widely used to study the spread of scientific beliefs. They have been used to ask: what is the best communication structure for science? Is there true wisdom of the crowds? Is experimentation good for scientific communities? How does cognitive diversity benefit science? In this section, we will briefly go through some variations on the model, connecting them with real cases from science.

3.1 Ulcers and the Zollman Effect

In the early 1900s, scientists were divided on the topic of stomach ulcers. Some thought that ulcers were caused by bacteria, and others that the cause was stomach acid. There was evidence supporting both theories. But in 1954 the gasteroenterologist E.D. Palmer, upon biopsying the stomachs of over 1000 patients, found no evidence of bacterial life at all (Palmer, 1954). Of course, stomachs do, in fact, have bacteria in them. Palmer's findings were misleading. But they were so influential, that an entire generation of scientists turned away from the bacterial theory of ulcers and focused on treatments for stomach acid. It wasn't until 30 years later that researchers Robin Warren and Barry Marshall revived the bacterial theory. They did so in dramatic fashion. After encountering skepticism from the scientific community, Marshall downed a petri dish of H. Pylori, the bacteria that causes stomach ulcers, developed a case of ulcers, and cured himself using antibiotics (Kidd and Modlin, 1998). In 2005, their worked earned them a Nobel prize.

Zollman (2007, 2010b) made a curious discovery about epistemic network models. One might think that in general communication among agents is a good thing. Each agent is gathering real evidence by testing the bandit arms. If each agent receives more of this good evidence from neighbors, surely they are better able to draw accurate conclusions. As it turns out, though, this is not always the case.

Scientific evidence is probabilistic, which means that sometimes it is misleading. For instance, many individual studies on the dangers of smoking find no link between tobacco and cancer. This is because not everybody who smokes gets cancer, and some people get cancer who do not smoke. If we do enough studies some, by chance, will fail to detect the increased risk that smoking causes. Likewise, agents pulling arms on a bandit will sometimes get spurious results. Suppose arm A is successful 50% of the time, and arm B 60% of the time. An agent who tests arm B 10 times may find that it is successful in only 4 of these tests. This data makes B looks worse than A. In these models, sometimes a single set of misleading results of this sort will be enough to convince an entire network, or a large portion of it, of a false belief. In other words, too much connectivity can be a problem.

Here is another way to put the issue: some transient disagreement in beliefs is generally a good thing for a scientific community. Without diversity of belief, a community might fail to ever investigate a promising theory. A community that is too unified, and communicates too much, may fail to maintain enough diversity of opinion to settle on the best option.

Zollman finds, in particular, that well connected groups of researchers tend to come to consensus quicker in network epistemology models. But sparsely connected networks are more likely to settle on a true consensus, rather than a false one. This has since been called the Zollman effect.[4] Something like

[4]Though, as Zollman (2013) and Rosenstock et al. (2017) point out, this finding is sensitive to parameters in the model. For many parameters, there is no such effect (though it never seems to reverse). See also Frey and Šešelja (2017a) for critiques of the robustness of these

this is arguably what happened in the H. pylori case. Palmer's findings were too influential, and they led an entire community to prematurely settle on a false consensus. If a sub-community had continued to test the bacterial theory, they may have managed to convince the larger community of the truth with the evidence they gathered.

3.2 Smallpox and Conformity

Lady Mary Whortley Montague was a British aristocrat born in the late 17th century.[5] As a young woman, she suffered a bout of smallpox. While she survived, she lost a brother to the disease and she was permanently scarred. A few years later, Montague traveled to Turkey with her husband. While there, she encountered the practice of smallpox variolation. Similar to vaccination, this involved exposing patients to pus from smallpox pustules via a scratch on the arm. The subsequent infection tended to be mild, and to prevent further, more virulent infection. With a Turkish nurse and an English physician named Charles Maitland, she had her own son variolated before returning to England.

Upon her return in 1721, Lady Montague attempted to spread the practice, but met resistance from English physicians. They were not particularly interested in a practice performed by Turkish women and advocated for by an English woman. Even Charles Maitland was unwilling to perform variolations once under the eye of his English peers.

Montague managed to get around their resistance in a particularly clever way. She was friends with Princess Caroline of Ansbach, who was married to the Crown Prince of England. Montague convinced the princess to publicly variolate her own two daughters. After this, the practice spread quickly among the nobility, especially those with personal connections to the princess and Lady Montague.

In O'Connor and Weatherall (2019) we argue that the spread of variolation in this case ultimately had relatively little to do with evidence and belief. Instead, it was a largely social phenomenon. Experimental psychologists, starting with Asch and Guetzkow (1951), have shown that humans have tendencies towards social conformity. We do not like to stick out from the crowd, and this leads to behaviors like publicly avowing a belief, even if we have good reason not to hold it (Bond and Smith, 1996). In the case of smallpox variolation, English physicians were likely influenced by conformist tendencies in assessing whether variolation might be a good practice. This is especially notable in the case of Maitland who was perfectly happy to perform a variolation in Turkey, but was hesitant to do so in England.

This tendency hurt the belief state of English physicians, but in this particular case conformity bias also ended up playing a positive role. Lady Montague made use of the desires of English nobility to conform to the practices of the most respected and influential members of their society. This desire helped convince many people to variolate once the princess did so.

models.

[5]This history is drawn from Grundy et al. (1999).

In Weatherall and O'Connor (2017) we consider what happens in epistemic network models when actors have a tendency to conform. We do this by supposing that they weigh two desires. The first is a desire to take the action best supported by their beliefs. The second is a desire to conform, by testing the theory that their network neighbors also test. Actors balance these desires in deciding which arm to pull.

In these models, the tendency towards conformity means that communities no longer necessarily converge to a consensus. Instead, sometimes there are cliques that hold stable, opposing beliefs. Imagine if one clique settles on one action (variolating) and another on a different action (not variolating). Some members of the non-variolating clique will learn the truth via evidence that comes from the other clique. But their desire to conform means that they never test the better action themselves, and thus they fail to spread evidence of its success to their network neighbors.

In general, the greater the tendency of actors in these models to conform, the worse the beliefs and practices of the agents. Conformity stops the natural spread of good practices to new parts of a social network. There are conditions that make these effects more or less serious, though. In particular, the models predict that conformity will be less of a problem when there is a big difference in the success of the theories under consideration. For instance, even strong social pressures will not protect a belief like 'it is safe to drink cyanide' from disconfirmation. The negative consequences will outweigh the desire for social conformity. But we might expect conformity to play a big role in the spread of something like evolutionary theory. For most people, there are few consequences one way or another to believing in evolution. Thus social pressures tend to determine what they espouse.[6]

3.3 Lyme Disease and Polarization

Polarization is a popular political buzzword. It usually refers to situations where groups of individuals fail to achieve consensus, or even move further apart in belief/opinion over the course of interaction. Typical cases of polarization happen along political lines, and involve differences in social and moral values. Consider, for instance, debates in the US about gun control or abortion. In some cases, though, polarization happens over matters of scientific fact, and among groups of individuals who share values.

In the mid-1970s, rheumatologist Allen Steere identified Lyme disease as a new tick-borne illness. The symptoms of Lyme disease are many and varied, but typically involve joint pain, nerve pain, headaches, fatigue, and brain fog. Because Lyme is caused by a spirochete, it can be treated by antibiotics. This

[6]Others have used different versions of network models to consider the role of conformity in belief. Zollman (2010a) points out that in cases where agents do not have better ways to combine data, conformity can play a beneficial role. Mohseni and Williams (2017) look at an epistemic network model where agents have expectations about how conformity influences their peers, and these expectations can change their social trust. They find that conformity bias hurts the ability of the network to develop accurate beliefs.

discovery thus radically improved the lives of thousands of sufferers whose symptoms were reduced or eliminated by antibiotic treatment. Despite this apparent success, however, by the late 90s Steere was receiving death threats from Lyme patients.

In the early 90s, Steere became concerned that Lyme was being treated as a catch-all diagnosis for anyone suffering from pain and fatigue. He worried that the long-term antibiotics these sufferers were prescribed were causing unnecessary harm, and began to advocate for stricter standards in Lyme treatment. This was the beginning of what is now called the 'Lyme wars'. At the heart of this debate is the question of whether Lyme is always cured by a short dose of antibiotics, or whether it sometimes persists in a chronic form that requires long-term antibiotic treatment.

What is striking about this debate is that the physicians involved seem to share goals and values. They want to learn the truth, and they want to successfully treat the Lyme patients they see. Nonetheless, the two sides of the 'war' are highly polarized. There is an enormous amount of mistrust. The physicians and patient groups who believe in chronic Lyme accuse the establishment camp of being influenced by money from insurers who do not want to pay for long-term treatment. The establishment physicians accuse doctors who treat chronic Lyme of taking financial advantage of vulnerable patients, and have often attempted to revoke their medical licenses.

In O'Connor and Weatherall (2018) we use epistemic network models to ask: how might scientific community might end up in such a polarized state? And might this happen without the influence of money or bad actors? There is a long literature using models to explore polarization, but these models are usually not well tuned to scientific communities. In previous polarization models, actors' opinions are typically determined solely on the basis of social influence, rather than evidence from the world.[7] And they usually change opinions in nonrational ways. Our goal is to consider agents who 1) collect evidence, 2) use this evidence to shape beliefs, and 3) share this evidence within their communities, but who nonetheless end up polarizing.

The key modification we make to the standard epistemic network model is to add a component of social trust. In particular, we assume that agents treat the evidence they receive as uncertain. If they see some set of data, they think there is some chance that this data obtained, and some chance it did not. Furthermore, their level of uncertainty is determined by how close their beliefs are to a peer in the network. So if two agents have very similar beliefs (both support chronic Lyme, say), they trust the evidence shared by the other. If two agents have disparate beliefs (one believes in chronic Lyme and one does not) they treat each others' evidence as uncertain. There is something reasonable about this—scientists should not treat all data as totally trustworthy given the presence of quacks in scientific communities.

What we find is that this uncertainty can lead to stable polarization. When

[7] We do not review this literature here for space reasons. See O'Connor and Weatherall (2018) and Bramson et al. (2017) for reviews.

this happens, two groups form. One group has good beliefs, and takes the more successful action. The other group has worse beliefs, and takes the less successful action. But because this second group does not trust evidence coming from the first group, they never learn about the better action.

It is interesting to note that polarization can appear in the conformity models discussed above, as well. In that case, polarization arises when networks exhibit a certain 'clique' structure, where there exist tightly knit groups that are only weakly connected to one another. In such cases, members of different cliques come to conform only with members of their own group, preventing outside information from flowing into the clique. This mechanisms is importantly different from the one that arises in the social trust models now under consideration. Consider how one might eliminate polarization in each case. If polarization arises from clique structure and conformity, increasing the degree of social connectedness of a community should lead to less polarization, because now information may flow more freely. But if social trust is responsible for polarization, new social connections will make no difference—unless agents have some reason to trust their new neighbors.

Across the polarization models we explore we find that this sort of mistrust makes the community as a whole worse at forming beliefs. In particular, many more individuals end up with bad beliefs, because they ignore the best data available to them.[8]

3.4 Industrial Selection

In the last two parts of this section, we will describe models that consider how outside forces, especially those from industry, can influence scientific beliefs.

In the late 1970s, medical researchers began to explore the arrhythmic suppression hypothesis, which states that because heart arrhythmias often precede a heart attack, suppressing arrhythmia might work to prevent them.[9] Bernard Lown, who first proposed this hypothesis, pointed out that it was not clear whether arrhythmia suppression would have the desired effect. He advocated for testing anti-arrhythmic medications by looking at whether they reduced heart attack deaths. Other researchers, though, tested the efficacy of these drugs by looking simply at whether they were successful at reducing arrhythmia. Pharmaceutical companies funded researchers only in the latter camp, and the results of this research led to the widespread prescription of anti-arrhythmics.

The problem was that these drugs in fact increased heart attack deaths, rather than preventing them. It wasn't until the late 80s that the large scale Cardiac Arrhythmic Suppression Trial showed this conclusively. It has been estimated that upwards of one hundred thousand deaths may have been caused

[8]In Weatherall and O'Connor (2018) we consider similar models but where actors consider multiple arenas of belief, and ground trust in all of these. For instance, in deciding whether to trust evidence about the safety of vaccines, an individual might compare beliefs about vaccines, but also about the safety of GMO crops in deciding who to trust. We find that groups holding multiple, polarized beliefs can endogenously emerge in these models.

[9]For the history of this episode, see Moore (1995).

by anti-arrhythmic medications in the intervening years.

There is a widespread idea that the way industry influences science or scientists is by paying scientists to get certain outcomes—that is, by scientific fraud. But as Holman and Bruner (2017) emphasize, the pharmaceutical industry did not need to buy off or even influence researchers in this case to have serious effects on the progress of science. Instead, they simply selected who was to receive funding. Holman and Bruner use the network epistemology framework to explore how industrial influence might shape a community in cases like this. In their model, not all scientists draw from the same bandit arms. Instead, the assumption that scientists use different methods means that some scientists pull arms that are biased in one direction or another. They assume that arm B is better than arm A, and most studies reflect this fact, but when some scientists pull arm B their payoffs are worse than A. Furthermore, the model assumes that funding levels influence the productivity of scientists, and that industry can shape these funding levels. In particular, industrial agents choose to fund only those scientists whose results tend to support A.

As they show, this can create a situation where results supporting the worse belief flood the community and convince many other researchers of the wrong thing. In addition, as they point out, various feedback loops can exacerbate the effect. Successful scientists who have received large grants are often better at placing students, and these students tend to use their (in this case faulty) methods. They are also more likely to get independent government grants to support their research. The result is a community with widespread beliefs in the worse theory.

What is striking about this model, and the historical case, is that the pharmaceutical industry did not do what we typically think of when it comes to industry influence on science. They did not change the research practices that individual scientists were engaged in. There was no corruption or fraud. Instead they made use of the natural variation within a community to shape outcomes in a much more insidious way. In the next section we will discuss another set of models looking at subtle and surprising industry influences on scientific belief.

3.5 The Tobacco Strategy

The historians of science Oreskes and Conway (2011) painstakingly document how, starting in the 1950s, tobacco companies managed to spread public doubt about the growing consensus that cigarettes were dangerous. Their strategy, which Oreskes and Conway call "The Tobacco Strategy", involved fighting science with science for the first time. There were various components to this strategy. In Weatherall et al. (2018) we use models to explore the workings of several of these components.[10]

In 1954, six major US tobacco companies started an organization called The Tobacco Industry Research Committee, which was headed by a prominent geneticist. The ostensible goal of the committee was to fund research into

[10]See also Lewandowsky et al. (2019).

whether tobacco smoking was dangerous to health. In fact, the group was a propaganda machine. In 1957, for example, they produced a pamphlet called "Smoking and Health" that emphasized independent research finding no link between tobacco and cancer while downplaying research that did find such a link. This pamphlet was distributed to hundreds of thousands of doctors and dentists.

We describe this sort of case as one of *selective sharing*. Industry propagandists were not producing biased research, and, in this particular case, were not even intervening in the scientific community in any way. Instead, they were taking advantage of the fact that scientific evidence is probabilistic. Remember, some studies on the link between cancer and cigarette smoke will not find any connection. Industry actors can widely publicize just these studies, while failing to mention the larger body of data supporting a link between smoking and cancer.

In our model, we supplement the basic network epistemology model with two new sorts of actors. The first we call *observers*. These actors do not test the world themselves, but they have credences, and they update these credences on the basis of evidence. They correspond to members of the public who are interested in developing true beliefs, but do not have the tools to gather evidence. Second, we add a *propagandist*. This agent scours the scientific community and in each round of the model shares only data that spuriously suggests arm A is better to each of the observers. In this way they engage in selective sharing, and bias the total body of data seen by each observer towards the worse belief.

We analyze these models and ask: suppose the scientific community settles on successful beliefs, do the observers do so as well? Can the propagandist confound them simply by sharing real, independent data from the community? What we find is that for many parameter settings the propagandist can, indeed, convince the public of the worse belief. The problem at hand and the behaviors of the scientists, though, determine how easy it is for the propagandist to do this.

In particular, we focus on the role of study power in this process. In the model, as discussed, agents can gather different amounts of data each round. They could pull arm B 10 times, for example, or 1000. Less data corresponds to real world studies that are lower powered. For these smaller studies, it is more likely that each one happens to support the worse theory. A large study on the other hand is very likely to support the better theory. What this means is that in communities where scientists run low powered studies, propagandists have more material to work with, and are better at deceiving the public. The take-away is that when public belief is at stake, scientists should maintain high standards.

This observation also tells us something about the best strategies for propagandists. Suppose that the tobacco industry was funding research themselves, but that the scientists involved were unwilling to commit fraud. The smaller their studies, the more likely they are to generate a spurious result that industry can use in their best interests.

4 Conclusion

In this short piece, we have outlined some of the ways the network epistemology framework has been used to explore the workings of scientific communities. In particular, there are a number of lessons here about false belief and how scientific communities can go wrong. Biases towards conformity and an inclination to ground scientific trust in shared belief can hurt the knowledge producing capacity of a community. In addition, pernicious influencers, such as industry propagandists, can make use of subtle strategies that do not subvert the norms of science, but nonetheless mislead.

In general, this overview shows some of the reasons that formal models are useful in studying scientific communities. First, they allow us to represent and explore processes that, in the real world, are very hard to get empirical access to because of how large and extended they are. Second, they allow us to engage in interventions that tell us something about causal effects. We can add more conformity to the model, and observe that actors do worse. Or we can see what happens when study power goes up, to figure out how study power is linked to outcomes. It is typically impractical to make interventions like these in real communities of scientists. Last, the models allow us to gain understanding about causal processes by removing factors that are at play in the real world. For instance, real world actors have many psychological biases and are sometimes bad at reasoning. In the real world, conformist biases and aspects of social trust are at play at the same time. By stripping away such factors, the models let us focus on just one aspect of the problem at a time.

Of course, stripping away factors that are important in the real world can also be dangerous in a model. It could be that other causal factors interact with things like conformity bias, or industrial selection, to negate, or seriously alter their consequences. For this reason social modeling results like those discussed here must always be taken with a grain of salt and supplemented with historical and empirical work.

References

Asch, Solomon E and H Guetzkow (1951). "Effects of group pressure upon the modification and distortion of judgments." *Groups, leadership, and men*, 222–236.

Bala, Venkatesh and Sanjeev Goyal (1998). "Learning from neighbours." *The review of economic studies*, *65*(3), 595–621.

Bond, Rod and Peter B Smith (1996). "Culture and conformity: A meta-analysis of studies using Asch's (1952b, 1956) line judgment task.." *Psychological bulletin*, *119*(1), 111.

Borg, AnneMarie, Daniel Frey, Dunja Šešelja, and Christian Straßer (2017). "Examining Network Effects in an Argumentative Agent-Based Model of Scientific Inquiry." *International Workshop on Logic, Rationality and Interaction*. 391–406.

Bramson, Aaron, Patrick Grim, Daniel J Singer, William J Berger, Graham Sack, Steven Fisher, Carissa Flocken, and Bennett Holman (2017). "Understanding Polarization: Meanings, Measures, and Model Evaluation." *Philosophy of Science*, *84*(1), 115–159.

Frey, Daniel and Dunja Šešelja (2017a). "Robustness and idealizations in agent-based models of scientific interaction.".

Frey, Daniel and Dunja Šešelja (2017b). "What is the Function of Highly Idealized Agent-Based Models of Scientific Inquiry?.".

Grundy, Isobel et al. (1999). *Lady Mary Wortley Montagu.* Clarendon Press.

Holman, Bennett and Justin Bruner (2017). "Experimentation by industrial selection." *Philosophy of Science*, *84*(5), 1008–1019.

Holman, Bennett and Justin P Bruner (2015). "The problem of intransigently biased agents." *Philosophy of Science*, *82*(5), 956–968.

Kidd, Mark and Irvin M Modlin (1998). "A Century of Helicobacter pylori." *Digestion*, *59*(1), 1–15.

Kummerfeld, Erich and Kevin JS Zollman (2015). "Conservatism and the scientific state of nature." *The British Journal for the Philosophy of Science*, *67*(4), 1057–1076.

Lewandowsky, Stephan, Toby D Pilditch, Jens K Madsen, Naomi Oreskes, and James S Risbey (2019). "Influence and seepage: An evidence-resistant minority can affect public opinion and scientific belief formation." *Cognition*.

Mayo-Wilson, Conor (2014). "The Reliability of Testimonial Norms in Scientific Communities." *Synthese*, *191*(1), 55–78.

Mayo-Wilson, Conor, Kevin JS Zollman, and David Danks (2011). "The independence thesis: When individual and social epistemology diverge." *Philosophy of Science*, *78*(4), 653–677.

Mohseni, Aydin and Cole Randall Williams (2017). "Truth and Conformity on Networks." (Working Paper).

Moore, Thomas J. (1995). "Deadly Medicine: Why Tens of Thousands of Heart Patients Died in America's Worst Drug Disaster.".

O'Connor, Cailin and James Owen Weatherall (2019). *The Misinformation Age: How False Beliefs Spread.* New Haven: Yale University Press.

Oreskes, Naomi and Erik M Conway (2011). *Merchants of doubt: How a handful of scientists obscured the truth on issues from tobacco smoke to global warming.* Bloomsbury Publishing USA.

O'Connor, Cailin and James Owen Weatherall (2018). "Scientific polarization." *European Journal for Philosophy of Science*, *8*(3), 855–875.

Palmer, ED (1954). "Investigation of the gastric spirochetes of the human.." *Gastroenterology*, *27*, 218–220.

Rogers, Everett M (2010). *Diffusion of innovations.* Simon and Schuster.

Rosenstock, Sarita, Justin Bruner, and Cailin O'Connor (2017). "In Epistemic Networks, Is Less Really More?." *Philosophy of Science*, *84*(2), 234–252.

Weatherall, James Owen and Cailin O'Connor (2017). "Do as I Say, Not as I Do, or, Conformity in Scientific Networks.".

Weatherall, James Owen, Cailin O'Connor, and Justin Bruner (2018). "How to Beat Science and Influence People." DOI: 10.1093/bjps/axy062. *British Journal for Philosophy of Science.*

Weatherall, James Owen and Cailin O'Connor (2018). "Endogenous Epistemic Factionalization.".

Zollman, Kevin James Spears (2010a). "Social structure and the effects of conformity." *Synthese*, *172*(3), 317–340.

Zollman, Kevin JS (2007). "The communication structure of epistemic communities." *Philosophy of science*, *74*(5), 574–587.

Zollman, Kevin JS (2010b). "The epistemic benefit of transient diversity." *Erkenntnis*, *72*(1), 17.

Zollman, Kevin JS (2013). "Network epistemology: Communication in epistemic communities." *Philosophy Compass*, *8*(1), 15–27.