

# The Problem of Intransigently Biased Agents

Bennett Holman and Justin P. Bruner\*†

---

In recent years the social nature of scientific inquiry has generated considerable interest. We examine the effect of an epistemically impure agent on a community of honest truth seekers. Extending a formal model of network epistemology pioneered by Zollman, we conclude that an intransigently biased agent prevents the community from ever converging to the truth. We explore two solutions to this problem, including a novel procedure for endogenous network formation in which agents choose whom to trust. We contend that our model nicely captures aspects of current problems in medical research and gesture at some morals for medical epistemology more generally.

---

**1. Introduction.** The emergence of social epistemology has provided both a new range of philosophical problems and new formal tools to address them. A salient aspect of this new approach has been the examination of how information is shared among agents in a group. Surprisingly, features that would intuitively seem to be epistemic virtues, such as free exchange of information, can turn out to inhibit the group from acquiring true beliefs (Zollman 2007). More generally, instead of one optimal communication structure, epistemic virtue depends crucially on the particular problem confronting the group

\*To contact the authors, please write to: Bennett Holman, Underwood International College, Yonsei University 50 Yonsei-ro, Seodaemun-gu, Seoul, 03722 Republic of Korea; e-mail: bholman@yonsei.ac.kr. Justin P. Bruner, School of Philosophy, The Australian National University, Coombs Building 9, Acton ACT 0200, Australia; e-mail: Justin.Bruner@anu.edu.au.

†This work was largely completed while the authors were graduate students at the University of California, Irvine. A special thanks to Kyle Stanford for his encouragement and the trust he placed in me (B. H.) to find my own way. We both owe a debt of gratitude to Brian Skyrms and the rest of the UCI Social Dynamics Seminar for providing direction during the formative stages of this work. Thanks also to audiences at the Formal Epistemology Workshop and the Understanding Epistemic Injustice Conference for their feedback on earlier versions of this work. Finally, our most sincere appreciation goes to Benjamin Rin for reading a final draft with characteristic persnickiness.

Philosophy of Science, 82 (December 2015) pp. 956–968. 0031-8248/2015/8205-0019\$10.00  
Copyright 2015 by the Philosophy of Science Association. All rights reserved.

(Zollman 2013). This article will consider a formal model of one problem that increasingly confronts diverse areas of scientific inquiry: the problem of intransigently biased agents.

Previous studies have assumed that research is conducted by agents who, broadly speaking, are interested in discovering the truth (e.g., Alexander 2013). But there are broad swaths of science where those who are financially backing research do so with the express aim of promoting a claim regardless of underlying facts. Tobacco companies funded work that delayed the establishment of a causal link between secondhand smoke and lung cancer, the potential consequences of regulation and taxes for the energy sector led the fossil fuel industry to fund studies that controverted the reality of anthropogenic climate change (Oreskes and Conway 2010), chemical companies fund research that minimizes the effects of exposure to toxic substances in order to reduce their legal liability (Elliott 2011), and so on. The presence of financial interests in these domains fundamentally alters the incentives that drive scientific inquiry. Consequently, epistemically motivated inquirers in these areas must contend with intransigently biased agents.

To illustrate this problem, we will first examine the use of diethylstilbestrol (DES) as a prophylactic for miscarriage. Next, we will review networks of agents confronting the bandit problem as a model of social learning. The static nature of communication in these networks exacerbates the problem of intransigently biased agents, suggesting that if agents were allowed to choose whom to trust, they might be able to avoid manipulation. Our article uncovers that such freedom can render biased agents ineffective at misleading the community.

**2. DES: Four Decades of Intransigence.** In the decades before antibiotics revolutionized medical care, endocrinology was in ascendance. But like many advances, endocrinology brought with it excess enthusiasm in both a host of legitimate products and similar-sounding quack remedies. It is somewhere in the penumbra of legitimacy that we find DES. The synthetic estrogen began with an excellent pedigree. In 1939, the British Medical Research Council reported favorably for its use in several conditions related to menstruation and menopause. Because no patent was sought, any drug company that wished to could manufacture and market DES.<sup>1</sup>

In 1941, 12 companies gained US Food and Drug Administration (FDA) approval to use DES to ameliorate the symptoms of menopause, and 7 years later it was approved as a prophylactic for miscarriage. Yet not all of the research on DES was favorable. Pervasive side effects, as well as animal studies demonstrating that DES was carcinogenic, led to some at the American Medical Association to recommend that it not be recognized for general use,

1. This account is heavily indebted to the excellent scholarship of Dutton (1988).

characterizing contemporary practices as “overzealous . . . indiscriminate and excessive” (Stoddard, quoted in Dutton 1988, 47).

Among the early proponents for the prevention of miscarriage were Harvard professors Olive and George Smith, whose research formed the substantive basis for FDA approval. Though well respected, the work was not without its critics, and within 5 years four separate (methodologically superior) studies had shown that DES was ineffective. Unfortunately, the FDA had come to the conclusion that it lacked the legal authority to remove ineffective products from the market. Although the FDA was explicitly awarded such legal authority in 1962, it would take until 1971 before officials concluded that DES was contraindicated for pregnant mothers.

Meanwhile, roughly 100,000 prescriptions per year were written throughout the 1960s. By the end of its period of use, at least 3% of the nation's children had been exposed to DES in utero, in addition to the millions of mothers who had ingested it (Meyers 1983). Ultimately, it fell from favor in large part as a result of the actions of patients who brought public attention to the increase in cancer, deformed genitalia, and fertility problems. But given that many of these problems were known or suspected from the start, a perennial question has been, what explains the continued use of an ineffective and dangerous drug? The following possible answers are considered below: studies published in medical journals, expert opinion, a doctor's and their colleagues' experience, and information provided by pharmaceutical companies.

We have already seen that DES was not supported by the medical literature. By 1954, over 2,000 women had already participated in four randomized clinical trials, all of which failed to support efficacy, and the largest of which showed that DES increased miscarriages.<sup>2</sup> As for experts, while the Smiths never recanted, their position was increasingly isolated. Internal memos document that the companies themselves were aware that use of DES became rejected by the medical elite (Dutton 1988).

Thirdly, there is the experience of the doctors themselves. Given that DES exacerbated the problems it was prescribed to ameliorate, a doctor's experience should lead her to the conclusion that DES was ineffective at best. This is too fast for two reasons. First, a doctor might simply encounter a random string of live births and mistake it for drug efficacy. Second, the doctor might be so sold on an intervention that failures are perceived as successes. However, the former reason would not explain such widespread use, and the latter possibility only pushes the question back as to where a doctor's enthusiasm came from. While some fervor might be attributed to the progress of medicine in general, the lion's share can be found in the information provided by pharmaceutical companies.

2. See Bamigboye and Morris (2003) for a retrospective analysis of the available literature.

One common and influential source of information for doctors was the *Physicians' Desk Reference (PDR)*. The information contained in the *PDR* was submitted by the manufacturer and then sent out to doctors free of charge. In 1960, over half of busy doctors consulted it daily (Dutton 1988). Beginning in 1947, DES was listed as indicated for "habitual or threatened abortions," with no mention of any disconfirming evidence until 1969, when the indication was dropped and a strong warning against use in pregnancy was added.

Pharmaceutical marketing was both passive and active, and all of it sang the praises of DES. Magazine ads ranged from relatively subdued pieces listing only the claims approved by the FDA to garish ads recommending DES for all pregnancies (see Langston 2010). More active marketing involved company spokesmen (detailers), whose job was to visit doctors and keep them up to date on the companies' products. Corporate memos clarified the approach that detailers were to take with doctors: "Tell 'Em Again and Again and Again—Tell 'Em Till They're Sold and Stay Sold" (Quoted in Dutton 1988, 58).

Many doctors would deny that such sources affect them, claiming they are men and women of science, moved by reason, not the same tricks used to sell soaps. Yet in the case of DES, no other source besides advertising appears to be a viable candidate for explaining such widespread and enduring use. In the face of detailers telling doctors and telling them again, many doctors were sold and stayed sold. Moreover, it seems that nothing that the doctors told the detailer could change their mind. Even a biased agent would have been moved to reconsider their position if they were in search of the truth, but detailers and other sources of information like the *PDR* were not simply biased; they were intransigently biased.

**3. Network Structure and the Bandit Problem.** We now move to a precise formal framework that can help us better understand the influence that biased agents have on a group of epistemically pure agents. In particular, we examine a network of individuals that are all confronted by the so-called bandit problem, a situation in which one is presented with two slot machines and must determine which to play. Zollman suggests that this is analogous to a doctor determining which of two medications to administer.

Doctors are modeled as Bayesian learners, who update their beliefs when presented with new evidence, and are myopic in the sense that they simply administer the drug they believe is more efficacious. Moreover, there is no guarantee that an individual doctor will correctly identify the more efficacious drug. Consider the following scenario: a doctor has observed 5,566 successes upon administering drug A 10,000 times, and only 10 successes upon administering drug B 20 times. In this case our agent will believe that drug A is superior, but clearly, since comparatively little is known about drug B, the optimal long-run strategy may include prescribing B to gain more information.

The myopic doctors considered in the course of this article, however, will only begin to prescribe B if the success rate of A falls under 50%.<sup>3</sup>

In our model, the doctors do not know the true success rates of drugs A and B. In each interval doctors administer the drug they believe to be superior to their  $N$  patients—where each patient has probability  $p_A$  (or  $p_B$ ) of recovering—and record what percentage recover. Doctors are embedded in a social network and treat results obtained by their neighbors on par with their own experience. As figure 1 indicates, epistemic agents are represented as nodes in a graph, and those nodes connected by a line are said to be “neighbors.”

With the society of knowers in view, we can now ask some interesting questions; chief among them, how should the group communicate in order to maximize the likelihood that every member will learn which drug is superior? While agents in the maximally connected graph reach consensus more quickly, the agents in the cycle are more likely to reach a true consensus (Zollman 2007). This counterintuitive finding occurs because, as connection density increases, the entire group is likely to be converted from the superior option by a chance wave of poor results. By contrast, the cycle promotes situations in which the group as a whole stays undecided for longer and there is at least one member collecting data on each option, a phenomenon Zollman (2010) calls “transient epistemic diversity.”

We find in the present article that these results only hold so long as agents are epistemically pure. Generally speaking, an impure agent is an agent interested in convincing the group of a view irrespective of the truth. An epistemically impure agent in the medical field is an agent, such as a pharmaceutical company representative, who attempts to encourage doctors to use a drug irrespective of which drug is more efficacious. In our simulations, epistemically impure agents administer only their favored drug, and the results they obtain are produced by a biased distribution. Hence, if the actual probability of success is 56% and the bias is 10%, the impure agent reports data as if the probability of success is 66%. Specifically, their data come from a binomial distribution with a mean of  $56\% + b$ , where  $b$  is the strength of the bias. This is our attempt to capture, in our idealized model of medical epistemology, the fact that pharmaceutical manufacturers find numerous ways to subtly bias their results.<sup>4</sup>

We focus primarily on the “worst-case scenarios” in which the pharmaceutical company promotes the inferior drug and is connected to all doc-

3. Myopia is plausible in several cases, including where a doctor feels ethically prohibited from giving a patient a drug perceived to be worse just to increase the doctor’s confidence that the drug is inferior.

4. For a nonexhaustive list see Safer (2002).




Figure 1. Each node is an agent, and each line represents a two-way communication channel between the agents. We refer to these three canonical structures as “the cycle” (*left*), “the wheel” (*middle*), and “the complete graph” (*right*).

tors.<sup>5</sup> In terms of Zollman’s canonical network structures, it is a wheel with the biased agent at the hub. Briefly, the exact setup is as follows. Agents are randomly assigned beliefs regarding the two available drugs. Doctors, as well as the biased agent, administer the drug they believe to be most efficacious. This generates  $N$  data points, and all share their data with those they are connected to. Doctors then update their beliefs in a fashion outlined by Zollman (2010) and then repeat this process.

**4. The Impossibility of Sustained Convergence to the Truth.** Consider the case in which drug A is successful with probability .51 and the pharmaceutical company’s drug (B) is slightly inferior ( $p_B = .5$ ). Assume that all doctors begin with true beliefs regarding both drugs. Given this belief profile, all will immediately begin administering drug A to their patients. We have convergence in the short run, but not in the long run. This is due to the bias of the pharmaceutical company (assume that the bias is .03). Since the pharmaceutical company is the only one conducting research on drug B, they alone influence the doctors’ perceptions about it. As the doctors’ impressions of drug B improve, one of the doctors will “cross over” and begin to administer drug B. By doing so, she is now running her own unbiased experiment. This, in turn, helps to mitigate the influence the pharmaceutical company has on everyone she is connected to, including herself. Thus, if she and two of her neighbors both switch over to drug B, the combined results of their experiments are sufficient to mitigate the influence of the pharmaceutical company and move back to the superior drug. Yet when none of the doctors investigate drug B, the only information they receive about the drug, once again, comes from the biased pharmaceutical company. We now

5. Given manufacturers’ ability to organize “educational events” and fund “key opinion leaders,” maximum connectivity is a reasonable approximation of reality (Elliott 2010; cf. Krimsky 2003).

see why convergence to the superior drug for a sustained amount of time is impossible in Zollman-type models with intransigently biased agents.

In order to quantify this effect, we look to the last 1,000 rounds of a 2,000-round simulation and determine how frequently the best drug was used. We find that six doctors arranged on the wheel use the superior drug 42% of the time (see fig. 2). Interestingly, this number increases as we add more connections to the network. In the complete network, doctors utilize the superior drug 63% of the time. Thus, in contrast to Zollman (2007), the more connections there are, the more likely the network as a whole is to adopt the more efficacious treatment. The reasons for this should be obvious. When doctors are better connected to each other, fewer doctors have to spend their time debunking the biased results because the unbiased results are more widely broadcast.

Though more connected networks provide a defense against intransigently biased agents, nothing short of eternal vigilance is required of the community—the community must constantly devote members to the investigation of the less successful drug. This keeps the biased agent at bay, but it is surely a second-best solution. We argue that what is primarily driving the need for eternal vigilance is the fact that experimental results from one agent are taken just as seriously as experimental results from another. If individual doctors could learn that the pharmaceutical company is severely biased, doctors may begin to discount the company's results. The above models are static—epistemic agents must listen to everyone they are connected to. If this assumption is relaxed and we consider dynamic networks, in which connections can change, our doctors may learn to ignore the pharmaceutical company. We now turn our attention to endogenous network formation and see that if individuals have some control over whom they listen to, then for a wide variety of parameters, the pharmaceutical company is unlikely to draw doctors away from the most efficacious drug.

**5. Choosing Your Neighbors: Endogenous Network Formation.** Modeling network formation is an active area of research in a number of disparate fields.<sup>6</sup> Unfortunately, none of the canonical models can be appropriately applied to our epistemic community because our agents are continuously generating data. As in our earlier model, a doctor  $i$  is connected to doctor  $j$  if  $i$  is somehow influenced by the experimental findings of  $j$ . However, in this model, network connections now vary continuously and are no longer symmetric, meaning that  $i$  can be strongly connected to  $j$ , while  $j$  is only weakly connected to  $i$ . In this case,  $i$  is strongly influenced by  $j$ , while  $j$  is only slightly influenced by  $i$ . Similar arrangements no doubt do occur, as when the work of a senior scientist is very influential on a junior scientist,

6. See Jackson (2005) for an overview.




Figure 2. Connectivity level refers to the proportion of possible connections between doctors (e.g., the wheel has 6/15 possible connections). Each plus sign represents one simulation. As the connectivity level increases (i.e., connections are added), the doctors prescribe the superior drug more often. Color version available as an on-line enhancement.

but this influence is not reciprocal. In general,  $j$  strengthens her connection to  $i$  if  $i$ 's experimental findings are somehow in line with  $j$ 's subjective beliefs. Likewise,  $j$  weakens her connection to  $i$  the more that  $i$ 's experimental findings seem to clash with  $j$ 's beliefs. Making this precise is difficult and highlights why many models of endogenous network formation used in economics and sociology are not applicable when thinking about our epistemic network.

We instead present a novel model of endogenous network formation that replicates basic hypothesis testing inside an epistemic community of agents that are continuously experimenting. Consider a network of  $D$  doctors. Each doctor has  $D + 1$  bins (one for each doctor and one for the pharmaceutical company) that initially have anywhere from 0 to 100 balls in them. Let  $B_i$  be the vector  $\langle b_{i1}, b_{i2}, \dots, b_{iD+1} \rangle$ , where  $b_{il}$  is the number of balls in agent  $i$ 's first bin. How strongly connected agent  $i$  is to agent  $j$  is determined by the proportion  $b_{ij} / \sum_{k \in D+1} b_{ik}$ . This connectedness determines how much weight  $i$  puts on the experimental findings of  $j$  (call this  $w_{ij}$ ). Agent  $i$  updates her beliefs regarding drug A as in Zollman (2010), except that the results are weighted as follows:

$$P(\text{drug A works} | \text{agent } j \text{ has } s \text{ successes in } N \text{ trials}) = (\alpha + w_{ij}s) / (\alpha + \beta + w_{ij}N),$$



where  $\alpha$  and  $\beta$  are the agent's values from the previous round. *Ceteris paribus*, the more balls agent  $i$  has in her  $j$ th bin, the more connected she is to agent  $j$  and thus the larger impact agent  $j$  has on  $i$ 's beliefs. Individuals adjust their connections in the following fashion. Upon receipt of  $N$  data points from agent  $j$ , agent  $i$  conducts a one-sample  $t$ -test based on her subjective beliefs. Let  $t_{ij}$  be the  $t$ -score agent  $i$  assigns to agent  $j$ 's experimental results in round  $r$ . The number of balls in  $b_{ij}$  is then updated by the following equation:

$$b_{ij}(r + 1) = b_{ij}(r) + f(t_{ij}), \text{ where } f(x) = \Lambda(1.96 - |x|) \text{ and } \Lambda > 0.$$

Here  $b_{ij}(r)$  is the number of balls in agent  $i$ 's  $j$ th bin at round  $r$ . Thus, a  $t$ -score with an absolute value of less than 1.96 results in an increase in the number of balls in the bin, while a  $t$ -score with an absolute value exceeding 1.96 results in a decline. How the strength of connection to  $j$  is affected can of course only be determined if we take into account the change in all bins. One intuitive property this update rule satisfies is the following: if you are connected to two individuals and they repeatedly provide you the same evidence, then in the long run you should expect to be equally connected to these two individuals. One's initial connectivity "washes out" in the end.

We find that the inclusion of network formation has drastic effects. One common outcome is for all doctors in the community to heavily discount the pharmaceutical company's experimental data. In this case, none of the doctors administer the inferior drug, and all have minimal connections to the pharmaceutical company. The biased agent is effectively squelched, thereby allowing doctors to converge on the superior drug. Less desirable arrangements are also possible. In some scenarios a minority of agents listen to both their fellow doctors and the pharmaceutical company. The level of connection these agents have to the company does not completely dissipate because the company had a hand in shaping their perception of the drug. The company's biased experimental results are thus not seen as particularly unusual, since they are in some sense already reflected in these doctors' subjective beliefs.

By and large, however, a dynamic network helps the community to better identify the superior drug. For example, in the static network with  $p_A = .51$ ,  $p_B = .50$ , and  $b = .08$ , doctors almost never come to prescribe the superior drug. In contrast, the superior drug is prescribed 80% of the time in the dynamic network. In general, dynamic networks are much more resistant to the influence of intransigently biased agents than static networks, and figure 3 drives this point home quite nicely.

Two variables are primarily responsible for ensuring that the more effective drug is taken up by the population:  $N$  and  $b$ . As  $N$  increases, the community becomes more likely to converge on the better drug. Surprisingly,




Figure 3. Proportion of times the more efficacious drug is administered in the last 1,000 rounds for different levels of  $b$  (pharmaceutical company bias), ranging from 0 to .15, as well as various levels of  $N$  (i.e., 500, 250, 100, 50). Color version available as an online enhancement.

convergence on the superior drug is also more probable when the company is highly biased. All else being equal, if the bias is outlandish, then even a small number of trials will be able to alert the community that something is awry. Introducing biased data can influence honest agents, but lies have to be subtle enough to go undetected. In a dynamic network, agents can simply stop listening if the bias becomes apparent.

### **6. The Problem of Intransigently Biased Agents and Epistemic Clarity.**

The problem posed by intransigently biased agents can be alleviated if agents learn to identify and trust good informants. We have seen that this is not possible in a static network, since by decree individuals cannot come to ignore their neighbors, thereby allowing a biased agent to mislead the community. Furthermore, Zollman's finding that "in small finite groups, the best graphs are minimally connected" (2013, 25) fails to obtain with the introduction of biased agents. Instead of promoting a virtuous transient epistemic diversity, the lack of communication forces sparsely connected agents to duplicate the debunking work—if they are able to resist the biased agent at all.

The introduction of our network formation rule yields desirable results. While other update rules may be superior, this simple rule prevents agents from being manipulated by a highly biased pharmaceutical company. It creates a point at which increasing the bias in one's results merely makes it easier to be identified as untrustworthy. Even in cases where the pharmaceutical company retains some influence with most doctors, groups virtually never converge to the wrong drug, and under most circumstances reviewed here, they prescribe the right drug more often than not. Indeed, one common result is that every doctor gives no weight to the pharmaceutical company and roughly equal weight to everyone else.

Returning to the DES case, doctors most closely approximate agents in the static wheel. Each doctor was in contact with a limited number of colleagues but maintained contact with the pharmaceutical company via advertisements, the *PDR*, and interactions with detailers. Thus, despite the experimental evidence, elite opinion, and the doctor's own experiences, use of DES continued apace. The models considered here suggest two possible responses: increase the number of connections, or learn to ignore biased agents (e.g., stop meeting with detailers).

It might be suggested that the DES disaster could have been averted if doctors had simply been trained to pay attention to the results of randomized clinical trials, as is currently recommended by the evidence-based medicine movement. As described above, these results showed a lack of efficacy, but note that this just pushes the problem back. As doctors have become more influenced by research, pharmaceutical companies have come to spend an increasing amount of their marketing budget on biased trials

(Angell 2004). A number of meta-analyses have found a large correlation between positive results and industry funding (Bekelman, Li, and Gross 2003). Rochon et al. (1994) found that 56/56 comparison trials funded by manufacturers of nonsteroidal anti-inflammatory drugs for arthritis concluded that the funder's product was as good as or better than the comparison drug. While this was particularly egregious, it is estimated that between 89% and 98% of trials yield results favorable to the company that funded the research (Cho and Bero 1996).

Given the severity of this problem, some commentators have suggested that pharmaceutical companies be prohibited from conducting such research. An alternative to such a fundamental change in the structure of scientific practice is to better exercise epistemic discrimination. Though it is rare, official bodies have occasionally considered devaluing the epistemic weight accorded to industry-funded studies, a proposal that the National Institute for Health and Care Excellence, a British advisory agency, considered but ultimately rejected. The present analysis suggests that something like our network formation rule may be preferable to the current practice of treating all equally well designed trials as equivalent regardless of their source.

## REFERENCES

- Alexander, Jason. 2013. "Preferential Attachment and the Search for Successful Theories." *Philosophy of Science* 80:269–82.
- Angell, Marcia. 2004. *The Truth about Drug Companies*. New York: Random House.
- Bamigboye, Anthony, and Jonathan Morris. 2003. "Oestrogen Supplementation, Mainly Diethylstilbestrol, for Preventing Miscarriages and Other Adverse Pregnancy Outcomes." *Cochrane Database of Systematic Reviews* 3, Art. No. CD004353.
- Bekelman, Justin, Yan Li, and Cary Gross. 2003. "Scope and Impact of Financial Conflicts of Interest in Biomedical Research." *JAMA* 289:454–65.
- Cho, Mildred, and Lisa Bero. 1996. "The Quality of Drug Studies in Symposium Proceedings." *Annals of Internal Medicine* 124:485–89.
- Dutton, Diana. 1988. *Worse than the Disease: Pitfalls of Medical Progress*. Cambridge: Cambridge University Press.
- Elliott, Carl. 2010. *White Coat Black Hat: Adventures on the Dark Side of Medicine*. Boston: Beacon.
- Elliott, Kevin. 2011. *Is a Little Pollution Good for You? Incorporating Societal Values in Environmental Research*. New York: Oxford University Press.
- Jackson, Mathew. 2005. "A Survey of Models of Network Formation: Stability and Efficiency." In *Group Formation in Economics: Networks, Clubs and Coalitions*, ed. Gabrielle Demange and Myrna Wooders, 11–57. Cambridge: Cambridge University Press.
- Krimsky, Sheldon. 2003. *Science in the Private Interest: Has the Lure of Profits Corrupted Medical Research?* Lanham, MD: Rowman & Littlefield.
- Langston, Nancy. 2010. *Toxic Bodies*. New Haven, CT: Yale University Press.
- Meyers, Robert. 1983. *DES: The Bitter Pill*. New York: Seaview/Putnam.
- Oreskes, Naomi, and Erik Conway. 2010. *Merchants of Doubt: How a Handful of Scientists Obscured the Truth on Issues from Tobacco Smoke to Global Warming*. New York: Bloomsbury.
- Rochon, Paula, Jerry Gurwitz, Robert Simms, Paul Fortin, David Felson, Kenneth Minaker, and Thomas Chalmers. 1994. "A Study of Manufacturer-Supported Trials of Nonsteroidal Anti-inflammatory Drugs in the Treatment of Arthritis." *Archives of Internal Medicine* 154:157–63.

- Safer, Daniel. 2002. "Design and Reporting Modifications in Industry-Sponsored Comparative Psychopharmacology Trials." *Journal of Nervous and Mental Disorder* 190:583–92.
- Zollman, Kevin. 2007. "The Communication Structure of Epistemic Communities." *Philosophy of Science* 74:574–87.
- . 2010. "The Epistemic Benefit of Transient Diversity." *Erkenntnis* 72:17–35.
- . 2013. "Network Epistemology: Communication in Epistemic Communities." *Philosophy Compass* 8:15–27.