

University of Groningen

How to Beat Science and Influence People

Weatherall, James Owen; O'Connor, Cailin; Bruner, Justin P.

Published in:
British Journal for the Philosophy of Science

DOI:
[10.1093/bjps/axy062](https://doi.org/10.1093/bjps/axy062)

IMPORTANT NOTE: You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.

Document Version
Publisher's PDF, also known as Version of record

Publication date:
2020

[Link to publication in University of Groningen/UMCG research database](#)

Citation for published version (APA):

Weatherall, J. O., O'Connor, C., & Bruner, J. P. (2020). How to Beat Science and Influence People: Policy-Makers and Propaganda in Epistemic Networks. *British Journal for the Philosophy of Science*, 71(4), 1157-1186. <https://doi.org/10.1093/bjps/axy062>

Copyright

Other than for strictly personal use, it is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license (like Creative Commons).

The publication may also be distributed here under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license. More information can be found on the University of Groningen website: <https://www.rug.nl/library/open-access/self-archiving-pure/taverne-amendment>.

Take-down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Downloaded from the University of Groningen/UMCG research database (Pure): <http://www.rug.nl/research/portal>. For technical reasons the number of authors shown on this cover page is limited to 10 maximum.

How to Beat Science and Influence People: Policymakers and Propaganda in Epistemic Networks

James Owen Weatherall , Cailin O'Connor,
and Justin P. Bruner

ABSTRACT

In their recent book, Oreskes and Conway ([2010]) describe the ‘tobacco strategy’, which was used by the tobacco industry to influence policymakers regarding the health risks of tobacco products. The strategy involved two parts, consisting of (i) promoting and sharing independent research supporting the industry’s preferred position and (ii) funding additional research, but selectively publishing the results. We introduce a model of the tobacco strategy, and use it to argue that both prongs of the strategy can be extremely effective—even when policymakers rationally update on all evidence available to them. As we elaborate, this model helps illustrate the conditions under which the tobacco strategy is particularly successful. In addition, we show how journalists engaged in ‘fair’ reporting can inadvertently mimic the effects of industry on public belief.

- 1 *Introduction*
 - 2 *Epistemic Network Models*
 - 3 *Selective Sharing*
 - 4 *Biased Production*
 - 5 *Journalists as Unwitting Propagandists*
 - 6 *Conclusion*
- Appendix*
-

1 Introduction

In December 1952, *Reader’s Digest*—the publication with the largest circulation in the world at the time—published an article with the title ‘Cancer by the

Carton' (Norr [1952]).¹ Coming on the heels of a period during which smoking rates had risen significantly, the article brought the growing body of evidence that cigarette smoking caused cancer to the attention of a wide audience. Six months later, a group of doctors at Sloan Kettering Memorial Hospital completed a study in which they demonstrated that mice painted with cigarette tar developed malignant carcinomas (Wynder *et al.* [1953]). To a public already primed by the *Reader's Digest* article, this article provided a direct and visceral link between a known by-product of smoking and fatal cancer. It produced a media frenzy, with articles reporting the result in a number of national newspapers and magazines.

The tobacco industry panicked. The heads of six major US tobacco companies contracted with the New York public relations firm Hill and Knowlton to combat claims that their product had serious adverse health effects. At a meeting in December of 1953, they agreed on a novel strategy: they would fight science with science. In 1954 they hired a prominent geneticist, C. C. Little, to run the newly create Tobacco Industry Research Committee (TIRC). The goal of this committee was ostensibly to support and promote research on the health effects of tobacco. But in fact, it was a propaganda machine whose principal activity was to find, fund, and promote scientific research that contradicted the growing consensus that smoking kills. The goal was to create a false sense of uncertainty and controversy, so as to delay policy change in the service of public health.

As historians Oreskes and Conway ([2010]) carefully document, drawing on documents released as part of a series of lawsuits in the 1990s, the tobacco industry initiated this strategy—which Oreskes and Conway dub the ‘tobacco strategy’—despite the fact that their own internal research had already established a link between smoking and various health risks.² The tobacco strategy had multiple parts. One part, which we will call selective sharing, consisted of identifying and promoting research produced independently of industry that happened to support the industry’s preferred position. For instance, in 1954 the TIRC distributed a pamphlet entitled ‘A Scientific Perspective on the Cigarette Controversy’ to nearly 200,000 doctors, journalists, and policymakers, in which they emphasized favourable research and questioned results supporting the contrary view. A second part of the tobacco strategy, which we will call ‘biased production’, consisted of funding their own research, either through in-house scientists or by offering research grants to

¹ The history presented in these opening paragraphs draws heavily on the account given by Oreskes and Conway ([2010], Chapter 1). Note that the chronology we describe differs slightly from theirs, as they report that the *Reader's Digest* article was published in the wake of the 1953 Wynder article, whereas in fact it appeared almost a year earlier. See also the discussion in (O'Connor and Weatherall [2019]).

² For other discussions of this strategy, see, for instance, (Michaels and Monforton [2005]; Michaels [2008]).

scientists pursuing research activities that the industry supported, and then publishing the results of that research only if it supported industry views.

Oreskes and Conway ([2010]) make a compelling case both that the tobacco strategy was self-consciously adopted by industry to influence public opinion, and that those goals were apparently achieved. They also argue that the tobacco strategy succeeded, in the sense of playing a causal role in realizing these ends. But the historical analysis they provide is necessarily *in situ*, and the details are complex. It is difficult to clearly distinguish the causal role played by the TIRC's activities from other cultural and psychological considerations that may have influenced public opinion and slowed policy change, such as the facts that smoking was culturally ingrained, including among doctors and legislators, and that the tobacco industry made significant political contributions. One might worry that it was over-determined that cultural and political change regarding tobacco would be slow, so that it is not clear that the tobacco strategy played a significant role.

In this article, we attempt to strengthen the causal link between the tobacco strategy and public opinion by proposing and analysing an agent-based model of the tobacco strategy. Our model is based on the network epistemology framework introduced by Bala and Goyal ([1998]), later adapted to philosophy of science by Zollman ([2007]), and used more recently by Holman and Bruner ([2015], [2017]) to study industry effects on science.³ This approach allows us to test counterfactual effects of industry propaganda on public belief, free from compounding factors.⁴ As we will argue, our model provides evidence that the tobacco strategy can, in principle, be very effective. We also provide important insight into how (and when) it works. As we will argue, the tobacco strategy trades on the fact that empirical research into difficult, probabilistic problems will yield results that, individually, support more than one position. It is only by viewing a maximal, unbiased sample of the output of a research community that one can hope to come to hold reliable beliefs about the problem at hand. This means that even subtle manipulations, not in the scientific record

³ Similar methods are used in (O'Connor and Weatherall [2018], [2019]; Weatherall and O'Connor [unpublished]) to discuss the propagation of false beliefs in social networks more generally.

⁴ To be clear, we take this ability to intervene in our model to be the key aspect that helps us strengthen the causal link between the tobacco strategy and public policy related to tobacco products. The model allows us to test where and how the presence of a propagandist can influence public belief. It also lets us explore the effects of a propagandist in isolation from other potential causal factors. That said, we do not claim that the factors we consider here are the only relevant causal factors, nor that other models could not be used to identify other counterfactual relationships and possible causal factors. For more on the interventionist account of causation that is lurking in the background here, see (Woodward [2003]); see also our Footnote 8 for more on the relationship between the model we discuss and how-possibly modelling.

itself, but in the sample of that record that one is exposed to, can have disproportionate impact on one's beliefs.⁵

An important take-away from these arguments is that what may seem like weak or subtle interventions are often most effective for the would-be propagandist. In particular, outright scientific fraud—intentional publication of incorrect, fabricated, or purposely misleading results—is not only unnecessary to influence public opinion on topics of scientific inquiry, it is also riskier and often less effective than other forms of manipulation. Biased production, which does not involve fabricating results, is a successful strategy for misleading the public. And in many cases, biased production is itself less effective than selective sharing—even though it involves more direct and explicit interventions on the production of science. These observations support claims by Holman and Bruner ([2017]) that more subtle industry influences—like their 'industrial selection', which we will discuss at length in the conclusion of this article—can be very effective. In fact, as we will also argue in what follows, even interventions with no malicious intent, such as those of journalists who aim to be 'fair' in their reporting on controversial science, can unwittingly mimic the effects of an active propagandist.

Before proceeding, we wish to emphasize what is at stake in this discussion. Recognizing how the tobacco strategy works is essential for understanding the relationship between science, public policy, and democracy. There is now a large literature in philosophy of science on what it means for science to be suitably 'democratic', in the sense of reflecting the needs and values of the members of a society affected by scientific results.⁶ This literature focuses on how democratic ideals should shape scientific research. But the other direction, of how to best understand the role that scientific knowledge plays in democratic institutions, is equally important—and requires us to reckon with the fact that there are powerful political and economic interests that wish to manipulate and interfere with the distribution of scientific knowledge. Understanding the mechanisms by which public understanding of science can fail is crucial to developing institutions and interventions that might be effective in defanging the tobacco strategy, and, in the case of 'fair' reporting, improving journalistic practice to protect public belief.

The remainder of the article will be structured as follows: In the next section we will introduce the model we consider in the article. Sections 3 and 4 will present results concerning selective sharing and biased production,

⁵ This conclusion is in line with prior work in social epistemology exploring how subjects come to form incorrect beliefs on the basis of selective reporting (for an overview, see Romero [2016]). We demonstrate how the public can be systematically misled even in those circumstances where all experiments are published and scientists exhibit no bias in favour of industry.

⁶ See, for instance, (Longino [1990], [2002]; Kitcher [2003], [2011]; Douglas [2009]; Elliott [2011]) and references therein.

respectively. Section 5 will describe results from a slightly modified model in which we consider ways that journalistic practice with no malicious intent can simulate the actions of a propagandist. We will conclude with some reflections on the results, including a discussion of possible interventions.

2 Epistemic Network Models

Bala and Goyal ([1998]) introduce a model where individuals in a community try to choose which of two theories, or practices, is the better one. In order to decide, they use evidence from two sources. First, each individual tests the theory they favour, and second, they share their results with those in their social network. Over time, they update their beliefs about the theories based on this evidence, until, in many cases, the whole group comes to believe that one or the other is better. Zollman ([2007], [2010]) shows how simulations of this model can be applied to understand scientific discovery and consensus. Under his interpretation, each individual in the model is a scientist testing a hypothesis. For instance, suppose a new treatment for cancer is introduced. How do researchers go about deciding whether to adopt the new treatment or stick with the old one?

The model starts with a network, defined by a set of scientists and a set of fixed links representing pathways of communication between them. Inquiry is modelled using what is called a two-armed bandit problem. Each scientist is aware that there are two actions they can perform. Action *A* is well understood, and is known to be successful 50% of the time. Action *B*, on the other hand, is less well understood, and the agents are uncertain whether it is more or less successful than *A*. In particular, the instantiation of the model we use here assumes that *B* is a bit better than *A*, so that its chance of success is $p_B = 0.5 + \epsilon$.⁷ Scientists, however, do not know whether it is a bit better, or a bit worse (success rate $0.5 - \epsilon$). Their aim is then to discover which of these possibilities is the better one, and so adopt the superior action. In what follows, we will refer to ‘theory *A*’ as the view that action *A* is preferable, and ‘theory *B*’ as the view that action *B* is preferable.

At the beginning of a simulation, scientists are given random credences about which theory is better, represented by numbers between zero and one. (For example, a scientist with credence 0.7 believes there is a 70% chance that theory *B* is right.) Each round, scientists apply their preferred theory to the world. Agents with a credence <0.5 believe that theory *A* is probably correct and employ it. Agents with credence >0.5 , on the other hand, use theory *B*. This might correspond to doctors who approve of a new treatment deciding to try it on patients, while others stick with a more conservative approach.

⁷ The payoff of success for both actions is the same, and failure yields no payoff, so action *B* has a strictly higher expected payoff.

In each round, scientists observe how their choices worked, and change their beliefs in light of new evidence. In particular, it is assumed that they test their theory some number of times, n . They then update their credence using Bayes's rule based on what they, and all their network neighbours, have observed. Simulations stop when either all agents believe in theory A —and so do not learn anything new about theory B —or when all agents have reached some certainty threshold (say, 0.99 credence) that B is correct, and so are exceedingly unlikely to end up changing their minds. Both outcomes occur in simulations. Sometimes initial beliefs and/or strings of data supporting theory A are enough to create a mistaken consensus around the worse theory, and sometimes the community gets enough good data to settle on the better one.

As noted in the introduction, we are particularly interested in the role of industry, and propagandists in general, in shaping policy based on science, and in influencing popular belief. For this reason, our models follow the framework developed by Bala and Goyal ([1998]) and Zollman ([2007]), but with a few key alterations. First, we introduce a new set of agents, which we will call policymakers. These agents have beliefs, and update these in light of evidence, but do not gather evidence themselves. Each one listens to some, or all, of the scientists in the network, but does not influence them or anyone else. We assume that at the beginning of each simulation, policymakers are sceptical about the new, uncertain theory, in the sense that their initial credences are sampled from a uniform distribution supported on the interval $(0, 0.5)$.

We also introduce one further agent that we call the propagandist. This agent observes all the scientists, and communicates with all the policymakers. Their goal is to convince the policymakers to prefer the worse theory A , by sharing evidence that suggests action A is preferable to action B . We consider two treatments, corresponding to different strategies by which the propagandist might do this. The idea is to represent an individual, or set of individuals, who are motivated to use all their resources to publicize data that is ultimately misinformation. The actions of the propagandist capture various aspects of the tobacco strategy: by sharing a biased sample of evidence with policymakers (and the public) they seek to divide public opinion from scientific consensus.⁸

⁸ It is common in discussions of idealized models of epistemic communities—and philosophical discussion of modelling in the social sciences generally—to suppose that such models provide, at best, how-possibly information about a target phenomenon. (For a discussion of this view, see Frey and Šešelja [2018].) We wish to explicitly reject this attitude, at least with regard to the inferences we draw from the present model. On our view, there is ample and convincing historical evidence that the tobacco industry has intervened to influence public beliefs using strategies very similar to the ones discussed here. The present model supplements this historical evidence by allowing us to ask counterfactual questions concerning how precisely this kind of intervention works. For this reason, we adopt the position that our model provides 'how-potentially' information concerning the causal role of this sort of intervention. What we mean by this is that

3 Selective Sharing

In the first treatment, which we call selective sharing, the propagandist adopts the following strategy: in each round of the simulation, they search the network of scientists for results that tend to favour theory A , and then they share only those results with the policymakers. The policymakers then update their credence in light of both the evidence they receive from the scientists to whom they are connected and the evidence received from the propagandist, which includes a subset of the results produced by the scientists.⁹ In other words, each policymaker receives, and updates on, a biased sample of evidence including a disproportionate number of results favouring theory A . However, all of the results the policymakers receive are ‘real’ in the sense that they are produced by members of the scientific community, without bias or influence by the propagandist or anyone else. In addition, the propagandist does nothing to influence the scientific community; they only seek to shape policymaker beliefs.¹⁰

We analyse this model by varying parameter values including the number of scientists in the network, K , the number of draws each round, n , the probability that action B is successful, $p_B = 0.5 + \epsilon$, and the number of scientists each policymaker listens to, k .¹¹ (Each policymaker could connect to one scientist, or to two, or to three, and so on.) We consider models where the network of scientists is either a cycle, where each scientist is connected to two peers in a ring structure, or a complete network, where every scientist communicates with every other. Figure 1 shows these two structures. They differ in that it will tend to take longer for the cycle to converge to a consensus, especially for large communities, because fewer scientists share data with each other. This network structure affects our results, as we will make clear, mainly by influencing the amount of time it

the model is realistic enough to explore processes that have the potential to really occur. In this sense, we take this model to lie somewhere on a spectrum between how-possibly and how-actually (that is, fully realistic) models. Exploring how-potentially models can direct further inquiry into whether these processes are really at play, or (as in this case) supplement existing empirical work on these processes (Rosenstock *et al.* [2017]). See also Footnote 4. We are grateful to an anonymous referee for pressing us to clarify our stance on these issues, and for emphasizing that these models lie on a spectrum (a point with which we agree).

⁹ As we discuss below, this model involves some double-counting of results favourable to theory A , as a given policymaker may receive the same evidence from both a scientist and from the propagandist.

¹⁰ In this way, the model we consider differs significantly from that of Holman and Bruner ([2015]), who consider an industry scientist who influences other scientific beliefs. We will return to this relationship in the next section, as the treatment discussed there is somewhat closer to that discussed by Holman and Bruner.

¹¹ We consider subsets of the following parameter values: population size $K = 4, 6, 10, 20, 50$, number of draws $n = 1, 2, 5, 10, 20, 50, 100, 500, 1000$, success of better theory $p_B = 0.51, 0.55, 0.575, 0.6, 0.7$, and number of scientists listened to by each policymaker ranges from one to all of them. Our choices were influenced, to some degree, by tractability considerations; simulations with low p_B and many scientists, in particular, took longer to run.




Figure 1. Cycle and complete network structures.

takes the communities of scientists to reach consensus. All reported results were averaged over 1000 runs of simulation that terminated when either correct or incorrect consensus was reached.¹²

We find that the presence of a single propagandist who communicates only actual findings of (ideal, unbiased) epistemically motivated scientists can have a startling influence on the beliefs of policymakers. No one in this scenario is doing bad science. No one is falsifying data, or even publishing it selectively. The propagandist communicates only real scientific findings. Nonetheless, under many scenarios we find that while the community of scientists converges on true beliefs about the world, the policymakers reach near certainty in the false belief. This point will become clearer as we share results in this section and the next—for a wide range of parameter values, public belief is highly accurate when no propagandist is at play, but as inaccurate as it could possibly be under the influence of the propagandist.

We find that the likelihood of this sort of disagreement between science and public belief depends on features of the community and the problem studied. First, unsurprisingly, when policymakers listen to more scientists they are harder to deceive. Figure 2 shows the average credence that policymakers have in the correct theory at the point where scientists reach a correct consensus, which we understand as a state in which all members of the scientific community have credence greater than 0.99. The x -axis tracks the proportion of scientists in the entire community that each policymaker consults. As we can see, without a propagandist, public beliefs lag behind scientific beliefs slightly, though the more scientists are consulted, the better the public does. With a propagandist, the public is convinced that the false theory is preferable, though if they are connected to more scientists, this

¹² We ignore outcomes where scientists do not converge to the correct belief. (In these cases, policymakers also tend to accept the worse theory.) This means that our results are averaged over fewer than 1000 rounds. In the worst case scenario, though, this only meant ignoring 10% of data points.




Figure 2. Policymaker beliefs as influenced by their connections to scientists.

effect is ameliorated.¹³ In related work, Holman and Bruner ([2015]) model the effects of industrial propagandists who share biased data with scientists and likewise observe that it is best for scientists to be highly connected to each other in such cases.¹⁴

It is also harder to deceive policymakers when the size of the effect—that is, the difference in expected payoff between action *A* and *B*—is larger. As p_B (the probability that the better action pays off) increases, so that the difference between the better and worse theories is more stark, the policymakers are more likely to learn the truth. This is because there are fewer studies that happen to support the worse theory, and thus less material for the propagandist to publicize. Figure 3 shows how as the difference between the success rates of the two theories increases, policymakers tend to end up holding the true belief regardless of propaganda.¹⁵ This observation is especially relevant to the biomedical sciences, because human bodies are highly variable, which means that data is often highly equivocal. We should expect studies of human subjects to provide plenty of material for potential propagandists looking to obscure true evidential trends.

The number of draws each round, n , also has a strong effect on the ability of the propagandist to mislead. This parameter may be thought of as the sample size for each experiment performed by the scientists. For fixed p_B (that is, fixed effect size), it is roughly analogous to the statistical power of each scientist's

¹³ These results are for a cycle network with $K=20$, $n=10$, and $p_B=0.55$. This trend occurred across all parameter values, as did the other trends reported unless otherwise specified.

¹⁴ This shows how the suggestion from Zollman ([2007], [2010]) that less connected scientific communities are more successful at reaching the truth does not hold when industry attempts to sway consensus.

¹⁵ Results are for $K=20$ scientists in a cycle, each policymaker listening to 50% of scientists ($k=10$), and $n=10$. Again the trend is stable across parameter values.




Figure 3. Policymaker beliefs as influenced by the difficulty of the scientific problem.

experiments: the higher it is, the less likely that a study will erroneously support theory A .¹⁶ Thus scientists who run the equivalent of well-powered studies and therefore produce fewer spurious results also produce less for the propagandist to work with. Figure 4 shows the final average beliefs of policymakers with and without a propagandist as n increases.¹⁷ Note that the x -axis is on a log scale to make the trend more visible. For high enough n , the propagandist can do almost nothing since they very rarely have the spurious results necessary to mislead the public. In these cases, policymakers figure out the truth despite the propagandist's efforts.¹⁸

One might think that increasing the number of scientists in a network, K , holding all else fixed, would decrease the effectiveness of the propagandist. More scientists means more data, and more epistemically motivated minds searching for truth. But this does not occur in general in these models. Assume that policymakers listen to some fixed number of scientists, k . This is not an

¹⁶ More precisely, if we take the null hypothesis to be that that $P_B < P_A$, and the alternative hypothesis to be that $P_B > P_A$, then for virtually any hypothesis test one considers, if we fix the effect size, $\epsilon = |P_B - 0.5|$, increasing n will increase statistical power in the sense that it will increase the probability that the false null hypothesis is rejected in favour of the true alternative hypothesis. Of course, our agents are Bayesians, and so frequentist notions such as power are not strictly relevant—the agents update their beliefs in light of any and all evidence that they receive. Still, thinking of n as roughly tracking power is useful, here, because it gives an intuitive sense of why small n studies are generally better for the propagandist.

¹⁷ This figure is for $K = 10$ scientists in a complete network, each policymaker listening to 90% of scientists, and $p_B = 0.55$.

¹⁸ In an interesting connection, Holman and Bruner ([2015]) find that when a propagandist shares biased data with scientists, and when these scientists are able to ignore data that is statistically unlikely given their credences, scientists do better when n is large. This is for a different reason than we describe. When n is large in their models, it is easier to detect a propagandist and ignore them. If n is small, more results are statistically likely given some agent's credence.




Figure 4. Policymaker beliefs as influenced by the amount of data gathered by scientists.

unreasonable assumption given limits on the amount of time and energy policymakers or the public devote to science. In such cases, as the total number of scientists increases, there is no benefit to policymakers, since they continue to listen to k scientists. On the other hand, the increased number of scientists increases the total number of spurious results that the propagandist can then share with policymakers, often worsening policymaker belief.¹⁹ In the next section, we will engage in some further discussion of the effects that the size of the scientific community has on policymaker beliefs.

The final consideration that influences the success of the propagandist has to do with how much time it takes for the scientific community to converge to consensus. The convergence time depends on various factors. For instance, the size of a network generally influences how long it takes for all scientists to reach agreement. (In a sparsely connected network, more scientists means a longer time until each one is convinced.) In addition, as mentioned above, for fixed population size the cycle and complete network arrangements differ in that cycle networks take longer to converge. This is because scientists listen to fewer of their colleagues, meaning that disagreements will tend to persist as actors in different parts of the network gather different sorts of data.

Previous results from (Zollman [2007], [2010]) have shown that in many cases factors that slow convergence can be beneficial for an epistemic community. These factors lead to transient diversity of opinion, which is important for preventing scientists from settling too quickly on a theory that might be wrong. Heavily connected communities, on the other hand, are easily

¹⁹ This is not always the case since, as we outline in the appendix, the size of the community also influences the speed of convergence to scientific consensus, which can influence policymaker beliefs.

swayed by strings of misleading evidence that are widely shared.²⁰ But our results show that other considerations can complicate this issue. When industrial actors have a stake in a community's findings, taking a long time to settle on consensus can provide more opportunity for the propagandist to find and promote results that went the wrong way under conditions where spurious results are likely.

To try and better understand the relationship between time to convergence of the scientific network and the effectiveness of the propagandist, we derive semi-analytic results predicting expected outcomes in this model. In general, we can think of each policymaker's belief as influenced by two separate factors. One comes from the community of scientists, which, on average, will tend to pull credences in theory B up. The other comes from the propagandist and will tend to pull credence in theory B down. In other words, a policymaker's belief in round $t + 1$ of simulation on average can be represented by:

$$B_{(t+1)} = Y_P \times Y_S \times B_t,$$

where B_t is their prior belief, and Y_P and Y_S represent the average effects of the propagandist and scientists. In the appendix, we derive expressions for these quantities and show that it will always be the case that $Y_P < 1$ and $Y_S > 1$.

We can thus think of the dynamics in this model as a sort of tug of war between science and industry over public belief. In general, for any set of parameter values in our model, and some state of scientist or policymaker beliefs, one side will be winning because either $Y_P \times Y_S < 1$ or $Y_P \times Y_S > 1$.²¹ This means that when the propagandist is winning on average, they do better during long periods of scientific debate. In such cases, policymakers do better when scientists quickly reach consensus. Conversely, when the $Y_P \times Y_S > 1$ more evidence will tend to lead the policymakers towards the true belief—with the propagandist's main effect being to slow this progress.²²

For an example of the effects of time to convergence on belief, consider Figure 5, which shows the beliefs of policymakers consulting cycle and complete networks. We show results in a scenario that favours true belief— $p_B = 0.7$ and $n = 20$ —and also one that disfavors it— $p_B = 0.55$ and $n = 5$. In the first case, the slower cycle network gives policymakers time to learn the correct theory from scientists, while relatively unencumbered by the propagandist.

²⁰ Rosenstock *et al.* ([2017]) show that this will only occur for the region of parameter space where the problem of selecting the better theory is hard (due to there being fewer scientists, less difference between theories, or lower n).

²¹ The situation where the two influences exactly cancel is a knife's edge case.

²² One interesting observation is that the size of the scientific network, besides influencing time to convergence, also influences the rate at which policymakers change their beliefs. Adding more scientists, increases the total amount of data both coming directly from scientists and coming from the propagandist. When the scientists are winning the tug of war, this increase benefits the policymakers. When the propagandist is winning, this simply increases the rate at which policymakers decrease their credence in P_B .




Figure 5. Influence of scientific network structure on policymaker beliefs.

In the second case, the faster complete network gives the propagandist less time to share spurious data, and so is better for policymakers.

To be clear, if we look again at the simulation results we have presented, we can observe that, since policymakers, on average, start simulation at credence 0.25, any outcome where their final beliefs are <0.25 is one where the propagandist is pulling harder than the scientists, and vice versa. If we imagined that the process kept going—scientists kept performing action *B*, and the propagandist kept publicizing spurious results—we would expect policymakers to become completely convinced in either *A* or *B* eventually, with the relative values of Y_S and Y_P determining which one. But the community can end up at a better or worse state depending how quickly this happens. The most important take-away here is that there are reasons to think that a long argumentative process between scientists might actually be a bad thing when public belief is on the line. Thus we find a curious trade-off with the Zollman effect: at least in some cases, maximizing convergence time is truth conducive for the scientists, but not for the public.

In the next section, we will move on to discuss the second part of the tobacco strategy—where industry funds their own scientists and disseminates the resulting research. Before doing so, we mention a few potential concerns about the models just discussed, and robustness checks meant to alleviate those concerns. First, as noted above, data in our model can be ‘double counted’ in the sense that a policymaker might update their belief on data from a scientist and then update on the same evidence as it comes from the propagandist. We take this to be a feature rather than a bug, because it captures how the prevalence and repetition of evidence can increase its impact. Nonetheless, we also ran the same simulations where data could not be double counted this way. Results

were qualitatively similar, though the strength of the propagandist was diminished, particularly for policymakers connected to many scientists. We also varied the pattern with which policymakers connected to scientists to ensure that this did not have a significant effect on the results.

4 Biased Production

As mentioned in the introduction, another part of the tobacco strategy involved producing industry-funded data that could be used to counter the research coming from the scientific community.²³ Our second treatment concerns this aspect of the strategy. To model this, we now assume that the propagandist has their own resources that they can allocate towards gathering data about the uncertain theory B . After gathering this data, the propagandist then shares only those studies that spuriously indicated that theory A was the better one. Notice that we do not allow the propagandist to do ‘bad science’ in the sense of falsifying data, or to use biased methodologies. Their data gathering process is, in fact, identical to scientists attempting to find the truth. The only intervention we allow on the scientific process is one that is currently widely practiced (and not always considered unethical): selectively choosing not to publish some results.

We varied several new parameter values in this treatment. First, we varied the total amount of resources available to the propagandist. (Think of this as the total number of pulls from a bandit arm the propagandist can perform.) We also varied how the propagandist chose to divide up these resources into studies. Did they run many low-powered studies, or just a few high-powered ones?²⁴

In this scenario, holding other parameters fixed, more scientists are unequivocally beneficial to policymakers’ beliefs. Scientists no longer unintentionally provide grist for the propagandist, and, instead, produce and publicize data that is more likely to point in the right direction. It is also consistently better, as in the last section, for policymakers to consult more scientists. The more data they receive from epistemically motivated actors, the

²³ There is a close connection here to the work of Holman and Bruner ([2015]), who consider biased agents in a scientific community. Like our ‘biased production’ propagandist, their ‘intransigently biased agent’ (IBA) produces and shares evidence that tends to favour the wrong theory. However, in their case the IBA draws from a biased arm relative to the other scientists in the community, whereas, as we describe, our propagandist performs the same actions as the scientists but selectively publishes results. In addition, they focus on the effect an IBA might have on beliefs within a scientific community, showing how industry can shape scientific consensus itself, rather than public belief in particular.

²⁴ We ran simulations for a subset of the following parameter values: population size = 4, 6, 10, 20, number of draws $n = 1, 2, 10, 25, 50, 100$, success of better theory $p_B = 0.55, 0.6, 0.7$, and number of scientists consulted, k , ranged from one to all of them. Propagandist total resources varied from equal to that of one scientist, to equal to those of the whole scientific community. Propagandist n ranged between 1 and their total possible pulls.




Figure 6. Propagandists with more resources are better able to mislead policymakers.

less the impact of the propagandists' spurious studies on their beliefs. And, again, the better the successful theory, the lower the chance that the propagandist generates enough spurious results to influence policymakers.

Unsurprisingly, more resources are better for the propagandist. If they have the wherewithal to run more studies, they have a greater chance of generating results supporting theory *A*. Suppose each scientist has the resources to collect ten data points per trial. Assume that the propagandist runs trials of the same size, but can run a different number of these depending on their level of resources. As Figure 6 shows, as this number increases, the likelihood that the policymakers believe the correct theory goes down.²⁵

Suppose that we hold fixed the resources a propagandist has at their disposal. The last question we ask is: given this, what is the best way for the propagandist to allocate these resources to different studies? The answer is that it is always better for them to break their data up into the smallest possible chunks. Suppose the propagandist has resources to collect 1000 data points. In the most extreme case, we could imagine a scenario where the propagandist runs 1000 studies with one draw each. They could then report only those where action *B* failed (on average this would be $(0.5 - \epsilon) * 1000$ data points), making it look to policymakers as if action *B* was never successful in these experiments. At the other extreme, the propagandist would run a single experiment with 1000 data points. The chances that this experiment would support the better theory *B* are quite high, leaving the propagandist

²⁵ These results are for ten scientists, $n = 10$ for scientists, $p_B = 0.55$, and policymakers listening to half of scientists.




Figure 7. Propagandists who use their resources to perform small studies are better at misleading policymakers.

with nothing to disseminate. In general, the propagandist should choose the smallest n they can get away with.

Figure 7 shows this effect. For propagandists with resources to gather fifty data points, the smaller the studies they run, the worse the eventual beliefs of the policymakers. If propagandists run a few, good studies, they are unable to mislead anyone.²⁶

This observation actually has an analogue in the previous set of results on propagandists who selectively share others' data rather than gathering their own. In that model, given a set of available resources for a community, it is better to use these to run fewer, better studies than to run many less well-powered ones. We can observe this by comparing how policymakers do given communities of scientists that have a finite amount of resources distributed differently. Figure 8 demonstrates this. The scientific network, in this figure, has the resources to gather a total of 100 data points each round. This is divided up by the number of scientists involved in exploring the phenomenon so that, for example, ten scientists gather ten data points each or two scientists gather fifty data points each. Without a propagandist policymakers develop successful beliefs regardless of how resources are allocated. As the individual studies get smaller, though, the propagandist is increasingly able to take advantage of the stochasticity involved in gathering data to forward their agenda. Small studies mean confused policymakers.

²⁶ These results are for ten scientists, $n = 10$ (for scientists), policymakers listening to all scientists, and $p_B = 0.55$.




Figure 8. Policymakers do better when fixed resources are allocated to a few, good studies rather than many with small samples.

5 Journalists as Unwitting Propagandists

In the two treatments we have discussed so far, the propagandist has adopted strategies intended to influence the policymakers' beliefs. In both cases, we found that the propagandist can succeed, at least under some circumstances. The basic mechanism at work in both strategies is that the propagandist endeavours to bias the total body of evidence that the policymakers see. They do not do this through fraud, or even by sampling from a different distribution than the other scientists (as in Holman and Bruner [2015]). Instead, they do it by choosing results so that the total evidence available, even though it was all drawn from a fixed distribution, does not reflect the statistics of that distribution. What these results show is that in many cases, particularly when evidence is relatively difficult to get or equivocal, it is essential that policymakers (and the public) have access to a complete, unbiased sampling of studies. This suggests that curation, or even partial reporting, of evidence can lead to epistemically poor results.

But propagandists are not the only agents who regularly curate scientific results. Non-malicious actors can also play this role. For instance, journalists often select scientific studies to feature in articles. Few journalists are experts in the relevant science, and in any case, the goal is rarely to give a uniform and even-handed characterization of all available evidence. To the contrary, journalistic practices often encourage, or even require, journalists to share either only the most striking, surprising, or novel studies; or else, when they choose to share a study on a controversial topic, to endeavour to be 'fair' by sharing research that supports both sides. Indeed, the US Federal Communications Commission (FCC) had an official policy, from 1949 to 1987, that required

broadcast news to present controversial topics of public interest in a ‘balanced’ manner. In practice, this meant giving roughly equal time to multiple perspectives. The policy was known as the ‘fairness doctrine’. Oreskes and Conway ([2010]) argue that this doctrine became a tool of the tobacco industry’s propagandists, who used it to insist that pro-industry science get publicized.

To study the effects of the fairness doctrine, we consider the following modification of the models we have considered thus far.²⁷ The network structures remain the same, but in place of the propagandist, we have a ‘journalist’, who like the propagandist can see all of the scientists’ results and can communicate with all of the policymakers. In this case, we suppose that the policymakers receive all of their information from the journalist, and consider different ways the journalist might share that data.

Suppose, for instance, that the journalist—abiding by the fairness doctrine—shares two results from the network of scientists each round, one of which supports theory *B*, and one of which supports theory *A*. More precisely: we suppose that each round, the journalist chooses randomly among all of the results that support theory *A* and all that support theory *B*. In rounds where no results support one of the two theories, the journalist shares the result they shared the previous round (capturing the idea that sometimes journalists need to search past data to find something to adequately represent the ‘other side’), unless there has never been a result favouring a given theory, in which case the journalist shares only one result. What we find is that this addition artificially slows the policymakers’ acceptance of theory *B*.

To show this, it is most informative to compare the policymakers’ beliefs under the action of this journalist (fair) with what one would find if (a) the journalist shared two randomly selected pieces of data each round (random) and (b) the journalist simply shared all available evidence each round (all). We find that (on average across runs) the fairness doctrine never improves the epistemic performance of the policymakers, and for many parameter values, policymakers receiving evidence from a journalist abiding by the fairness doctrine end up with beliefs substantially less accurate than under conditions (a) or (b). In all cases we consider, the policymakers’ beliefs still tend towards the truth (whenever the scientists’ beliefs do), but they do so much more slowly.²⁸ In this sense, then, the ‘fair’ journalist mimics the propagandist under the

²⁷ We also considered a version of the model in which the journalist only published results that were sufficiently ‘surprising’ in the sense that the journalist judged them to have had a low probability of occurring. For some parameter values, this sort of journalistic practice led policymakers to accept the wrong theory. But we deemed the results to be ambiguous, because they depended sensitively on the value of *n* chosen.

²⁸ When policymakers see one result supporting *B* and one supporting *A* each round, on average the result supporting *B* is still stronger since this is in fact the better arm. So they learn the truth eventually.




Figure 9. ‘Fair’ reporting slows policymakers from reaching accurate beliefs.

conditions where they are losing the tug of war for policymakers’ beliefs, but nonetheless slow the policymakers’ progress.

Figure 9 shows average policymaker beliefs when scientists reach consensus on the true belief for complete networks under different parameter values.²⁹ Data points are for simulations where journalists share fair, random, or all data as described above. The trend line shows the average beliefs across parameters under each condition. As is evident, when journalists share all the data, policymaker beliefs are very accurate. With less data available—two random sets of data—their convergence to true beliefs is slowed. Under fair reporting, policymakers converge to true beliefs much more slowly, since they see a result that supports theory *A* every time they see a result supporting *B*. The strength of this effect is dependent on parameter values, but, as mentioned, fair reporting only worsens, and never improves, policymaker beliefs in comparison to reporting randomly or reporting all data.

Journalists are not the only actors who perform data curation. University faculty, when teaching courses, may endeavour to give a panoptic and even-handed view of a field, but they will not be able to share all available relevant data with a class. Literature reviews and survey articles also curate data, commenting on its quality and selecting what seems most relevant or significant. The mechanisms by which researchers and university instructors select material to present are arguably motivated by epistemic considerations—they

²⁹ We chose a representative set of data points including all combinations of $K=6, 10, 20$, $p_B = 0.55, 0.6, 0.7$, and $n=5, 10, 20$. As in the previous models, initial policymaker beliefs were sceptical. Complete networks demonstrated the effect more strongly because cycle networks are slow to converge and thus there is more time for policymakers to learn true beliefs even under ‘fair’ reporting. But even with cycle networks the effects of fair reporting were always neutral or negative.

seek to give a compact overview that supports the views they believe are best supported, given the total evidence that they have seen—but it is hard to see how they can do this without using their own beliefs—which may be influenced by various factors—to guide the choices. And so, it is not clear that even this sort of curation is certain to accurately track the statistics of the total available evidence. In this way, even independent researchers, in their teaching and writing, can unwittingly act as propagandists. In addition, it has been widely observed that scientists regularly fail to publish negative or null results. This ‘file drawer’ effect biases the pool of evidence available to the entire community, and has already been shown to potentially have negative effects on scientific beliefs (Romero [2016]).³⁰

6 Conclusion

We have now considered two aspects of the tobacco strategy as described by Oreskes and Conway. We take the results we have presented to provide strong evidence that the tobacco strategy can work, in the sense that biasing the total evidence presented to a group of agents can lead those agents to converge to the incorrect belief. Of course, modelling alone cannot establish that this did happen in the case of tobacco industry propaganda efforts. But when you combine the modelling results we have presented with the careful historical analysis and arguments offered by Oreskes and Conway ([2010]), there are strong reasons to believe that the tobacco strategy was an important causal factor in the long delay between when a link between cigarettes and cancer was established by the medical research community, and large-scale changes in regulation and public attitudes. Moreover, we take it that if these strategies have succeed in the past, they are very likely to succeed now and in the future.³¹

What is perhaps most interesting about the results we have presented is not that they show what can work, but rather the insight they provide into how those strategies work. As we have emphasized above, in our discussion on unwitting propagandists, the basic mechanism at play in our models is that the propagandist biases the total evidence on which the policymakers update their

³⁰ See also (Ioannidis [2005]).

³¹ Returning once more to the themes of Footnotes 4 and 8: we emphasize that, at very least, these models provide how-possibly information: the processes we describe with these models can possibly occur. Furthermore, as we argue here, the theoretical and historical knowledge developed by Oreskes and Conway increases our confidence that these processes have been important ones in the real world. As noted, we cannot be sure that this is the case. But we do not take this to be a special problem for models of this sort. With almost any sort of scientific exploration, including experimental and other empirical work, conclusions are drawn from data via theoretically supported inference. (For more on this point, see Longino and Doell [1983].) In this case, we draw (conservative) inferences based on a combination of modelling and historical work.

beliefs. The fact that each result that is shared is, in some sense, ‘real’ ends up being irrelevant, because it is the statistical properties of the total available evidence that matter. For this reason, we see that producing results favourable to industry—say, through industry funding—is not necessary to manufacture favourable public beliefs, at least in cases where the evidence is probabilistic and spurious results are possible.

On this point: one might have expected that actually producing biased science would have a stronger influence on public opinion than merely sharing others’ results. But when one compares the two treatments we consider above, there are strong senses in which the less invasive, more subtle strategy of selective sharing is more effective than biased production, all things considered, particularly when the scientific community is large and the problem at hand is difficult (or the power of generic experiments is low). The reason is that such a scientific community will produce, on its own, plenty of research that, taken without proper context, lends credence to falsehoods. Merely sharing this already-available evidence is cost-effective, and much less risky than producing one’s own research—which, after all, can cost a great deal to produce, fail to generate the desired results, and ultimately be disqualified or ignored because of the association with industry. From this point of view, producing industry science (or even outright fraud) is simply not worth it. In many cases, another more effective, less expensive, and more difficult to detect strategy is available.

Indeed, if industrial interests do wish to spend their money on science, perhaps rather than producing their own results and publishing them selectively, they would do better to adopt a different strategy, which Holman and Bruner ([2017]) have called ‘industrial selection’. The idea behind industrial selection is that there are many experimental protocols, methodologies, and even research questions that scientists may adopt, often for their own reasons. Some protocols may tend to produce more industry-friendly results than others.³² Industrial selection involves identifying methods already present among the community of scientists that tend to favour one’s preferred outcome, and then funding scientists who already use those methods. This extra funding works to amplify these scientists’ results, by allowing them to publish more, perform higher powered studies, and train more students who will go on to produce consonant results. Indeed, there is a strong potential feedback effect between industrial selection and selective sharing: by increasing the number of ambient industry-favourable results in a scientific literature, and

³² Holman and Bruner ([2017]) consider, for instance, the case of a class of drugs intended to stop heart failure (Moore [1995]). In that case, studies that used reducing arrhythmia as a proxy for preventing heart failure found that the drugs were highly successful, whereas studies that directly measured death rates from heart failure later showed that the drugs were substantially increasing the likelihood of death. Industry swayed scientific practice by funding the former type of study, but not the latter.

then further amplifying those results by sharing them selectively, propagandists can have an even stronger effect.

It is worth emphasizing that industrial selection does not require some scientists to have adopted bad or unreliable research methods; it may also just be that they have focused on research topics that promote an industrial agenda. For instance, Oreskes and Conway ([2010], p. 17) describe how the Tobacco Industry Research Council supported the work of noted asbestos researcher Wilhelm C. Hueper of the National Cancer Institute. Hueper's research was not methodologically flawed; to the contrary, it very convincingly established the (real) link between lung cancer and asbestos. But for the tobacco industry, it was valuable to have evidence that other environmental factors played a role in lung cancer, since that could be used to obfuscate the role their own product played. Taken together, then, selective sharing and industrial selection suggest those interested in dampening industry influence should turn their attention to the more subtle ways propagandists can come to bias epistemic communities. This will to a certain extent require philosophers of science to reorient their focus, as extant work tends to address explicit, direct ways industry can impact inquiry.

It is natural to ask given the past success of the tobacco strategy and the results we have described, whether there are ways policymakers or members of the public can protect themselves against a propagandist using the methods described. Of course, one natural suggestion is to stop listening to propagandists. But this is difficult to do, in large part because propagandists pose as (or, in some cases, are) credentialed scientists.³³ Another possibility, perhaps more realistic, is to gather evidence from as many scientists as possible: as we discuss above, policymakers connected to more scientists tend to do better, because the propagandist has to produce more evidence to effectively influence such policymakers' beliefs.³⁴ Connecting to many scientists may also allow policymakers to identify propagandists on statistical grounds: indeed, as Holman and Bruner ([2015]) point out, in some cases it may be possible to identify biased agents by comparing the statistics of their results with those of other agents in their network. Finally, and along similar lines, it might also be possible, at least in some circumstances, for policymakers to adjust their updating rule to accommodate the possibility that some of the evidence they receive is exaggerated, filtered, or otherwise intended to mislead.³⁵ On the

³³ Indeed, perhaps the most striking message of (Oreskes and Conway [2010]) is that the most effective propagandists have been prestigious scientists, often coming from different fields, with a political axe to grind.

³⁴ We note that connecting to more scientists was particularly effective in the selective sharing treatment when double-counting was eliminated, suggesting that policymakers would do well to attend to cases where a small number of studies are used repeatedly to prove a point, when other studies are also available.

³⁵ We are grateful to Manheim ([2018]) for emphasizing this latter possibility, which is also quite close to what Holman and Bruner ([2015]) propose.

other hand, as O'Connor and Weatherall ([2018]) show, some apparently reasonable heuristics for updating belief in light of unreliable sources can lead to bad epistemic outcomes, including stable polarization. We think that it would be valuable, in future work, to study how these latter two reactions might be implemented, and how and when they would be effective.

We will conclude by noting a few ways the incentive structure of science contributes to the success of the would-be propagandist, particularly when employing the selective sharing strategy. As we noted above, selective sharing is most effective when (i) the problem is difficult (in the sense that p_B is close to 0.5); (ii) the number of data points per scientist per round is small; (iii) there are many scientists working on a given problem; and (iv) policymakers are connected to few scientists. This situation interacts with how scientists are rewarded for their work in complex, troubling ways.³⁶ For instance, although there are some incentives for scientists to make their work accessible and to try to communicate it with the public and policymakers, it continues to be the case that academic scientists are evaluated on the basis of their research productivity—and there are sometimes reputational hazards associated with ‘popularizing’ too much. This sort of effect can limit the number of scientists to whom policymakers are connected.

But a more troubling way that propaganda and incentives intersect concerns the role of low-powered studies. Scientists are strongly encouraged to publish research articles, and to do so often. Moreover, publishing positive results is strongly preferred to null results. Finally, scientists work with limited funds that they must use to meet many demands. All of these considerations lead to strong incentives to run studies with relatively low power (particularly, fewer data points) whenever possible (Smaldino and McElreath [2016]).³⁷ Of course, most scientific fields recognize that, all else being equal, higher powered studies are more rigorous, more reliable, and make for all-around better science. But there is little incentive, and often few resources, for individual scientists to produce consistently high-quality studies. And as we have seen, low-powered studies are more likely to spuriously support erroneous conclusions. In cases where industry is not involved, this is not as worrying: spurious results can be detected via meta-analysis. But in the presence of a propagandist, they become a powerful tool in shaping public belief.

³⁶ Previous authors have assumed that scientists are motivated by a quest for credit, in the same way normal people seek money or goods. ‘Credit economy’ models have been used, for example, to show why scientists might divide labour (Kitcher [1990]) or commit fraud (Bright [2017]).

³⁷ The ubiquity of low-powered studies in some fields—such as neuroscience (Button *et al.* [2013]; Szucs and Ioannidis [2017]), ecology (Lemoine *et al.* [2016]), and evolutionary biology (Hersch and Phillips [2004]), for instance—has been remarked upon and widely discussed within those fields. Smaldino and McElreath ([2016]) perform a meta-analysis of work in the behavioural sciences over the last sixty years and argue that despite repeated calls to increase power, little change has occurred. In addition, they show why scientific communities might move towards producing low quality, low effort studies.

Discussions of low-powered science often focus on their role in controversies about particular results, in the context of the so-called replication crisis in the behavioural sciences (and especially psychology) (Open Science Collaboration [2015]). And the regular publication and acceptance of ultimately incorrect results is certainly worrying. But the arguments and results we have given here suggest that even more is at stake: the same sorts of conditions that lead to a replication crisis also provide extra fodder for industrial propaganda. This makes reshaping scientists' incentives to publish many, low-powered 'least-publishable units' crucial to foiling propagandists. For instance, scientists should be granted more credit for statistically stronger results—even in cases where they happen to be null. Scientists should not be judged on the number of publications they produce, but on the quality of the research in those articles, so that a single publication with a very high-powered study, even if it shows a null result, should be 'worth' more than many articles showing surprising and sexy results, but with low power. Similar reasoning suggests that given some fixed financial resources, funding bodies should allocate those resources to a few very high-powered studies, rather than splitting it up into many small grants with only enough funding to run lower-powered studies. Or a more democratic alternative might involve scientists who work together to generate and combine smaller studies before publishing, making it more difficult for the propagandist to break off individual, spurious results.

Appendix: Semi-analytic results

Here we describe a semi-analytic treatment of the expected progress of the models presented in Section 3.

Consider a population of K scientists each performing n trials each round, solving a problem with $p_B = 0.5 + \epsilon$. Consider a policymaker connected to k scientists, with credence B_t at time t in the true proposition that arm B is advantageous. This observer will update her beliefs by conditionalization in light of two sources of evidence, from the scientists she is connected to and from the propagandist. Thus, at the end of the round, we will have:

$$B_{(t+1)} = Y_P \times Y_S \times B_t, \quad (1)$$

where Y_S represents the contribution from scientists and Y_P represents the contribution from the propagandist. What are these contributions on average?

In a generic round of play, the average result for each scientist performing the informative action should be $E = n * p_B$, that is, the expectation value for n draws from a binomial distribution with probability p_B . We estimate Y_S by supposing that the policymaker receives, from each scientist performing action B that they are connected to, the average (that is, expected) result, and update. This yields:

$$Y_S = \left(\frac{P(n * (0.5 + \epsilon) | H)}{P(n * (0.5 + \epsilon))} \right)^{r * k}, \tag{2}$$

where $r \leq 1$ is the fraction of scientists in the population performing the informative action, so $r * k \leq k$ approximates the number of scientists connected to the observer who are performing the informative action. Here, H is the policymaker’s hypothesis about the success rate of arm B . When multiplied by the current belief, B_t , as in Equation (1), this is just Bayes rule for the expected evidence. Notice that this equation does not account for the fact that B will change as an agent updates on different sets of evidence within a time step. We should think of this approximation as most useful for cases where B is changing slowly.

One can show by direct computation that Y_S reduces to:

$$Y_S = \left(\frac{1}{B_t + (1 - B_t)X} \right)^{r * k}, \tag{3}$$

where

$$X = \left(\frac{1 - 2\epsilon}{1 + 2\epsilon} \right)^{2n\epsilon}.$$

Observe that $X < 1$, and so for any value of B_t , $Y_S > 1$. It follows that updating on evidence from the scientists will, on average, increase credence that arm B is preferable. Moreover, one can see that as n and ϵ each increase, X will decrease, and thus Y_S will increase. Likewise, Y_S will increase as r , the number of scientists testing the preferred arm, increases, and as k , the number of scientists a policymaker listens to, increases. Of course, in the actual model stochastic effects can defeat this; we are describing average behaviour.

Note that as B_t decreases, Y_S also becomes larger, so as observers become more sceptical, evidence from scientists will become more powerful. Importantly this means that a set of parameter values alone will not determine the average behaviour of the model: the belief state of a policymaker B_t and the belief states of the scientists, as represented by r , will influence which way policymaker beliefs trend.

Now consider Y_P . Here we consider all (and only) spurious results produced by the population of scientists each round. First, observe that we can calculate a rate at which these will occur if all scientists perform the informative action, given by $z = P(x < n/2) \times K$, where x is the number of times arm B pays off. This expression is just the probability of getting a spurious result—one that suggests $P_B < .5$ —times the number of scientists in the population. For instance, for even n , we have:

$$P(x < n/2) = \sum_{i=0}^{n/2-1} \frac{n!}{i!(n-i)!} (P_B)^i (1 - P_B)^{n-i}. \tag{4}$$

This is the sum of the probabilities of drawing each possible $x < n/2$ from a binomial distribution with n draws and P_B .

We can also calculate an expectation value for such spurious draws, \tilde{E} , which is given by (again for even n):

$$\tilde{E} = \frac{1}{P(x < n/2)} \sum_{i=0}^{n/2-1} \frac{in!}{i!(n-i)!} (P_B)^i (1 - P_B)^{n-i}. \tag{5}$$

This equation multiplies the value of each possible spurious draw, i , by the probability of that draw, conditional on the assumption that the draw will be spurious.

This leads to an estimate for Y_P of:

$$Y_P = \left(\frac{1}{B_t + (1 - B_t)\tilde{X}} \right)^{r * z}, \tag{6}$$

where now

$$\tilde{X} = \left(\frac{1 - 2\epsilon}{1 + 2\epsilon} \right)^{2\tilde{E} - n}.$$

Once again, r is the fraction of the population of scientists performing the informative action, while z , remember, is the rate of spurious results. Observe that, by definition $\tilde{E} < n/2$, and so whatever else is the case, $(2\tilde{E} - n) < 1$, so $\tilde{X} > 1$, which means Y_P is always less than one. Hence the propagandist will always tend to decrease the observer’s confidence in the true proposition, and this effect will be stronger when more scientists take the informative action and the rate of spurious results is higher (more on this shortly).

Beyond these observations, it is harder to analyse the dependence of Y_P on other parameters, on account of the complexity of the function \tilde{E} . An analytic simplification of the expression for \tilde{E} is apparently not available, but a numerical analysis reveals that for small ϵ , \tilde{E} is very well approximated for sufficiently large n by a linear function in n , of the form $\tilde{E} = -0.62 + 0.47n$. (This expression is determined by taking a linear best fit in Mathematica.) Thus we have:

$$\tilde{X} = \left(\frac{1 - 2\epsilon}{1 + 2\epsilon} \right)^{-0.06n - 1.24}. \tag{7}$$

We see that as ϵ becomes larger, \tilde{X} becomes smaller (that is, approaches one), and thus Y_P approaches one, so that the propagandist is less effective. Notice though that as n increases in this equation, \tilde{X} becomes larger, thus Y_P gets

smaller, so the propagandist gets more effective. Thus n influences Y_P in two, countervailing ways. A large n makes z , the probability of a spurious result, smaller. On the other hand, when spurious results do occur, they are more convincing for large n . In addition, \tilde{X} tends to grow slowly in n and whether the propagandist ends up swaying the observer depends on the relative scaling of Y_P and Y_S with n , which in turn depends on ϵ . We also have that, as B_i approaches one, Y_P approaches one, so that the propagandist becomes less effective. (Thus, both the propagandist and the scientists are more effective when policymakers have low credence in arm B .)

A final consideration concerns z , the rate of spurious results. Clearly as z increases, the propagandist becomes more effective. This will happen as K , the number of scientists, increases. (Though if policymakers listen to all the scientists, k will increase as K does, strengthening Y_S as well.) The rate z will also increase as r , the fraction of scientists performing the informative action, increases, though an increase in r will increase both Y_S and Y_P . There is also a dependence in z on $P(x < n/2)$. It is difficult to get a good first order approximation for this function, but by studying plots one can characterize some qualitative features. In general, for fixed n , increasing ϵ decreases $P(x < n/2)$, and thus decreases z , making the propagandist less successful.

For sufficiently large n , we find that as n increases, $P(x < n/2)$ decreases, again making the propagandist less effective. This is what one would expect: the probability of spurious results decreases as n increases. But what is surprising is that for small n , there is a regime in which increasing n actually increases $P(x < n/2)$. For sufficiently large ϵ (say, >0.1 or so), this regime is a small part of n -space. But for small ϵ (say, $\epsilon \approx 0.001$), $P(x < n/2)$ actually increases with n for all values of n that we study. So increasing n under these circumstances makes the propagandist more effective. Again, an increase in n also makes scientists more effective, so the overall influence of n will depend on whether it increases Y_P or Y_S to a greater degree.

Acknowledgements

This material is based upon work supported by the National Science Foundation under grant number 1535139 (O'Connor). The majority of the work on this project was completed while O'Connor and Weatherall were Visiting Fellows at the Australian National University; we are grateful to the Research School for Social Science and the School of Philosophy for support. We are also grateful to audiences at the ANU, University of Groningen, UC Irvine, UC Merced, UC San Diego, and the 2018 PPE Society Meeting for helpful feedback on talks related to this work, and to

Jeff Barrett, Craig Callender, Bennett Holman, David Manheim, Aydin Mohseni, Mike Schneider, and Chris Smeenk for helpful discussions.

James Owen Weatherall
Department of Logic and Philosophy of Science
University of California, Irvine
 Irvine, CA, USA
 weatherj@uci.edu

Cailin O'Connor
Department of Logic and Philosophy of Science
University of California, Irvine
 Irvine, CA, USA
 cailino@uci.edu

Justin P. Bruner
Department of Theoretical Philosophy
University of Groningen
 Groningen, The Netherlands
 j.p.bruner@rug.nl

References

- Bala, V. and Goyal, S. [1998]: 'Learning from Neighbors', *Review of Economic Studies*, **65**, pp. 595–621.
- Bright, L. K. [2017]: 'On Fraud', *Philosophical Studies*, **174**, pp. 291–310.
- Button, K. S., Ioannidis, J. P. A., Mokrysz, C., Nosek, B. A., Flint, J., Robinson, E. S. J. and Munafò, M. R. [2013]: 'Power Failure: Why Small Sample Size Undermines the Reliability of Neuroscience', *Nature Reviews Neuroscience*, **14**, pp. 365–76.
- Douglas, H. [2009]: *Science, Policy, and the Value-Free Ideal*, Pittsburgh, PA: University of Pittsburgh Press.
- Elliott, K. C. [2011]: *Is a Little Pollution Good for You?* Oxford: Oxford University Press.
- Frey, D. and Šešelja, D. [2018]: 'What Is the Epistemic Function of Highly Idealized Agent-Based Models of Scientific Inquiry?', *Philosophy of the Social Sciences*, **48**, pp. 407–33.
- Hersch, E. I. and Phillips, P. C. [2004]: 'Power and Potential Bias in Field Studies of Natural Selection', *Evolution*, **58**, pp. 479–85.
- Holman, B. and Bruner, J. P. [2015]: 'The Problem of Intransigently Biased Agents', *Philosophy of Science*, **82**, pp. 956–68.
- Holman, B. and Bruner, J. [2017]: 'Experimentation by Industrial Selection', *Philosophy of Science*, **84**, pp. 1008–19.
- Ioannidis, J. P. [2005]: 'Why Most Published Research Findings Are False', *PLoS Medicine*, **2**, p. e124.

- Kitcher, P. [1990]: 'The Division of Cognitive Labor', *The Journal of Philosophy*, **87**, pp. 5–22.
- Kitcher, P. [2003]: *Science, Truth, and Democracy*, New York: Oxford University Press.
- Kitcher, P. [2011]: *Science in a Democratic Society*, Amherst: Prometheus Books.
- Lemoine, N. P., Hoffman, A., Felton, A. J., Baur, L., Chaves, F., Gray, J., Yu, Q. and Smith, M. D. [2016]: 'Underappreciated Problems of Low Replication in Ecological Field Studies', *Ecology*, **97**, pp. 2554–61.
- Longino, H. and Doell, R. [1983]: 'Body, Bias, and Behavior: A Comparative Analysis of Reasoning in Two Areas of Biological Science', *Signs*, **9**, pp. 206–27.
- Longino, H. E. [1990]: *Science as Social Knowledge: Values and Objectivity in Scientific Inquiry*, Princeton, NJ: Princeton University Press.
- Longino, H. E. [2002]: *The Fate of Knowledge*, Princeton, NJ: Princeton University Press.
- Manheim, D. [2018]: 'Comment on "How to Beat Science and Influence People"', available at <medium.com/@davidmanheim/comment-on-how-to-beat-science-and-influence-people-https-arxiv-org-abs-1801-01239-5e96f0ad49bd>
- Michaels, D. [2008]: *Doubt Is Their Product*, Oxford: Oxford University Press.
- Michaels, D. and Monforton, C. [2005]: 'Manufacturing Uncertainty: Contested Science and the Protection of the Public's Health and Environment', *American Journal of Public Health*, **95**, pp. S39–48.
- Moore, T. J. [1995]: *Deadly Medicine: Why Tens of Thousands of Heart Patients Died in America's Worst Drug Disaster*, New York: Simon and Schuster.
- Norr, R. [1952]: 'Cancer by the Carton', *Reader's Digest*, **61**, pp. 7–8.
- O'Connor, C. and Weatherall, J. O. [2018]: 'Scientific Polarization', *European Journal for Philosophy of Science*, **8**, pp. 855–75.
- O'Connor, C. and Weatherall, J. O. [2019]: *The Misinformation Age: How False Beliefs Spread*, New Haven, CT: Yale University Press.
- Open Science Collaboration. [2015]: 'Estimating the Reproducibility of Psychological Science', *Science*, **349**, p. aac4716.
- Oreskes, N. and Conway, E. M. [2010]: *Merchants of Doubt*, New York: Bloomsbury Press.
- Romero, F. [2016]: 'Can the Behavioral Sciences Self-correct? A Social Epistemic Study', *Studies in History and Philosophy of Science Part A*, **60**, pp. 55–69.
- Rosenstock, S., Bruner, J. and O'Connor, C. [2017]: 'In Epistemic Networks, Is Less Really More?', *Philosophy of Science*, **84**, pp. 234–52.
- Smaldino, P. E. and McElreath, R. [2016]: 'The Natural Selection of Bad Science', *Open Science*, **3**, p. 160384.
- Szucs, D. and Ioannidis, J. P. A. [2017]: 'Empirical Assessment of Published Effect Sizes and Power in the Recent Cognitive Neuroscience and Psychology Literature', *PLoS Biology*, **15**, p. e2000797.
- Weatherall, J. O. and O'Connor, C. [unpublished]: 'Conformity in Scientific Networks', available at <arxiv.org/abs/1803.09905>.
- Woodward, J. [2003]: *Making Things Happen: A Theory of Causal Explanation*, Oxford: Oxford University Press.

- Wynder, E. L., Graham, E. A. and Croninger, A. B. [1953]: 'Experimental Production of Carcinoma with Cigarette Tar', *Cancer Research*, **13**, pp. 855–64.
- Zollman, K. J. [2007]: 'The Communication Structure of Epistemic Communities', *Philosophy of Science*, **74**, pp. 574–87.
- Zollman, K. J. [2010]: 'The Epistemic Benefit of Transient Diversity', *Erkenntnis*, **72**, pp. 17–35.