

# Experimentation by Industrial Selection

Bennett Holman and Justin Bruner\*†

---

Industry is a major source of funding for scientific research. There is also a growing concern for how it corrupts researchers faced with conflicts of interest. As such, the debate has focused on whether researchers have maintained their integrity. In this article we draw on both the history of medicine and formal modeling to argue that given methodological diversity and a merit-based system, industry funding can bias a community without corrupting any particular individual. We close by considering a policy solution (i.e., independent funding) that may seem to promote unbiased inquiry but that actually exacerbates the problem without additional restrictions.

---

**1. Introduction.** Industry funding has become the major source of income for scientific research. In addition to economic forces causing this trend, a number of nonprofit and government agencies (e.g., the Gates Foundation, the Wellcome trust, the President’s Council of Advisors on Science and Technology, the Food and Drug Administration) are recommending more collaboration between industry and academia (Drazen 2015). The entire field of biomedicine is undergoing a massive shift along these lines as pharmaceutical companies shift their research and development work into universities under the rubric of translational medicine (Robinson, forthcoming). Similarly, most of the work assessing the safety of industrial chemicals is conducted by the companies that produce them (Conrad and Becker 2011). Yet concurrent with this shift toward industry funding, there has been growing concern with the ways in which industry “bends science” (McGarity and Wagner 2008).

\*To contact the authors, please write to: Bennett Holman, Underwood International College, Yonsei University; e-mail: bholman@yonsei.ac.kr. Justin Bruner, School of Politics and International Relations, Australian National University; e-mail: justin.bruner@anu.edu.au.

†We would like to thank the audiences at Values in Medicine, Science, and Technology 2015; the Munich Center for Mathematical Philosophy; The Science of Evolution and the Evolution of the Sciences; International Conference on Applied Ethics; Rationality, Agency, and Responsibility in Social Groups; and the biennial meeting of the PSA 2016.

Philosophy of Science, 84 (December 2017) pp. 1008–1019. 0031-8248/2017/8405-0017\$10.00  
Copyright 2017 by the Philosophy of Science Association. All rights reserved.

The focus of this concern has been with the conflicts of interest created by industry partnership and how to manage them (Lo and Field 2009; Elliott 2014). Such concerns have led some to suggest that industry-funded research should be assigned less weight (Resnik and Elliott 2013; Holman and Bruner 2015). Others find this antithetical to the very nature of science, in which facts are assessed independent of the personal qualities of the authors (Conrad and Becker 2011). Moreover, industry scientists react very strongly to the idea that their work is suspect. For example, Brain, Stavely, and Ortego (2016, 1071) note that while society is distrustful of research produced by researchers affiliated with industry, “this is a peculiar and seemingly illogical behavior given that employees in government, industry and academia are all sourced from the same educational system, from the same academic institutions, and often from the same laboratories, mentors and advisors, which instill commensurate scientific training, education, experience, integrity, morals, and ethics. Consequently, an insinuation of affiliation bias categorically ignores educational equality and job opportunity/availability, and is therefore completely irrational.” Whether or not concern with affiliation bias is irrational, one shared assumption is that conflicts of interest operate by causing researchers to bias their results in favor of a secondary interest (e.g., corporate profits) at the expense of a primary interest (e.g., truth). The insinuation of those distrustful of industry is that researchers with industry ties have either willfully compromised their integrity or been unwittingly corrupted. It is no wonder why industry researchers bristle at such distrust.

While not disregarding problems caused by conflicts of interest, the pursuit offered by social epistemology can highlight entirely distinct mechanisms that produce industry-favorable outcomes without any individual researcher being corrupted by industry. In order to highlight such effects, we first consider the antiarrhythmic drug disaster. We argue that while conflicts of interest were present, they were not the primary cause of the disaster. Instead, the tragedy resulted from the following features of the scientific community: (1) an initial diversity of views, (2) a merit-based structure, and (3) the ability of industry to selectively distribute resources.

Without a countervailing force, we argue that these features are sufficient to tip an otherwise honest epistemic community toward pro-industry conclusions. This effect occurs not by altering the view of any particular individual but as a result of the selective pressures that naturally exist in the scientific community. In section 2, we formalize this intuition using a formal model of social inquiry put forth by Zollman (2010) and show that the above three conditions produce a distortion in group belief. Finally, we consider a number of possible countermeasures. Here we show how some practices that may seem to promote unbiased inquiry (i.e., independent funding) actually exacerbate the problem if not implemented with additional restrictions.

**2. Class 1 Antiarrhythmic Drugs: A Confluence of Interests.** The Cardiac Arrhythmia Suppression Trial (CAST) was convened to determine whether antiarrhythmic drugs prevented heart attacks. An arrhythmia (i.e., an irregular heartbeat) precipitated the majority of heart attacks, and it was believed that if arrhythmias could be suppressed, then heart attacks could be prevented as well. The CAST randomized clinical trial showed how dangerous relying on such theories is, and this historical episode has been used by philosophers as a paradigm case for the use of randomized clinical trials (e.g., Howick 2011). The trial showed that not only were antiarrhythmic drugs ineffective, but they dramatically increased the odds of a patient dying in general and dying from a heart attack in particular (CAST Investigators 1989). Because antiarrhythmic drugs were both harmful and widely prescribed, it is estimated that their use caused tens, if not hundreds, of thousands of deaths (Moore 1995).

Yet we can push the problem back and ask why a causal explanation was a plausible basis for widespread prescription. This is even more perplexing when one realizes that when Harvard professor Bernard Lown (1979) first proposed the arrhythmia suppression hypothesis, he did so with caution. Indeed, he specifically warned that “commonly used antiarrhythmic agents may sometimes aggravate arrhythmia and thus threaten survival” (317) and claimed that “in the vast majority of patients, [heart arrhythmias] require no treatment” (321). Because of such concerns, the cardiological research team around Harvard adopted a number of cautious positions. They employed the Lown grading system to stratify patients by risk (limiting treatment to high-risk patients) and tested drugs using death as a primary end point rather than rely on the arrhythmia suppression hypothesis in evaluation (Graboyes et al. 1982).

Lown’s caution was not universally shared. There was a diversity of opinion on what methods should be used to judge efficacy, how dangerous antiarrhythmic drugs were, and what populations stood to benefit. Two researchers in particular, Joel Morganroth and Robert Winkle, were enthusiastic about the therapeutic potential of antiarrhythmic drugs and were making their reputations by using a surrogate endpoint to measure antiarrhythmic efficacy (Morganroth et al. 1978; Winkle 1978). While there were intermediate positions between Lown’s Harvard group and researchers like Morganroth and Winkle, these positions represented ends of the spectrum in multiple respects.<sup>1</sup> First and foremost, the measures of efficacy (arrhythmia suppression)

1. As an example of an intermediate position, some researchers felt that efficacy should be demonstrated by the ability of antiarrhythmic drugs to block an electrically induced heart attack. Such trials would also be short term but more invasive, and they tended to produce lower estimates of efficacy than trials that just used arrhythmia suppression (e.g., Anderson, Lutz, and Allison 1983). For a fuller description of this case, see Holman (forthcoming).

promoted by Winkle and Morganroth could be demonstrated in a matter of days or weeks, whereas long-term survival would take years to demonstrate and would cost far more to investigate.

It should not be terribly surprising that when pharmaceutical companies looked for academic researchers to partner with, they reached out to researchers like Morganroth and Winkle. Likewise, when pharmaceutical manufacturers convened a conference of cardiac researchers to discuss how such drugs should be evaluated by the FDA (Food and Drug Administration), it was Morganroth whom they arranged to chair the conference. The conference accepted three presentations that addressed the measurement of efficacy. They first argued that measuring prevention of cardiac death was "unreasonably expensive" and detailed how arrhythmia suppression could be used as a surrogate endpoint (Hoffman 1981). A second speaker argued that the Lown classification system was inappropriate for research purposes (Campbell 1981). Finally, Morganroth himself elaborated how arrhythmia data should be analyzed (1981). While researchers with alternative views were in attendance, they were not provided with the stage.

During the early 1980s the industry continued to fund studies that used arrhythmia suppression as a measure of efficacy, and it slowly became the default methodology in the community, used by researchers both with industry funding (e.g., Morganroth 1986) and without (e.g., Pratt et al. 1986). Rather than actively oppose researchers that held unfavorable views, industry simply chose not to fund their work. For example, as Moore (1995) details in his account of the antiarrhythmic drug disaster, although Winkle originally received industry funding, after he identified a potential for deadly side effects and began much more rigorous investigations of antiarrhythmic drugs, pharmaceutical companies merely canceled their contracts with him. In contrast, researchers that continued to hold industry-friendly views, continued to receive generous funding and as a consequence published prolifically. In so doing, they became highly influential members of the cardiology community; like Morganroth, they were asked to write textbooks, organize symposia, and sit on FDA advisory panels (Moore 1995).

It is worth noting that the researchers such as Morganroth were not hacks. The views they put forward after they began working with industry were no different from before.<sup>2</sup> Moreover, many industry-funded researchers played lead roles in the CAST that eventually showed the danger of antiarrhythmic drugs (CAST Investigators 1989). While industry-funded researchers faced conflicts of interest, they could honestly say that they did not hold their views because they worked for industry. In fact, it was the other way around.

2. For example, compare the research methodology advocated by Morganroth et al. (1978) with Morganroth (1986).

The community did become biased, but as a result of an *industrial selection effect*: it was because of the views researchers antecedently held that industry contracted them in the first place. Because the method they advocated yielded industry-friendly results, industry provided them with resources to promote their position, and because these resources allowed them to be immensely productive, they attained positions of influence within the community and built a scientific consensus around their view. Most crucially, to bias the scientific consensus, industry did not need to corrupt a single individual; they merely changed the selection pressures inside of a merit-based system.

**3. Network Epistemology.** As illustrated in the previous section, industry can play a large role in shaping an epistemic community by providing institutional and financial support to select researchers. This, we contend, can result in industry-favorable outcomes, despite the fact that all agents in the epistemic community are just motivated by a desire to discover the truth. In this section, we turn to a simple formal model of group inquiry pioneered by Kevin Zollman (2010) that has been previously used to investigate the medical research community (Holman and Bruner 2015). Above, we argued that the eventual scientific consensus among medical researchers was the result of industrial selection; however, readers may harbor a concern that even if researchers were not intentionally skewing results, they must have been unconsciously affected by the funding they received. In the simulations below, we allay such concerns by examining a community of (myopically) rational Bayesian agents whose beliefs and decisions regarding what experiments to run are unquestionably unaffected by whether they receive such funding. After briefly introducing Zollman's model, we make three alterations to explore the effect industry has on scientific inquiry.

Zollman's model consists of a group of individuals who all have the choice of taking one of two actions (call them  $A$  and  $B$ ). These actions can be taken to represent a variety of things, such as the administration of different drug regimens or even the application of competing scientific theories. Each action leads to success with some probability ( $p_A$  and  $p_B$  for  $A$  and  $B$ , respectively) and results in failure otherwise. To make this concrete, if the action in question corresponds to administering a drug regimen to a patient, then a 'success' would hold if the individual recovers. During each time period of the model, individuals perform  $E$  times the action they believe is more likely to be successful. Individuals then update their beliefs on the basis of these data.<sup>3</sup> Furthermore, individuals take into account not just the data they themselves produce but also the data generated by those they are con-

3. In particular, the agents apply Bayes's rule to their degree of belief taking into account evidence gathered by the agent and her neighbors. See Zollman (2010) for the exact details.

nected to on a social network. In other words, individuals share information with others in their community. Our model makes three alterations to Zollman's.

First, we relax the assumption that all researchers are equally productive. Instead, individuals vary with regard to how many experiments or trials they conduct per period. Some are able to perform a large number of trials in a given time period, while others are more limited in their productivity.

Second, and perhaps most important, we assume that individuals performing the same action need not have the same probability of success. In other words, individual  $i$ 's success rate with action  $A$  ( $p_A^i$ ) need not be equal to  $j$ 's success rate with action  $A$  ( $p_A^j$ ). We determine agent  $i$ 's initial level of success when performing action  $A$  by a draw from a normal distribution centered at  $p_A$  with variance  $\sigma^2$ . Likewise,  $p_B^i$  is determined by a draw from a normal distribution centered at  $p_B$  with a variance of  $\sigma^2$ . This modification to the baseline model captures the fact that while individuals are ostensibly performing the same action when they both, say, administer drug regimen A, they may be using different metrics to determine whether the outcome was in fact a 'success.' Thus, the value of  $p_A^i$  and  $p_B^i$  specifies the *methodological bias* of agent  $i$ .

Finally, we incorporate an exit-entry dynamic into our model to capture the turnover that routinely takes place in real epistemic communities. In each round, with probability  $e$ , one individual will be randomly selected to leave the community. This agent will be replaced by a new individual,  $j$ . Agent  $j$  is randomly assigned beliefs regarding the efficacy of actions  $A$  and  $B$ . Additionally, agent  $j$ 's productivity level (i.e., the number of actions  $j$  takes per round) is also randomly assigned. Finally, the values of  $p_A^j$  and  $p_B^j$  will be inherited from another individual in the epistemic community as if being trained in proper methods by a researcher in the existing community. Furthermore, the chance that agent  $j$  adopts the methodological bias of another individual,  $k$ , is determined by the productivity of agent  $k$ . In particular, the probability that  $j$  adopts the methodological bias of agent  $k$  is  $E_k / \sum_{n \neq j} E_n$ , where  $E_k$  is the productivity level of agent  $k$ . We believe this exit-entry dynamic is descriptive of real epistemic communities. Newcomers to an academic community are influenced by the previous generation of researchers, and more influential researchers are more likely to train new members joining the community.

The exit-entry dynamic described above is an instance of the so-called Moran process from evolutionary biology. The Moran process is a standard model of evolutionary change for a finite population. Traulsen, Claussen, and Hauert (2005) show that as the number of individuals in the population increases, the dynamics of the Moran process approximates the behavior of the widely used replicator dynamics, a standard model of cultural and biological evolution. Thus, the Moran process is an apt model for our purposes

and has been used to model the process of cultural evolution in small populations (e.g., Wagner 2015).

## 4. Results

*4.1. Baseline Results.* We now run a number of simulations of the above model to determine the likelihood that the community converges on the more successful action. Recall that the likelihood individual  $i$  achieves success when performing action  $A$  is  $p_A^i$  and that this value is initially determined by a draw from the normal distribution  $N(p_A, \sigma^2)$ .

We find that as the initial level of methodological diversity in the community ( $\sigma^2$ ) increases, the community is less likely to converge to the superior action. In fact, many simulation runs result in all individuals performing the less successful action when  $\sigma^2$  takes on moderate values. Details regarding our simulations are illustrated in figure 1. To better understand the connection between  $\sigma^2$  and the reliability of the network, consider the case in which action  $A$  is superior to action  $B$  (i.e.,  $p_A > p_B$ ). If  $\sigma^2$  is set to zero, then  $p_A^i > p_B^i$  for any agent  $i$  in the community. However, when  $\sigma^2$  is greater than zero, it is possible for an individual to adopt a methodology that suggests  $B$  is superior (i.e.,  $p_B^i > p_A^i$ ). This possibility becomes all the more likely as the

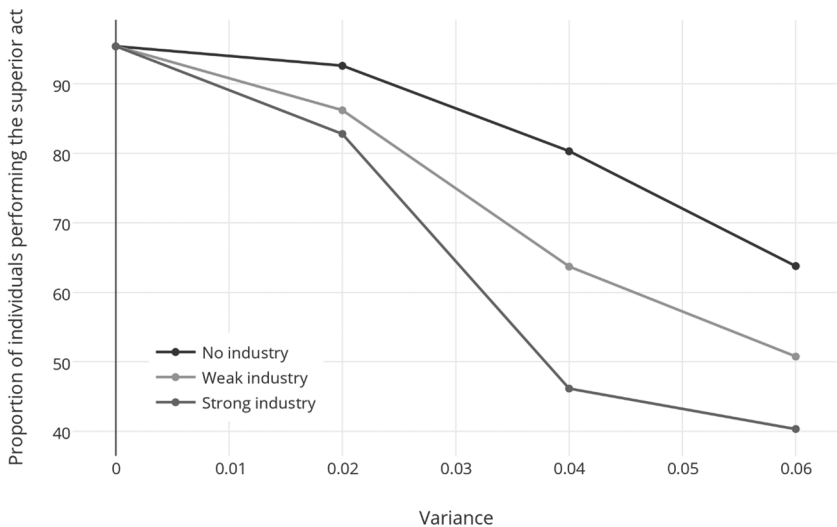


Figure 1. Proportion of individuals performing the superior act as a function of the variance of the normal distribution that initially sets the methodological biases of agents in the community. Simulations are of 20-person epistemic communities,  $p_A = .5$ ,  $p_B = .45$ ,  $T = 0.03$ ,  $F = 0$  (top),  $F = 20$  (middle),  $F = 100$  (bottom). Color version available as an online enhancement.



value of  $\sigma^2$  increases. Thus, for moderate to large values of  $\sigma^2$  it is possible that a large number of individuals adopt a methodology that favors *B*. This significantly affects not only the beliefs of others in the network (they receive information from their peers suggesting that *B* is superior to *A*) but also shapes the methodological bias of newcomers. Accordingly, as methodological diversity increases, the community is less likely to converge on the superior act.<sup>4</sup>

*4.2. Industry Funding.* We now allow for the possibility of industry funding. We assume that when individuals receive funding from industry their productivity each round increases by  $F$ .<sup>5</sup> This has two effects. First, it makes the recipient more influential in the short run—now that she is conducting more trials, she will have a larger impact on her peers. Second, however, it means that she is more likely to train new members of the community (recall that newcomers adopt the methodological bias of others in proportion to their productivity). In this way industry can indirectly influence beliefs as well as the composition of the epistemic community.

In line with the antiarrhythmic drug case study discussed in section 2, we consider a situation in which there are two options: treatment with an antiarrhythmic drug or no treatment. Throughout, we assume that withholding “treatment” is in fact superior to an antiarrhythmic drug ( $p_A > p_B$ ). Industry chooses to fund an individual,  $i$ , if  $p_B^i > p_B + T$ . In other words, industry has an established threshold, and if individuals are using a methodology that increases the apparent efficacy of the drug above this threshold, they will receive industry funding.

Figure 1 illustrates how the inclusion of industry funding affects the reliability of the epistemic community. We find that even low values of  $F$  and  $T$  significantly lower the chances that the network converges to the superior action and that as the amount of funding doled out ( $F$ ) increases, a true consensus becomes less likely. Also note that industry can only effectively manipulate the community when there is methodological diversity in the community (i.e., when  $\sigma^2 > 0$ ). If there is no methodological diversity to begin with, then industrial selection does not occur.

*4.3. Industrial and Independent Sources of Funding.* As we have seen, industry can have a sizable effect on the accuracy of an epistemic community. One countervailing force that could potentially nullify the effect of industry funding is for an alternative independent agency (e.g., the National

4. In general, the qualitative results reported in this article held across a variety of different network structures (i.e., both dense and sparse communication networks).

5. Productivity is initially assigned by a draw from the uniform distribution from 50 to 100. The probability of exiting the epistemic community is 0.02 in each round.



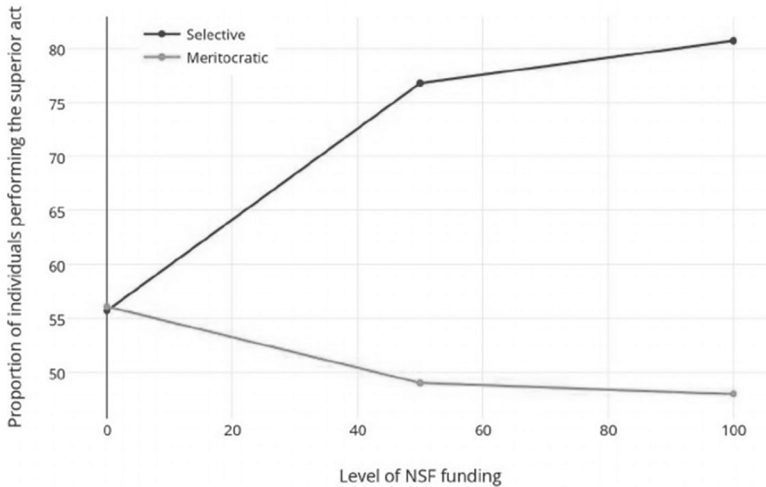


Figure 2. Proportion of individuals performing the superior act as a function of the level of NSF funding and NSF funding policy. Simulations are for 20-person epistemic communities,  $p_A = .5$ ,  $p_B = .45$ ,  $T = 0.04$ ,  $F = 20$ ,  $\sigma^2$ . Color version available as an online enhancement.

Science Foundation [NSF], National Institutes of Health [NIH]) to fund members of the community. In this section we consider the influence an ‘NSF’ funding body has on an epistemic community that also receives industry funding.

Just as we did with industry, we must specify how, exactly, the NSF allocates funds. We consider a meritocratic policy in which the NSF funds the most productive researchers and a selective policy in which the NSF neglects work done for industry when assessing a researcher’s credentials. First, we consider the meritocratic policy in which the NSF funds any researcher above a specified ‘productivity threshold’. Note that any help agents may have received from industry will go toward their meeting the threshold set by the NSF. Thus, those with industry-favorable views may luck out twice—once due to industry and yet again if industry funding allows them to qualify for additional help from the NSF. We contrast the meritocratic policy with a selective policy, in which individuals are once again funded if they meet the ‘productivity threshold’, but this time additional trials provided by industry are not counted when determining whether the agent qualifies for NSF support.<sup>6</sup>

Figure 2 shows the results of simulations comparing the effect these two policies have on the reliability of the epistemic community. As is evident, when industry is discriminatively funding researchers with favorable method-

6. In both of these cases the productivity threshold was set at 75.

ological biases the NSF does best to adopt the selective policy. A meritocratic policy actually compromises the epistemic community because it results in the NSF disproportionately funding individuals with industry-favorable biases. The implementation of the selective policy, however, significantly improves the epistemic community.

**5. Conclusion.** The primary concern about industry funding has heretofore been that a researcher will face a conflict of interest and either intentionally or unintentionally place corporate profits over truth. This same conception also leads industry-funded researchers to vehemently assert that their research has not been corrupted by corporate sponsorship. Our claim is that it is entirely possible for every individual researcher to remain unaffected by industry funding and for the community as a whole to be biased. Indeed, our simulation results show that in a merit-based system, in which researchers must make methodological decisions that affect their estimates of efficacy, a community can be biased even if it is made up of rational agents who are completely unaffected by receiving industry funding. Furthermore, the effect of “industrial selection” is all the more insidious because individuals may truly believe that their work has not been compromised.

Industrial selection biases a community without compromising the integrity of any individual. Yet the effect on individuals has been the focus of concern with industry funding. In the policy document put together by Douglas et al. (2014) on guidance for responsibly accepting industry funding, the primary concern is with integrity and credibility of the researcher (or group) accepting the funds. To address these issues they recommend policies to protect whistle-blowers, maintenance of a diversity of funding, freedom to publish, regular interaction with a broader range of stakeholders, and securing oversight from an independent third party. While we are not disparaging any of these recommendations, even perfect compliance with such guidance does not address the threat caused by industrial selection.

In the case of the antiarrhythmic drug disaster, there was legitimate scientific disagreement about proper methodology. There was no whistle to be blown, conferences with dissenters regularly occurred, and the entire process was overseen by the FDA. Each individual researcher maintained a credible position, and we can charitably grant that every scientist preserved the integrity of his or her research. What happened was not that individual results were corrupted but that certain ways of probing the world were favored over others. Industry was able to survey the community, identify whose research methods were most economically beneficial, and then increase the amount of research done with those methods and, thus, the probability that those researchers would attain positions of influence.

Douglas et al. (2014) suggest that individuals should seek a diverse array of funding. While this would no doubt address some problems at the individual level, our results suggest that simply having an independent agency fund top researchers may actually make things worse. It is worth noting that many industry-funded researchers in the antiarrhythmic drug disaster also received grants from the NIH (Moore 1995). While the prospect of funding agencies ignoring scientific work done during grant assessment might face opposition, it is clear that to address the problem of industry funding we must also consider the community-level effects of policies.

Finally, while a policy that would address industrial selection may be unpopular, awareness of the phenomena itself may provide an avenue for productive dialogue with industry scientists, as it opens up the possibility of discussing industry bias without impugning the credibility of individuals. Kitcher (1993) showed that a scientific community can be arranged such that it makes “good epistemic use of the grubbiest motives” (305). We take our results to show the inverse, that some community structures will systematically yield biased results even if they are populated by epistemic angels.

#### REFERENCES

- Anderson, J., J. Lutz, and S. Allison. 1983. “Electrophysiologic and Antiarrhythmic Effects of Oral Flecanide Patients with Inducible Ventricular Tachycardia.” *Journal of the American College of Cardiology* 2:105–14.
- Brain, R., J. Stavely, and L. Ortego. 2016. “Resolving the Perception of Bias in a Discipline Founded on Objectivity: An Industry Perspective.” *Environmental Toxicology and Chemistry* 35 (5): 1070–72.
- Campbell, R. W. 1981. “Evaluation of Antiarrhythmic Drugs: Should the Lown Classification Be Used.” In *The Evaluation of New Antiarrhythmic Drugs*, ed. J. Morganroth, E. N. Moore, L. S. Dreifus, and E. L. Michelson, 113–22. Boston: Nijhoff.
- CAST Investigators. 1989. “Preliminary Report: Effect of Encainide and Flecaïnide on Mortality in a Randomized Trial of Arrhythmia Suppression after Myocardial Infarction.” *New England Journal of Medicine* 321:406–12.
- Conrad, J., and R. Becker. 2011. “Enhancing Credibility of Chemical Safety Studies: Emerging Consensus on Key Assessment Criteria.” *Environmental Health Perspectives* 119:757–64.
- Douglas, H., K. Elliott, A. Maynard, P. Thompson, and K. Whyte. 2014. “Guidance on Funding from Industry.” SRPoiSE.org. <http://srpoise.org/wp-content/uploads/2014/06/Guidance-on-Funding-from-Industry-Final.pdf>.
- Drazen, J. M. 2015. “Revisiting the Commercial-Industry Interface.” *New England Journal of Medicine* 372:1853–54.
- Elliott, K. 2014. “Financial Conflicts of Interest and Criteria for Research Credibility.” *Erkenntnis* 79:917–37.
- Graboyes, T. B., B. Lown, P. J. Podrid, and R. DeSilva. 1982. “Long-Term Survival of Patients with Malignant Ventricular Arrhythmia Treated with Antiarrhythmic Drugs.” *American Journal of Cardiology* 50:437–43.
- Hoffman, B. F. 1981. “Relationship between Effects on Cardiac Electrophysiology and Antiarrhythmic Efficacy.” In *The Evaluation of New Antiarrhythmic Drugs*, ed. J. Morganroth, E. N. Moore, L. S. Dreifus, and E. L. Michelson, 5–16. Boston: Nijhoff.
- Holman, B. Forthcoming. “Philosophers on Drugs.” *Synthese*.
- Holman, B., and J. Bruner. 2015. “The Problem of Intransigently Biased Agents.” *Philosophy of Science* 82:956–68.

- Howick, J. 2011. *The Philosophy of Evidence-Based Medicine*. West Sussex: British Medical Journal Books.
- Kitcher, P. 1993. *The Advancement of Science: Science without Legend, Objectivity without Illusions*. Oxford: Oxford University Press.
- Lo, B., and M. Field. 2009. *Conflicts of Interest in Medical Research Education and Practice*. Washington, DC: National Academies Press.
- Lown, B. 1979. "Sudden Cardiac Death: The Major Challenge Confronting Contemporary Cardiology." *American Journal of Cardiology* 43:313–28.
- McGarity, T., and W. Wagner. 2008. *Bending Science: How Special Interests Corrupt Public Health Research*. Cambridge, MA: Harvard University Press.
- Moore, T. 1995. *Deadly Medicines: Why Tens of Thousands of Heart Patients Died in America's Worst Drug Disaster*. New York: Simon & Schuster.
- Morganroth, J. 1981. "Long-Term Ambulatory Electrocardiographic Recording in the Determination of Efficacy of New Antiarrhythmic Drugs." In *The Evaluation of New Antiarrhythmic Drugs*, ed. J. Morganroth, E. N. Moore, L. S. Dreifus, and E. L. Michelson, 103–12. Boston: Nijhoff.
- . 1986. "Encainide for Ventricular Arrhythmias: Placebo Controlled and Standard Comparison Trials." *American Journal of Cardiology* 58:74C–82C.
- Morganroth, J., E. L. Michelson, L. N. Horowitz, M. E. Josephson, A. S. Pearlman, and W. B. Dunkman. 1978. "Limitations of Routine Long-Term Electrocardiographic Monitoring to Assess Ventricular Ectopic Frequency." *Circulation* 58:408–14.
- Pratt, C., et al. 1986. "Efficacy and Safety of Moricizine in Patients with Ventricular Tachycardia: Results of a Placebo-Controlled Prospective Long-Term Clinical Trial." *Circulation* 73:718–26.
- Resnik, D. B., and K. C. Elliott. 2013. "Taking Financial Relationships into Account When Assessing Research." *Accountability in Research* 20:184–205.
- Robinson, M. Forthcoming. "Financializing Epistemic Norms in Contemporary Biomedical Innovation." *Synthese*.
- Traulsen, A., J. Claussen, and C. Hauert. 2005. "Coevolutionary Dynamics: From Finite to Infinite Populations." *Physical Review Letters* 95:238701.
- Wagner, E. 2015. "Conventional Semantic Meaning in Signaling Games with Conflicting Interests." *British Journal for the Philosophy of Science* 66 (4): 751–73.
- Winkle, R. A. 1978. "Antiarrhythmic Drug Effect Mimicked by Spontaneous Variability of Ventricular Ectopy." *Circulation* 57:1116–21.
- Zollman, K. 2010. "The Epistemic Benefit of Transient Diversity." *Erkenntnis* 72:17–35.

Copyright of Philosophy of Science is the property of Philosophy of Science Association and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.