

# What is Epistemically Wrong with Research Affected by Sponsorship Bias?

## The Evidential Account

Alexander Reutlinger (LMU Munich)

### 1. Introduction

Philosophy of science typically addresses ‘foundational’ philosophical questions about instances of successful scientific research. Although this is certainly one crucial task for philosophers of science, our discipline should not exclusively focus on successful science. Philosophers of science should also critically assess philosophical issues motivated by research that does not constitute scientific success – first and foremost, *biased research*. Because biased research occurs frequently, analyzing and criticizing biased research is becoming a central and urgent topic in the present philosophy of science.<sup>1</sup>

In this paper, I will focus on one particular kind of biased research: research that is subject to *sponsorship bias*. A scientific study is affected by a sponsorship bias iff the study is funded by a financially interested sponsor, and the outcome of the study is significantly and systematically wrong (or, alternatively, distorted or erroneous) in a way aligning with the sponsor’s financial interests. I will present and discuss examples of three types of sponsorship bias, as these types are perceived as being widespread and paradigmatic in the recent literature<sup>2</sup>:

1. a type of sponsorship bias regarding the choice of the experimental design

---

<sup>1</sup> See, for instance, Biddle (2007); Brown (2008); Sismondo (2008); Wilholt (2009); Radder (2010); Carrier (2013, 2017); Biddle and Leuschner (2015); Harker (2015); Elliott (2016); Biddle et al. (2017); de Melo-Martín and Intemann (2018); Holman and Elliott (2018); Steel (2018); Stegenga (2018: chapter 10).

<sup>2</sup> These types are usually included into lists of paradigmatic biases in science; see Resnik (2000: 265-276), Sismondo (2008: 1911-12), Wilholt (2009: 93), Holman and Elliott (2018: 4), Steel (2018: 129-35), and Stegenga (2018: 154-9). Holman and Elliott (2018: section 3) provide further useful references. Moreover, I refer to the relevant literature on each type of bias in Section 2.

- of a study (with the ‘Bisphenol A case’ as an example);
2. a type of sponsorship bias regarding the selection and representation of the data obtained in a study (illustrated by the ‘Celebrex case’);
  3. a type of sponsorship bias regarding the choice of the scientific concepts used to interpret the results of a study (exemplified by the ‘Tobacco case’).

In relation to these three types of sponsorship bias, I aim to provide an analysis of the notion of ‘wrongness’ (or, alternatively, of distortion or error) figuring in the characterization of sponsorship bias just above. In this paper, I take such an analysis to consist in answering the following *epistemological* question:

What is *epistemically wrong* (that is, *unjustified*) with examples of research affected by sponsorship bias?<sup>3</sup>

My answer to this question and the main claim of this paper is, what I call, the *evidential account of epistemic wrongness*:

Research affected by sponsorship bias is epistemically wrong if and only if the researchers in question make false claims about the (degree of) evidential support of some hypothesis H by data E.

In this paper, I have a rather modest goal: to argue that the evidential account captures the epistemic wrongness of the three types of sponsorship bias mentioned above. Of course, there might be further types of sponsorship bias, as indicated in the recent literature: for instance, a type of sponsorship bias regarding the choice of topic for empirical research (Resnik 2000: 267-9; Holman and Elliott 2018: 4) and different types of sponsorship bias that involve p-hacking in the analysis of data (Stegenga 2018: 157-8). For this reason, it is a fruitful task for future research to discuss whether the evidential account can be also be applied to further types of sponsorship bias.

The plan of the paper is as follows:

In Section 2, I will present concrete examples illustrating the three

---

<sup>3</sup> I do not mean to suggest that this is the only question worth asking with respect to sponsorship bias. Other fruitful questions include, for instance, what is politically and morally wrong with biased research. For work on moral and political wrongness of biased research see, for instance, Shrader-Frechette (1994); Kitcher (2011); Kourany (2016). Moreover, Biddle et al. (2017: 166) advocate a related distinction between epistemically, morally, and politically “detrimental” criticism of scientific research.

mentioned types of sponsorship bias.

In the central Section 3, I will argue that the evidential account provides a single satisfactory answer to the question as to what is epistemically wrong with all of the three types of sponsorship. In Section 3.1, I will apply the evidential account to the Bisphenol A case. Furthermore, in order to motivate my application of the evidential account to examples of the remaining two types of sponsorship bias, I will specify how the evidential account builds on Carrier's "false advertising" account (Carrier 2013, 2017, 2018). I will argue that the evidential account can be portrayed as expanding and making more precise the core idea of Carrier's account in a way that renders it applicable to the Celebrex case (Section 3.2) and the Tobacco case (Section 3.3). In section 3.4, I will rebut potential counterexamples to the evidential account stemming from Steel's (2018) recent work on sponsorship bias. Finally, in Section 3.5, I will present two virtues of the evidential account.

In Section 4, I will discuss two prominent (potential) alternatives to the evidential account in the recent literature: Wilholt's (2009) conventionalist view (my terminology), according to which epistemic wrongness is defined in terms of violating methodological conventions (Section 4.1), and Oreskes and Conway's (2010) account that is primarily based on the notion of disagreement with expert consensus (Section 4.2). I will argue for two claims: (1) neither of the two views provides necessary and sufficient conditions for epistemic wrongness; (2) the evidential account should be regarded as a fruitful complement to and not necessarily as a competitor of these two views.

In Section 5, I will summarize the results of my discussion. Section 6 is an appendix providing a formal Bayesian argument in support my assessment of the Bisphenol A case presented in Section 3.1.

## **2. Types of Sponsorship Bias: Three Examples**

In this section, I will present three examples of sponsorship bias in science: an example of sponsorship bias regarding the choice of experimental design of a study, the Bisphenol A case (Section 2.1), an example of sponsorship bias regarding the selection and representation of data, the Celebrex case (Section 2.2), and an example of a sponsorship bias regarding the choice of the scientific concepts used to interpret the results of a study, the Tobacco case (Section 2.3).

### **2.1 The Bisphenol A case**

The Bisphenol A case illustrates the type sponsorship bias regarding the choice

of experimental design of a study. In this particular case, the choice concerns the animals used in a study as a specific part of the experimental design.

Bisphenol A is an environmental estrogen that is used to manufacture plastic materials that are, for instance, contained in food and beverage cans. The question arises whether Bisphenol A compromises the health of people exposed to it. In particular, research was done to answer the question whether Bisphenol A – even in low doses – causes certain types of cancer such as prostate cancer. In the scientific and philosophical literature, many of the industry-funded studies on the effects of Bisphenol A are taken to be paradigm cases of sponsorship bias (for an influential survey and meta-study of the scientific literature, see vom Saal and Hughes 2005; philosophical analyses include, for instance, Wilholt 2009: 93; Carrier 2013: 2560-2561; Biddle and Leuschner 2015: 271-272). In my exposition of the example, I will rely on vom Saal and Hughes' (2005).

Publicly funded studies and industry-funded studies on whether low doses of Bisphenol A cause cancer arrived at strikingly different conclusions: 90% of the publicly funded studies on low-dose exposure to Bisphenol A reported significant effects; however, none of the industry-funded studies reported such effects (vom Saal and Hughes 2005: 928). As vom Saal and Hughes argue, the presence of a sponsorship bias is perhaps the best explanation for the striking divergence in the results: “Source of funding is highly correlated with positive and negative findings in published articles.” (vom Saal and Hughes 2005: 928).

One important manifestation of sponsorship biases in industry-funded studies on Bisphenol A concerns the choice of the experimental design. A case in point is the choice of animal used in the studies. Industry-funded studies used rats that were known to be insensitive to low-doses of substances similar to Bisphenol A such as DES or ethinylestradiol – that is, rats of the CD-SD strain. DES, however, was already known to be a carcinogen. Hence, the fact that DES did not cause cancer in the CD-SD rats was generally taken to be evidence for the claim that rats of this strain are an inappropriate model organism for studies on the carcinogenicity of Bisphenol A. As vom Saal and Hughes point out, one had good reasons for not using CD-SD rats: “The very low sensitivity of the CD-SD rat strain to BPA [i.e. Bisphenol A] was predicted by its low sensitivity to ethinylestradiol when it was included as a positive control.” (vom Saal and Hughes 2005: 929) For this reason, it is accurate to say that the insensitivity of CD-SD rats is an empirical lesson from cancer research closely related to research on Bisphenol A. Unsurprisingly, all studies on Bisphenol A involving CD-SD rats report the absence of significant effects. Moreover, industry-funded studies did not only use rats of the CD-SD strain but they also omitted positive controls indicating that CD-SD rats are insensitive:

“Two industry-funded studies (Ashby et al. 1999; Cagen et al. 1999) were designed with DES included as a positive control, which was reported by industry spokesmen (Toloken 1998) at a public news briefing about the Cagen et al. (1999) study. A critique [Environmental Data Services (ENDS) 1998] pointed out that the positive control, DES, failed to show a difference from the negative controls in each of these studies (Ashby et al. 1999; Cagen et al. 1999); *however, the authors did not indicate in their published articles that DES had been used as the positive control. Subsequent studies funded by chemical corporations, all of which have reported the absence of significant effects for low doses of BPA [i.e. Bisphenol A], avoided this problem by simply not including a positive control in the experiment.*” (vom Saal and Hughes 2005: 929; emphasis added)

Vom Saal and Hughes disclose the interesting fact that a small proportion of *publicly funded* studies on the effects of Bisphenol A also used CD-SD rats *initially*. However, unlike industry-funded studies, publicly funded studies stopped using CD-SD rats (for the reason outlined in the previous paragraph) and relied on sensitive animals, such as the fetal male CF-1 mouse, in subsequent studies: “If the studies that used the CD-SD rat are eliminated from consideration, 94 of 98 (96%) government-funded studies report significant effects of low doses of BPA [i.e. Bisphenol A], whereas 0 of 8 (0%) industry-funded studies report significant effects with the same low doses.” (vom Saal and Hughes 2005: 929)

In sum, the Bisphenol A case is an instructive but also an extreme case of a sponsorship bias (regarding the choice of animals used in the study). But although the case is extreme, there are further examples of this type of sponsorship bias regarding experimental design (Bero and Rennie 1996; Lexchin et al., 2003; vom Saal and Hughes 2005; McGarity and Wagner, 2008; Carrier 2017, 2018). Now, let me take a look at two further examples of sponsorship bias.

## **2.2 The Celebrex case**

Another common type sponsorship bias concerns the selection and representation of data. In order to illustrate this type, I will use the Celebrex case, involving a choice of data segments in a study on the efficacy of the drug Celebrex.

Brown describes the Celebrex case as follows:

“Celebrex, which is used in the treatment of arthritis, was the subject of a year-long study sponsored by its maker, Paramacia [...]. *The study purported to show that Celebrex caused fewer side effects than*

*older arthritis drugs*. The results were published in JAMA (Journal of the American Medical Association) along with a favorable editorial. *It later turned out that the encouraging results were based on the first six months of the study. When the whole study was considered, Celebrex held no advantage over older and cheaper drugs*. On learning this, the author of the favorable editorial was furious and remarked on ‘a level of trust that was, perhaps, broken’”. (Brown 2008: 193; emphases added)

The Celebrex case is just one among many examples of a sponsorship bias concerning the selection and representation of data (for a discussion of further cases, see Biddle 2007; Brown 2008; Carrier 2017; Douglas 2000; Doucet and Sismondo 2008; Matheson 2008; Michaels 2008; Wilholt 2009; Elliott, 2010). Examples of this type of sponsorship bias have the following general structure: first, researchers gather experimental data in the course of a study. Second, researchers select a part of the complete available data. The partial data, when considered in isolation from the rest of the complete available data, is indeed ‘favorable’, that is, supporting evidence for some hypothesis H. However, researchers omit, or ignore, the ‘unfavorable’ remainder of the complete available data in the publication of the results of the study – typically without making the omission explicit. Third, such a partial omission of data is characteristically in line with the economic interests of the sponsor funding the study in question.

### **2.3 The Tobacco case**

I will use the Tobacco case to illustrate a sponsorship bias regarding the choice of the scientific concepts used to interpret the results of a study. In the Tobacco case, the choice of concept concerns a particular concept of causation.

The Tobacco case involves the attempt to “create doubt” about a result of scientific research. Here, a financially interested sponsor employs scientists in order to “create doubt” with respect to established results of scientific research – results that are in conflict with the sponsor’s financial interests. The notion of creating doubt (or, alternatively, producing or manufacturing doubt) is inspired by the infamous expression “doubt is our product” used in an internal memo entitled “Smoking and Health Proposal” of the tobacco industry (in this case, the company Brown & Williamson) from 1969, as documented by Glantz et al. (1996: 171), Michaels (2008a: 11), Proctor (2008: 17; 2011: 289, 617n), and

Oreskes and Conway (2010: 34).<sup>4</sup>

One notorious and politically relevant cascade of examples of creating doubt consists in the fact tobacco companies hired scientists in order to undermine medical research supporting the claim that smoking cigarettes causes lung cancer – a claim that was, and still is, in conflict with the financial interests of tobacco companies. These scientists “created doubt” about the causal link between smoking cigarettes and lung cancer.<sup>5</sup>

The creation of doubt takes many different forms and strategies. It sometimes takes the form of biased counter-studies intended to question established research results (for instance, on the effects of smoking cigarettes). As other philosophers, historians of science and scientists have argued, such counter-studies make use of inadequate statistics, ad hoc adjustments, and inconsistent data fitting (Oreskes and Conway 2010; Biddle and Leuschner 2015), are permeated by incorrect empirical statements (Lewandowsky et al. 2018), and often data has been manipulated or deliberately ignored. Although I suspect that the epistemic wrongness of such counter-studies might often turn out to be of the same kind as the Bisphenol A case and the Celebrex case, it is not my goal in this paper to provide a general and exhaustive account of what is epistemically wrong with the creation of doubt, if it takes the form of biased counter-studies. Instead, my focus will be on another influential type of creating doubt that does *not* manifest itself in conducting biased counter-studies.

I will focus on one way of creating doubt consisting in a sponsorship regarding the choice of the *scientific concepts* used to interpret the results of established medical studies on the effects of tobacco smoke. Consider the following concrete historical example where the sponsorship bias concerns the

---

<sup>4</sup> Appealing to Proctor’s (2008) concept of agnotology (the study of ignorance), the concept of “agnotology” is also sometimes used to refer to the creation of doubt (for instance, Carrier 2018; see also Fernández Pinto 2017 for a helpful classification of different types of agnogenesis and a survey to the recent literature on agnotology.) But I think one has to be cautious with Proctor’s notion of agnotology, because it is considerably more inclusive than the notion of creating doubt. For instance, military research is a subject for agnotology, according to Proctor, because it is secret research. It is at least not obvious (to me) that secret military research is also an instance of creating doubt. For this reason, I use the more restrictive concept of the creation of doubt.

<sup>5</sup> For case studies of creating doubt beyond the tobacco cases (such as regarding climate science and medical research), see, for instance, Michaels (2008), Proctor and Schiebinger (2008), Oreskes and Conway (2010), Biddle and Leuschner (2015), and Biddle et al. (2017), de Melo-Martín and Intemann (2014, 2018), Carrier (2017).

choice of a particular concept of causation – a choice that aligns with the sponsors’ financial interests. The scientist Sheldon Sommers was paid by the Council of Tobacco Research (an organization funded directly by the tobacco industry) to create doubt about medical studies supporting the hypothesis that smoking causes lung cancer. Sommers gave the following answers to questions he was confronted with in court:

“Q: Doctor, do you have an opinion presently as to whether cigarette smoke is a cause of lung cancer?

A: In the *scientific sense*, I believe it is not a cause.

Q: When you qualify your answer to say ‘in the scientific sense’, what do you mean by such a qualification?

A: *Scientific evidence of a causative agent* involves that it should be *both necessary and sufficient to produce a condition*. (Sommers 1985: 65-6, emphases added; Proctor 2011: 270)

In this case, the creator of doubt, Sheldon Sommers, claims that a statistical correlation between smoking cigarettes and lung cancer (E) does not, and indeed cannot, provide confirming evidence for the hypothesis that smoking *causes* lung cancer (H), if one presupposes the “scientific sense” of cause. According to Sommers’ “scientific sense” of cause, a cause (or a “causative agent”) is defined as being a necessary and sufficient condition for its effect (Sommers 1985: 66-8). In other words, Sommers appeals to a simple regularity theory of causation. If one adopts this theory of causation, then smoking does not cause cancer because some smokers do not get lung cancer, and some people who suffer from lung cancer never smoked in their lives. Moreover, a mere correlation between smoking cigarettes and lung cancer cannot be supporting evidence for the hypothesis that smoking causes lung cancer, since any correlation fails to live up to the expectation that a cause has to be sufficient for the effect. Sommers’ choice of this particular theory of causation constitutes a sponsorship bias, because the choice and its consequence (that is, smoking does not cause lung cancer) aligns with the financial interests of Sommers’ employer (the tobacco industry), and the choice rests on a mistake, or is at least questionable, as I will argue in Section 3.3.

Sommers’ testimony in court is just one example of the creation of doubt among numerous other cases (see Glantz et al. 1996; McGarity and Wagner 2008; Micheals 2008a, 2008b; Proctor 2008, 2011; Proctor and Schiebinger 2008; Oreskes and Conway 2010). Regarding this type of creating doubt, the arena for the scientists paid by the tobacco industry often was not a scientific one (such as a peer-reviewed journal or conference) but a courtroom or a talk show.



The scientists hired by the tobacco industry did not directly engage with scientists doing medical research but with judges, state attorneys, and journalists instead (Oreskes and Conway 2010; Proctor 2011). Also for this reason, creating doubt contributed significantly to shaping the public opinion on the health effects of cigarette smoking.

### **3. Applying the Evidential Account of Epistemic Wrongness**

According to the evidential account, research that is subject to a sponsorship bias is epistemically wrong if and only if the researchers in question make false claims about the (degree of) evidential support of some hypothesis *H* by data *E*. In this section, I will argue that the evidential account adequately captures what is epistemically wrong with three paradigmatic types of sponsorship bias. In Section 3.1, I will apply the evidential account to the Bisphenol A case. Moreover, I will highlight the relationship between the evidential account and Carrier's false advertising account in order to motivate my application of the evidential account to the remaining two types of sponsorship bias. I will argue that the evidential account can be portrayed as expanding and making precise the core idea of Carrier's false advertising account in a way that renders it applicable to the Celebrex case (Section 3.2), and the Tobacco case (Section 3.3). In section 3.4, I will rebut a potential counterexample to the evidential account stemming from Steel's (2018) recent work. Finally, in Section 3.5, I will present two virtues of the evidential account.

#### **3.1 The Bisphenol A case – discussion**

For a discussion of the Bisphenol A case, it will be useful to introduce the following terminology:

- *Experimental test results E*: low-dose exposure to Bisphenol A does not have any “significant effects” on cancer – that is, there is no significant correlation.
- *Hypothesis H*: low doses of Bisphenol A do *not* cause cancer in humans.
- *Background knowledge K*: experimental design *D* of the industry-funded studies involves insensitive rats of the CD-SD strain. (This background knowledge is an empirical lesson from earlier studies, as pointed out at the end of Section 2.1.)

With this terminology at hand, I will now argue that standard theories of

confirmation and evidence support the evidential account. In particular, I will rely on the two major accounts of confirmation and evidence in the current philosophical literature: *Bayesian confirmation theory* and *frequentist hypothesis testing*.<sup>6</sup> The crucial upshot is that both Bayesian confirmation theory and frequentist hypothesis testing provide arguments in favor of the evidential account.

Let me stress two points to avoid possible misunderstandings about the dialectic role of Bayesian confirmation theory and frequentist hypothesis testing in this section:

- My reconstruction of the two accounts of confirmation and evidence will be rather non-formal because my goal is a non-formal one: to argue that the evidential account is applicable to three types of sponsorship bias. The applicability does not always depend on the details of mathematical machinery built into Bayesian and frequentist approaches, or so I will suggest. Only in Section 3.1, I believe it is necessary to provide a more formal kind of argument to support the evidential account.
- My goal is to (re)describe the case studies in both Bayesian and frequentist terms. The crucial motivation for such a (re)description consists in arguing for the claim that the evidential account can be maintained independently of whether one is an advocate of Bayesian confirmation theory or of frequentist hypothesis testing. Also because of this motivation, I will treat Bayesianism and frequentism as tools. In this paper, I will not defend, or assess the merit of, Bayesian confirmation theory or frequentist hypothesis testing.<sup>7</sup>

*Bayesian confirmation theory* is one of the most widely accepted accounts of evidential support and confirmation in the current literature in philosophy of science and (formal) epistemology (Horwich 1982; Earman 1992; Bovens and Hartmann 2003; Howson and Urbach 2006; Sprenger and Hartmann 2019). In the framework of Bayesian confirmation theory, experimental result E confirms H relative to K if and only if  $P(H|K, E) > P(H|K)$ .

---

<sup>6</sup> These two philosophical accounts mirror the methodological divide between Bayesian and frequentist frameworks in statistics.

<sup>7</sup> A referee identified a further premise I am relying on: that is, Bayesian and frequentist approaches permit that epistemic agents can communicate and come to agree upon what the (degree of) evidential support in a given situation is (for a helpful survey of the literature and a recent assessment of this topic, see Sprenger and Hartmann 2019: chapter 11).

Let us assume for the sake of simplicity that, in the Bisphenol A case, K deductively entails E. (In the paragraph just below, I will turn to a strategy for relaxing this simplifying assumption.) If K deductively entails E, then this has consequences: learning E (given K) does not raise the probability of H. Formally put,  $P(H|K, E) = P(H|K)$ . Hence, according to Bayesian confirmation theory, E does not confirm H relative to K. If this is correct, the scientists involved in industry-funded studies on Bisphenol A falsely assert that E confirms H relative to K. Thus, the evidential account captures the Bisphenol A case.

In the previous paragraph, I worked with the simplifying assumption that the available background knowledge K (including knowledge about the fact that the rats used in the industry-funded study were insensitive to Bisphenol A) deductively entails E (that there is no correlation between low-dose exposure to bisphenol A and cancer). If K deductively entails E, then  $P(E|K)=1$ . What happens if one relaxes this assumption by assigning (perhaps more realistically) a *high probability* to the proposition that CD-SD rats are insensitive to low doses of Bisphenol A? In other words, what if  $P(E|K)$  is high but not equal to 1? Let us call this the *probabilistic version of the Bisphenol A case*.

To deal with the probabilistic version of the Bisphenol A case, I will now introduce two premises (that will also turn out to be useful for analyzing the Celebrex case in Section 3.2):

1. Adopting a Bayesian theory of confirmation, one can use the notion of *comparative confirmation*: the confirmation that H receives from a body of evidence  $E_1$  is *higher* than the confirmatory power that H receives from another body of evidence  $E_2$ . Formally, Bayesians can express this situation in terms of the following inequality:  $P(H|E_1) > P(H|E_2)$  (see, for instance, Horwich 1982; Fitelson and Hawthorne 2010).
2. I take the following *principle of complete local evidence* as a premise: if one wants to assess whether some hypothesis H receives evidential support from a study, then – in order to guard against defeating evidence – one ought to consider *all* of the data produced by the experiments involved in that study (or perhaps in a series of studies on the same subject matter). The principle of complete local evidence is weaker than Carnap’s well-known requirement of total evidence. According to the latter, “the total evidence available must be taken as basis for determining the degree of confirmation” (Carnap 1950: 211). Carnap’s notion of total evidence is global in the sense that it goes beyond the data produced by an experimental study and includes the entire

“knowledge situation” (ibid.) of a person. The principle I employ allows for a ‘local’ restriction to the data produces by a particular study.<sup>8</sup>

Making use of the notion of comparative confirmation and the *principle of complete local evidence*, I will now apply the evidential account to the probabilistic version of the Bisphenol A case.

Suppose once more that  $P(E|K)$  is high (but not equal to 1), that is,  $P(E|K) \approx 1$ . Moreover, suppose that  $P(K) \approx 1$  – that is, we are highly confident about  $K$  but not certain. These two assumptions allow me to provide an informal Bayesian argument for the evidential account (for the formal Bayesian argument, see the Appendix, Section 6).

$$P(E|K) \approx 1, P(K) \approx 1 \Rightarrow P(E) \approx 1$$

This inference is licensed by the law of total probability:

$$P(E) = P(E|K) \cdot P(K) + P(E|\neg K) \cdot P(\neg K)$$

Now, if  $P(E) \approx 1$ , then also  $P(E|H) \approx 1$ . This inference is justified if we request that  $P(E) \approx 1$  for all possible values of  $P(H)$ . This line of reasoning has consequences for the posterior probability  $P(H|E)$ :

$$P(H|E) = \frac{P(E|H) \cdot P(H)}{P(E)} \approx P(H)$$

Therefore,  $E$  does not confirm  $H$ , or, at most, to a very small degree. This conclusion is intuitive because the evidence  $E$  is simply not surprising.

Even if  $E$  confirms  $H$  to a very small degree, it is – to say the least – *plausible* for Bayesians to describe the probabilistic version of the Bisphenol A case in the following way: the confirmatory power that  $H$  receives from  $E$  given  $K$  is considerably smaller than the confirmatory power that  $H$  receives from  $E$  if one suppresses  $K$ .<sup>9</sup> We can express this formally as follows:  $P(H|E, K) < P(H|E)$ .

---

<sup>8</sup> For this argument to succeed one might retreat to an even weaker principle proposed by Coney and Feldman (2004: 93) requiring that there be “no more inclusive body of evidence  $[E^*]$  had by [subject]  $S$  at  $t$  such that having [doxastic attitude]  $D$  toward  $p$  does not fit  $[E^*]$ ” (Coney and Feldman 2004: 93; notation altered)

<sup>9</sup> Note that the claim that  $P(H|E, K) < P(H|E)$  is plausible, given our (non-formal) knowledge about this (kind of) case study.

Since the biased researchers suppress information about relevant available background knowledge (particularly concerning what strain of rat is used), they falsely assert that  $P(H|E)$  is the relevant probability for assessing the evidential support that H receives from the experimental results in this case study. Applying the *principle of complete local evidence*,  $P(H|E, K)$  is the relevant probability for assessing evidential support in the context of this case study. Hence, the evidential account is also able to capture the probabilistic version of the Bisphenol A case.

*Frequentist hypothesis testing* is the second major account of evidence and confirmation in the current literature on confirmation in philosophy of science. Statisticians pioneered this approach in the first half of 20<sup>th</sup> century (see, for instance, Fisher 1935; Neyman and Pearson 1967). In this paper, I will focus exclusively on one present-day, relatively non-technical philosophical articulation of frequentist hypothesis testing, Mayo's error-statistical version (Mayo 1996, 2010; Mayo and Spanos 2006). According to Mayo's version, the experimental result (or outcome) E provides evidence for hypothesis H relative to K if and only if H passes severe test T with outcome E (Mayo 1996: 178-80; Mayo 2010: 32; Mayo and Spanos 2006: 328-30). Mayo defines the central concept of passing a severe test T as follows: H "fits"<sup>10</sup> E, and T is "a procedure which would have, at least with very high probability, uncovered the falsity of, or discrepancies from H, and yet no such error is detected" (Mayo 2010: 32).

The experimental design involving insensitive rats of the CD-SD strain in the industry-funded study does not constitute a severe test for the hypothesis that Bisphenol A does not cause cancer, because a test based on this experimental design is not "a procedure which would have, at least with very high probability, uncovered the falsity of" (ibid.) this hypothesis. Indeed, Mayo would classify the test of the hypothesis that Bisphenol A does not cause cancer as a "zero-severity" test, if the test involves rats of the CD-SD strain. She argues that

"such a test is no test at all. It has no power whatsoever at detecting the falsity of H. If it is virtually impossible for H to receive a score less than E on test T, even if false, then H's receiving score E provides no reason for accepting H; it fails utterly to discriminate H being true from H being false" (Mayo 1996: 182).

---

<sup>10</sup> Mayo discusses alternative notions of fit including "H does not fit [E] if [E] is improbable under H" and "evidence [E] might be said to fit H if [E] is more probable under H than under all (or certain specified) alternatives to H" (Mayo 1996: 179; notation altered). A discussion of the merit of these and other notions of fit is however beyond the scope of this paper.

Hence, according to Mayo's error-statistical version of frequentist hypothesis testing, E fails to be evidence for H relative to K.

In sum, I have defended the evidential account of epistemic wrongness with respect to the bisphenol A case by relying on two major philosophical theories of evidential support and confirmation. Let me now turn to the remaining examples of biased research.

In order to motivate my application of the evidential account to the remaining two types of sponsorship bias, I will highlight the relationship between the evidential account and Carrier's false advertising account of, what I call, epistemic wrongness.

Carrier illustrates his account by drawing on the Bisphenol A case (Carrier 2013: 2560-2561, 2017: Sections 4-5, forthcoming: Sections 4-5). This case is an instance of false advertising because the industry-funded scientists claim that the study involving CD-SD rats can be used to show that Bisphenol A does not cause cancer. But, Carrier argues, the biased scientists *falsely assert* that the experimental design employed in the study is appropriate for this use. Carrier diagnoses a "discrepancy between design and use" (Carrier 2017: Section 4).

Carrier's false advertising account and the evidential account agree on a common core idea: epistemic wrongness, at least in the context of research affected by sponsorship bias, should be understood as consisting in the researchers' *false assertions* about a study and its results. By contrast, epistemic wrongness should not be taken to refer to an 'intrinsic' feature of an empirical study itself (for instance, to the experimental design or the obtained data) independently of what researchers assert about their study. To take a concrete example, Carrier and I agree that it is somewhat inaccurate to say that an experimental design involving CD-SD rats is biased. Instead, we take the bias to consist in the fact that researchers *falsely assert* that this particular experimental design can be put to the use of generating data that confirm (or disconfirm) the hypothesis that Bisphenol A causes cancer.

The evidential account can be understood as an attempt to make more precise and to expand the scope of the common core idea just described in two ways: (1) in advocating the evidential account, I try to make more precise Carrier's ambiguous notion of "discrepancy" in terms of false assertions about the actual evidential support that some hypothesis receives from the data (given the principle of complete local evidence). (2) The evidential account is intended to have a broader scope than Carrier's original false advertising account and can thus be understood as an attempt to extend the latter. Carrier's notion of false advertising seems to be restricted to the type of sponsorship bias regarding the choice of experimental design. The evidential account, however, is supposed to

apply to all three types of sponsorship biases presented in Section 2, only one of which concerns a sponsorship bias regarding the choice of experimental design (the Bisphenol A case). In Sections 3.2 and 3.3, I will argue for the claim that the evidential account does indeed apply to two kinds of sponsorship bias that Carrier's approach fails to cover.<sup>11</sup>

### 3.2 The Celebrex case – discussion

From a Bayesian point of view, in the Celebrex case (and cases with a similar structure), there are two possible scenarios of interest:

- (i) the complete data do not confirm H to any degree,
- (ii) the complete data do indeed confirm H, but considerably less than the partial “favorable” data.<sup>12</sup>

In both scenarios, the industry-funded researchers make a false claim about the degree of evidential support of H by E.

In scenario (i), the industry-funded researchers who select and represent the data falsely claim that the complete data are evidence for H. That is, the situation is ultimately analogous to the ‘deductive’ version of Bisphenol A case in one important respect: E does not provide evidential support of H to any

---

<sup>11</sup> A referee asked me to clarify whether the evidential account is committed to “epistemic realism” – the view that there are facts about evidential support relations. I can only provide a tentative answer here. If one bases the evidential account on frequentist hypothesis testing, then facts about evidential support relations seem to be grounded in mind-independent facts about frequencies. This is compatible with epistemic realism. However, I suspect that a Bayesian version of the evidential account is not compatible with epistemic realism, if one takes “facts about evidential support relations” to exclusively refer to mind-independent facts, since all forms of Bayesianism contain a subjective, or cognitive, element. But, even for a Bayesian version of the evidential account, there might nonetheless be facts about evidential support relations, if one allows for a more liberal version of epistemic realism that includes subjective, or cognitive, facts. Bayesians differ on what these cognitive facts are. For subjectivist Bayesians, facts about whether E is supporting evidence for H are determined by facts about the prior probability assigned to H, the laws of probability, a probability space, the application of Bayes’ theorem (and perhaps the principle of complete local evidence). For objectivist Bayesians, facts about evidential support are also determined by further objective principles (such as the principle of indifference, etc.).

<sup>12</sup> Of course, there is also a third possible scenario in which the complete data confirm H to *the same degree* as the partial data do. However, this scenario is uninteresting because it is not a case of biased research.

degree. Despite this analogy, there is, of course, also an important disanalogy between the two cases: in the Bisphenol A case, the bias concerns the choice of experimental design, while the bias regards the selection and representation of data in Celebrex case.

Now consider scenario (ii) that makes the Celebrex case genuinely interesting and challenging. Using the Bayesian notion of comparative confirmation and the principle of complete local evidence (introduced in Section 3.1), it is now possible to defend the evidential account. In Brown's example, the industry-funded study "purported to show that Celebrex caused fewer side effects than older arthritis drugs" (Brown 2008: 193). This hypothesis is simply not justified in light of the available evidence because it rests on the researchers' false assertion that the probability of H given the partial evidence  $E_{\text{partial}}$  (in Brown's example, the data obtained in the first six months of the study) is the relevant probability for assessing the evidential support that H receives from the experimental results of the study. However, applying the principle of complete local evidence, the researchers should have considered the complete local evidence  $E_{\text{complete}}$  (and not  $E_{\text{partial}}$ ) for updating their subjective probability of H but they failed to do so. Thus, the evidential account applies to the Celebrex case.

Let me now apply Mayo's version of frequentist hypothesis testing to the Celebrex case. In Mayo's framework, it is difficult to define a concept of comparative confirmation (unlike in the Bayesian framework). However, it is not necessary to do so to defend the evidential account with respect to Celebrex case. Mayo can apply the principle of complete local evidence introduced above: if one wants to assess whether some hypothesis H receives evidential support from the outcome of a study, then one ought to consider *all* of the data produced by the experiments (the severe tests) involved in that study (and perhaps in a whole series of studies on the same subject matter). Hence, in the Celebrex case, all one needs to look at is the actual complete outcome of severe tests, the actual complete data set  $E_{\text{complete}}$ . If  $E_{\text{complete}}$  is, as a matter of fact, the complete outcome of a severe test procedure T, then, given the principle of complete local evidence,  $E_{\text{complete}}$  (and not  $E_{\text{partial}}$ ) has to be considered for determining whether some hypothesis H passes the severe test T. In the Celebrex case, the hypothesis defended by the industry-funded researchers does, on closer scrutiny, not pass the test.

It might well be true that *if* the outcome of *had been*  $E_{\text{partial}}$  (and not  $E_{\text{complete}}$ , the actual outcome), then the relevant hypothesis *would have* passed the test and, thereby, *would have* gained evidential support from  $E_{\text{partial}}$ . However, this is an assessment of the counterfactual outcome of a study that an advocate of Mayo's error-statistical approach need worry about.



Hence, if one applies Mayo's approach and the principle of complete local evidence to the Celebrex case, then the biased researchers falsely claim that the actual data obtained in the study support the hypothesis in question. This conclusion supports the evidential account.

### **3.3 The Tobacco case – discussion**

In the Tobacco case presented in Section 2.3, the creator of doubt, Sheldon Sommers, claims that a correlation between smoking cigarettes and lung cancer (E) does not, and indeed cannot, provide confirming evidence for the hypothesis that smoking *causes* lung cancer (H), *if* one presupposes a “scientific sense” of “cause”. (Just to recall, according to Sommers, a cause is a necessary and sufficient condition for the occurrence of its effect.) For simplicity's sake, I will refer to this claim as the “Sommers' claim”, which is a particular claim about the (degree of) evidential support that H receives from E.

However, Sommers' claim is false, because its crucial presupposition regarding the “scientific sense” of cause is at odds with our best philosophical theories of confirmation and causation in science.

*First*, Sommers' claim is at odds with extant theories of confirmation, such as Bayesian confirmation theory and Mayo's version of frequentist hypothesis testing. It is a central feature of such theories of confirmation that leads to an inconsistency with Sommers' claim: such theories explicate evidential support in probabilistic terms and they are certainly compatible with correlations playing the role of confirming evidence (indeed, such theories of confirmation are designed specifically for this purpose).<sup>13</sup>

*Second*, according to our best theories of causation in science, there is no good reason to accept Sommers' decisive presupposition that causes in the “scientific sense” have to be necessary and sufficient for their effects.<sup>14</sup> Scientists use causal notions that are not defined in the way that Sommers' assumes. For instance, they use the notion of a probabilistic cause and the notion of a contributing (or partial) cause, independently of the specific smoking-cancer case. The use of these causal concepts (that is, probabilistic and contributing cause) are also reflected and acknowledged in the philosophical literature on causation in the natural and social sciences: for instance, in various probabilistic theories of causation (see Hitchcock 2016 for a survey) and in broadly

---

<sup>13</sup> The same point applies to rules for inferring causal models from correlations as developed by Pearl (2000) and Spirtes et al. (2000).

<sup>14</sup> Proctor (2011: 33-4, 275) advocates a similar line of argument against the claim that causes are necessary conditions for their effects (see also Oreskes and Conway 2010: 33-4).

counterfactual and interventionist theories of causation (for instance, Woodward 2003).

If one accepts the mentioned causal concepts and if one accepts some probabilistic theory of evidence, then it is, at least, conceptually possible that (a) smoking causes lung cancer *and* that (b) some smokers did not get lung cancer (that is, smoking is not a sufficient cause of lung cancer) and that some people who got lung cancer did not smoke (that is, that is, smoking is not a necessary cause of lung cancer). Therefore, the creators of doubt, such as Sommers, falsely assert that correlations between smoking and lung cancer do not count as supporting evidence for the hypothesis that smoking causes lung cancer.<sup>15</sup> Hence, the evidential account applies to the Tobacco case.

### **3.4 A potential counterexample?**

Steel (2018) analyzes a type of sponsorship bias that, at least *prima facie*, poses a challenge to the evidential account. In this section, I will argue that the sort of examples that Steel is interested in can indeed be captured by the evidential account. Hence, these examples do not pose a threat to the evidential account.

Steel argues that in some examples of sponsorship bias the scientists carrying out and publishing the research do not endorse false claims about evidential support relationships, contrary to the evidential account (Steel 2018: 135). The falsehoods (or the systematic distortions or errors) are not located on the side of the researchers and their claims but rather on the side of their audience (that is, on the side of the consumers or users of the research in question). Steel holds that this is what happens if sponsorship bias takes the form of researchers endorsing “misleading claims” about the results of a study. According to Steel’s proposal, misleading claims are “not necessarily false, but are likely to lead to false beliefs among those who encounter them” (Steel 2018: 129; see also p. 119, 135) That is, misleading claims “true or ambiguously interpretable as true” (Steel 2018: 139), and such claims “encourage others” to infer false claims about the results of a study. In the context of medical research, the “others” typically include patients, physicians, and health insurers who have to make decisions based on the published results of studies (Steel 2018: 137).

---

<sup>15</sup> Sommers’ choice of defining a cause in terms of a necessary and sufficient condition for the effect leads to introducing an extremely (perhaps absurdly) high standard of what counts as evidence – that is, only perfect regularities but not correlations may play the role of evidence. Douglas provides a detailed discussion of this problematic kind of raising evidential standards by using the concept of “universal-standards bullshit” (Douglas 2006: 220-221).

Steel's argues that, at least *prima facie*, examples of misleading claims raise complications for the Cochrane definition of sponsorship bias. According to the Cochrane definition, a bias consists in a deviation from the truth (that is, in falsehoods). However, Steel holds that "misleading claims lead to bias not necessarily because they systematically deviate from the truth themselves, but because they encourage others to draw inferences that do." (Steel 2018: 135) For this reason Steel argues: "misleading communication raises complications for such definitions, because misleading claims are naturally regarded as biased but may nevertheless be true." (Steel 2018: 140) Steel's point also carries over to the evidential account and I will discuss it as such.

Consider one of Steel's own examples of a misleading claim appearing in an article about the adverse effects of paroxetine (Paxil). This article reports the results of a study sponsored by GlaxoSmithKline, the financially interested producer of paroxetine:

"The results section of the article reported that 11 serious adverse effects occurred in the Paxil (paroxetine) group compared to only 2 among those treated with a placebo (Keller et al. 2001, 769). One might regard a 450% increase in serious adverse effects as a red flag. Yet the conclusion of the article states, '*Paroxetine is generally well tolerated and effective for treating depression in adolescents*' (Keller et al. 2001, 762)." (Steel 2018: 132; emphasis added)

According to Steel, the conclusion of the article (that Paroxetine is generally well tolerated) is a misleading claim for two reasons: first, the phrase "generally well tolerated" could be "understood to mean that *most* did not suffer from serious adverse effects. Out of 93 patients in the paroxetine group, only 11 or about 12% suffered serious adverse effects, so *most* did not" (Steel 2018: 132-3; emphases added). If one adopts this interpretation, then the sentence "Paroxetine is generally well tolerated" is true. Second, as Steel's characterization of misleading claims requires, the conclusion of the article "encourages others" to infer falsehoods (for instance, that it is safe to use Paxil), even if one chooses the interpretation under which the sentence comes out as true. As Steel stresses, misleading claims figuring in empirical studies exploit "epistemic vulnerabilities" of patients, physicians, and health insurers, because they lack the "key information, analytical skills, or time to adequately process the information presented" (Steel 2018: 139).<sup>16</sup>

---

<sup>16</sup> See Steel (2018: 135-140) for a detailed discussion of such "epistemic vulnerabilities" in terms of "inferential asymmetries".

So, does Steel's example provide a counterexample to the evidential account? I do not think so. I grant Steel that it is possible to describe the example he discusses as an instance of a misleading claim (in his sense of the term). Because of the vagueness of the expression "generally well tolerated" it is indeed *possible* to evaluate the sentence "Paroxetine is generally well tolerated" as true, if one presupposes that the relevant interpretation of "generally well tolerated" is that most patients do not experience severe adverse effects. Let us call this the "most patients"-interpretation of "generally well tolerated". I agree with Steel that this is a *possible* interpretation of the vague sentence "Paroxetine is generally well tolerated". However, I also think there is no good reason to *accept* the "most patients"-interpretation in the scientific context of Steel's example.<sup>17</sup> Instead, a more natural (or at least equally convincing) interpretation of the sentence "Paroxetine is generally well tolerated" can be expressed by the following truth conditions: only a very small percentage of people in the treatment group suffer from severe adverse effects (that is, a percentage considerably smaller than "11 or about 12%" of patients in the treatment group) and there is only a small difference between the treatment group and the control (that is, some difference considerably smaller than a "450% increase in serious adverse effects"). If one adopts these truth conditions (or, this interpretation), then the resulting description of Steel's example is in line with the evidential account: the biased researchers falsely assert that the experimental results of the study provide confirming evidence for the hypothesis that "Paroxetine is generally well tolerated". Thus, Steel's example does not provide a counterexample to the evidential account.

Finally, I propose to classify the sort of examples that Steel discusses simply as an interesting case of a sponsorship bias regarding the choice of a scientific concept used to interpret the results of a study. To opt for the "most patients"-interpretation of "generally well tolerated" is a mistaken (or, at least, questionable) choice of a concept that aligns with the sponsor's financial interests. This choice is in perfect analogy with Dr. Sommers' choice of the "necessary and sufficient for its effect"-interpretation of causation (in the Tobacco case).

---

<sup>17</sup> One problematic feature of Steel's view is that he fails to draw a clear distinction between (i) a vague claim *having* (or *admitting*) a possible interpretation under which it is true and (ii) *accepting* such an interpretation for certain (good) reasons. For instance, Steel seems to accept the "most patients"-interpretation to raise his worry about standard definitions of bias and, indirectly, the evidential account in the context of the paroxetine example. I have to admit that it is unclear to me *why* he accepts this interpretation.

One might worry, as two reviewers did, that there are other examples that might support Steel's view in a better way.<sup>18</sup> For instance, one could imagine an experiment comparing the efficacy of drug A and drug B.<sup>19</sup> Suppose further that the pharmaceutical company funding the research has a financial interest in producing and selling drug A. The study correctly reports that the empirical results strongly support the hypothesis that drug A is more effective than drug B, using a *relative* outcome measure to express the effect size.<sup>20</sup> The effect size according to an *absolute* outcome measure is, however, not reported at all (for a presentation and discussion of such cases Stegenga 2015; Sprenger and Stegenga 2017).

Now, is this imagined scenario a counterexample to the evidential account? Don't the industry-funded researchers report the truth about the results of their study? This scenario does not necessarily constitute a counterexample to the evidential account. To begin with, the scenario might involve a sponsorship bias regarding the choice of the kind of outcome measure used to interpret the data. If so, I think the evidential account is able to capture this kind of sponsorship bias correctly. Such a bias might be a part of the scenario if there is a stark difference between the effect sizes described by the two kinds of outcome measures: for instance, if the effect size according to the relative outcome measure is considerably large (say, 30%) and it is small relative to an absolute outcome measure (say, 3%) (for real cases with this structure see Stegenga 2015: 67; Sprenger and Stegenga 2017: 843). According to the evidential account, the problematic feature of such a scenario consists in the fact that the industry-funded researchers omit a highly relevant piece of available information: here, they omit information about the effect size according to the absolute outcome measure, while they correctly report the effect size according to the relative outcome measure. Clearly, this omission serves to promote their sponsor's financial interests. If the researchers take into account the effect size according to the absolute outcome measure, as they should according the principle of complete

---

<sup>18</sup> Steel (2018: 125-9) also discusses a rather complex case study, criticizing Bero (2014), that involves the strategy of "gaming the dose" in the course of a study. I believe this is another interesting instance of what Steel calls "misleading claims". Addressing this case and others is a task for future research.

<sup>19</sup> One could enrich the scenario by imagining that such an experiment also involves a biased choice of the dose of drug B ("gaming the dose"), as one reviewer suggested. But I will keep things simple here.

<sup>20</sup> Roughly put, outcome measures are formal descriptions of the relationship between the values measured in the experimental group and in the control group of an experimental study (for a detailed formal exposition of different outcome measures, see Stegenga 2015: 65-6; Sprenger and Stegenga 2017: 841-3).

local evidence, then it is false to assert that the results of the study provide strong supporting evidence for the hypothesis that drug A is more effective than drug B (as the industry-funded researchers in our scenario claim). This is just as the evidential account would have it. But why should the researchers consider the effect size according to the absolute outcome measure? If one follows the compelling arguments of Stegenga (2015: 67-8) and Sprenger and Stegenga (2017: sections 3-5), reporting the effect size according to an absolute outcome measure is *the* relevant concept in the context of medical interventions (including assessments of the efficacy of drugs) and, hence, this measure has to be considered when interpreting the results of a study. Thus, if Sprenger and Stegenga are right about the relevance of the absolute outcome measure, then the evidential account is able to capture the kind of scenario just described.

To sum up this section, I have attempted to rebut two potential counterexamples to the evidential account. The first example is drawn directly from Steel's recent work, the second was suggested by two reviewers. Opposing Steel, I have argued that, in both examples, the industry-funded scientists endorse false claims about evidential support relationships, which is in accordance with the evidential account.<sup>21</sup>

To avoid misunderstandings, let me add two clarificatory remarks. *First*, this rebuttal does not force me to reject Steel's plausible diagnosis that a claim such as "Paroxetine is generally well tolerated" encourages "others" (that is, non-experts) to draw false conclusions and that this can be explained by "inferential asymmetries" which industry-funded researchers and their sponsors exploit. Similarly, the evidential account is compatible with empirical research in psychology showing that, as Sprenger and Stegenga (2017: 884-5) point out in the context of the second example, reporting (only) relative outcome measures leads to "proper reasoning fallacies" on the side of patients and physicians (that is, a non-expert audience).

*Second*, if one adopts the evidential account, it is at least plausible to distinguish two distinct projects: (i) providing an account of what is epistemically wrong with the scientific research in the context of both examples discussed in this section, and (ii) providing an explanation of why non-expert audiences draw false conclusions from published scientific work. Providing an account of epistemic wrongness is the intended purpose of the evidential account. By contrast, providing such an explanation is ultimately a scientific rather than

---

<sup>21</sup> Although I believe that my rebuttal of these two potential counterexamples is successful, I acknowledge that it is a fruitful research question to explore and discuss additional kinds of misleading claims and their role in various types of sponsorship bias.

philosophical task (for instance, a task for psychologists and sociologists; Sprenger and Stegenga 2017: 844-5). However, one might even be optimistic that the evidential account could contribute a little piece to an explanation of why non-experts arrive at false conclusions: non-experts might draw false conclusions because they base their reasoning on false information about evidential support relationships after having read the publications of industry-funded scientists.

### 3.5 Two Virtues

Let me present two virtues of the evidential account:

*First*, although many types of sponsorship bias might involve researchers who *intentionally* assert falsehoods about evidential support relations (that is, researchers who lie to their audience), the evidential account does not necessarily depend on knowing the intentions of the biased scientists in question. I take this to be a virtue of the evidential account because it is notoriously difficult to determine what the intentions of scientists are or were (see also Biddle and Leuschner 2015). Surely, there are cases of sponsorship bias where historians and philosophers of science have access to documents revealing the intentions of biased researchers. Research sponsored by the tobacco industry exemplifies a class of paradigmatic cases where such documents are available (Glantz et al. 1996; Proctor 2011). However, acquiring knowledge about the intentions of researchers is often simply beyond our reach.

*Second*, I want to remain neutral with respect to the controversy regarding the value-free ideal – that is, the normative claim that scientists should minimize appeal to non-epistemic (that is, moral, political, and economic) values and only appeal to epistemic values. By claiming neutrality, I mean to say that my argument based on the evidential account should be acceptable for the proponents of the value-free ideal (such as Jeffrey 1956; Levi 1960; Lacey 1999) and its opponents (for instance, Rudner 1953; Longino 1990; Douglas 2000; Wilholt 2009; Kitcher 2011; Biddle and Leuschner 2015). Let me provide a brief sketch of what this might amount to.

When adopting the evidential account, proponents and opponents of the value-free ideal can (and, I believe, should) agree on two points with respect to the examples of sponsorship bias I present and discuss in Sections 2 and 3: (1) the industry-funded researchers in question endorse false claims about evidential support relations and there is good evidence (in these cases) that the non-epistemic interests of the sponsor causally contribute to the fact that the scientists assert such falsehoods. (2) If the sponsors' non-epistemic interests lead to falsehoods of this sort, then we are facing a situation in which non-epistemic interests do not play a legitimate role in science.

However, proponents and opponents of the value-free ideal disagree

about what follows from their two points of agreement. Proponents of the value-free ideal take points (1) and (2) above to vindicate their view that science should be free of non-epistemic values. By contrast, opponents of the value-free ideal would argue that points (1) and (2) do not entail that there is no legitimate place for non-epistemic values in science.

It is not the purpose of the evidential account to settle this disagreement. Instead, the evidential account should be viewed as being logically independent of the controversy regarding the value-free ideal (as a general claim about the role of values in science). I take this neutrality to be another virtue of the evidential account, because it can be adopted by both friends and foes of the value-free ideal.

#### **4. Relation to Other Views**

In this section, I will discuss how the evidential account relates to two prominent views of research affected by sponsorship bias articulated in the recent literature: first, Wilholt's (2009) conventionalist view, and, second, Oreskes and Conway's (2010) account that is based on the notion of disagreement with expert consensus. I will argue for two claims: (1) neither of the two views provides necessary and sufficient conditions for epistemic wrongness; (2) the evidential account should be regarded as a fruitful complement to and not necessarily as a competitor of these two views, at least under certain conditions.

According to Wilholt's (2009) conventionalist view, research affected by sponsorship bias is epistemically wrong – or, an “epistemic failure”, as Wilholt (2009: 92, 99) puts it<sup>22</sup> – because such research violates conventionally accepted methodological standards within a specific research community. For instance, in the Bisphenol A case, the standard is violated that one should not use CD-SD rats if one investigates the effects of Bisphenol A (Wilholt 2009: 99).<sup>23</sup>

Wilholt develops his account of the epistemic wrongness of sponsorship bias under the assumption that the value-free ideal of science has got to be rejected. In other words, he accepts the currently received view that even ‘unbiased’, excellent science (at least sometimes) needs to draw on non-epistemic values, as articulated by opponents of the value-free ideal in the recent

---

<sup>22</sup> Equivalently, Wilholt uses expressions such as “epistemic shortcoming” and “epistemic mistake” (Wilholt 2009: 92, 99).

<sup>23</sup> Biddle and Leuschner hold the conventionalist view as one subcondition of a sufficient condition for “epistemically detrimental dissent” (Biddle and Leuschner 2015: 273; see also Biddle et al. 2017: 167).



debate on science and values (see Section 3.5 for some representative references). As a consequence, Wilholt assumes that an answer to the question “what is epistemically wrong with research affected by sponsorship bias?” cannot simply be that biased research involves reference to non-epistemic values, because the same is true of instances of unquestionably good science. Wilholt argues that his conventional account is compatible with value-laden science and it also captures the “intuition that preference [or sponsorship] bias constitutes *epistemic* failure (rather than just being a matter of differing value judgments)” (Wilholt 2009: 99; emphasis added).<sup>24</sup> Moreover, it is also helpful to make Wilholt’s ‘working assumption’ explicit because it uncovers a difference between Wilholt’s conventional account and the evidential account: the former is committed to rejecting the value-free ideal (at least in the version Wilholt presents), while the latter is intended to be neutral with respect to the value-free ideal (see Section 3.5 above).

Now, suppose one is willing to make the conventionalist assumptions that (a) methodological standards are indeed conventions<sup>25</sup>, and that (b) epistemic wrongness in cases of sponsorship bias consists in the fact that scientists do not conform to some relevant convention. However, even if one agrees with the conventionalists on these assumptions, one is (or, at least, I am) still left wondering what evidence *justifies* the methodological standard in question and why it is *unjustified* not to conform to some methodological standard.

The evidential account is intended to provide an answer to precisely this sort of question – in terms of assessing researchers’ true or false assertions about evidential support relationships. For instance, a proponent of the evidential account might say that it is a justified convention to avoid using CD-SD rats in studies on Bisphenol A (supposing that there is such a convention), because, roughly put, studies involving this strain of rats do not provide any evidence for

---

<sup>24</sup> Although Wilholt accepts that scientific research is value-laden, he argues for an analysis of biases from an *epistemological* perspective. He aims at an account of “epistemic failure” instead of a “counter-intuitive moral analysis of biases” (Wilholt 2009: 95) that would reduce a critical discussion of biases to a purely moral (or political) dispute about diverging non-epistemic value judgments. I agree with Wilholt that the epistemological project is worth pursuing – of course, without dismissing the question of what is morally or politically wrong with sponsorship bias (see footnote 3).

<sup>25</sup> See Wilholt (2009: 98) for a more detailed reflection on the epistemological function of conventional methodological standards in science. In more recent work, Wilholt further elaborates a social epistemological framework that is based on the notion of epistemic trust in research communities (Wilholt 2013: sections 3-4).

addressing the question whether Bisphenol A causes cancer. Hence, I picture the evidential account *not necessarily* as a competitor of the conventionalist view but rather as an independent foundation of the conventionalist view in that it provides an explication of the justification of methodological norms.

Although the evidential account may (at least sometimes) function as a complement to Wilholt's approach, I am not (unlike Wilholt) committed to the claim that the epistemic wrongness *always* consists in the violation of some methodological standard. I believe this is a virtue of the evidential account. Consider the following case as an illustration: violating a convention might not be necessary for epistemic wrongness, because the convention might be unjustified. Imagine the following counterfactual scenario to illustrate an unjustified convention. Suppose it were a convention in medical research on the effects of smoking cigarettes to define a cause as a necessary and sufficient condition for its effect, as Sommers does (Sections 2.3 and 3.3). In such a scenario, the conventionalist is not able to classify Sommers' choice of notion of causation as biased, because he does not fail to conform to any convention in the relevant field of inquiry. However, proponents of the evidential account hold an advantage in this situation because they can raise the question whether this (counterfactual) convention is justified. For them, the most natural strategy is to argue that such a convention is unjustified because, for instance, defining causation in this manner does not cohere with how the notion is fruitfully used in other subfields of medical research and in other scientific disciplines (for reasons outlined in Section 3.3).<sup>26</sup>

Let me now turn to a discussion of Oreskes and Conway's view. At first glance, Oreskes and Conway (2010) *appear* to propose an alternative account of the epistemic wrongness of research that is subject to sponsorship bias. In particular, they appear to focus on analyzing the epistemic wrongness of creation of doubt cases. In the epilogue of their book, Oreskes and Conway argue for a "new view of science" – that is, their positive account of scientific knowledge and justification. This account is mainly based on the notion of expert consensus. For instance, take the following quotes expressing the core idea of the new view of science: science "provides only the consensus of experts, based on the organized accumulations and scrutiny of evidence" (Oreskes and Conway 2010: 268), "there is simply the consensus of expert opinion on [a] particular matter. That is what scientific knowledge is" (*ibid.*), and "what counts as knowledge are

---

<sup>26</sup> See de Melo-Martín and Intemann (2018: 50-3, 68-9) and Carrier (2018: section 3) for further objections to conventionalism.

the ideas that are accepted by the fellowship of experts” (Oreskes and Conway 2010: 269).<sup>27</sup>

Oreskes and Conway use the new view of science to argue that the claims of creators of doubt do not constitute scientific knowledge and are not justified. Hence, *one might be tempted* to interpret Oreskes and Conway as directly proposing the following account of what is epistemically wrong with creation of doubt cases: (a) the creators of doubt are non-experts in the field F of scientific research they attack (for instance, medical research or climate science), (b) the creators of doubt oppose the expert consensus in field F (for instance, the expert consensus that smoking causes lung cancer), and (c) we have good reason to believe that the creators of doubt have non-epistemic interests (such as financial interests or an anti-communist<sup>28</sup> political agenda) increasing the likelihood of systematic errors (Oreskes and Conway 2010: 271-3).

I partly agree with Oreskes and Conway: if someone is a non-expert in some field of inquiry, s/he opposes the expert consensus in this field, and s/he has potentially biasing non-epistemic interests, then her/his arguments should be reviewed with particular suspicion and care. However, I think that Oreskes and Conway should *not* be taken to provide an account of, what I call, epistemic wrongness. If one decided to read Oreskes and Conway as presenting an account of the epistemic wrongness, this reading would seem to exclude – on a priori grounds – the possibilities that (i) non-experts might be able to contribute valuable non-biased research that runs counter to the expert consensus a field of inquiry, and that (ii) even experts might be biased. To adopt a reading of Oreskes and Conway’s view that excludes such possibilities on a priori grounds strikes me as mistaken, because it is overly restrictive. Moreover, if understood as an account of epistemic wrongness, Oreskes and Conway’s approach appears to be limited to (certain) creation of doubt cases, because condition (a) above (the lack of expertise in the relevant field F) does not seem to be violated in the other two types of sponsorship bias discussed in this paper.<sup>29</sup>

For these reasons, I believe a better and more charitable interpretation of Oreskes and Conway’s view is that disagreement with expert consensus

---

<sup>27</sup> Oreskes and Conway’s new view of science should be distinguished from another topic also discussed in the epilogue of their book: trust in experts. Their concept of informed or rational trust in experts is supposed to describe the epistemic attitude non-experts should adopt w.r.t. scientific experts (Oreskes and Conway 2010: 271-4). However, unlike the “new view of science”, their account of trust is not constitutive of criticizing creators of doubt as being epistemically wrong.

<sup>28</sup> See, for instance, Oreskes and Conway (2018: 27-8, 186, 248-9).

<sup>29</sup> See de Melo-Martín and Intemann (2018: 53-8) for additional objections.

(especially when in combination with lack of expertise and having certain non-epistemic interests) is a valuable *indicator* of epistemic wrongness rather than an analysis or definition of epistemic wrongness itself. For present purposes, I take an indicator to be a piece of information increasing our confidence that there is something epistemically wrong with the results of some study.

To be sure, there is no contradiction between this interpretation of Oreskes and Conway's view and the evidential account. If one takes a closer look at Oreskes and Conway's detailed arguments to the conclusion that industry-funded tobacco research is epistemically flawed (Oreskes and Conway 2010: 32, 241, 268), they typically identify cases of cherry picking (for instance, Oreskes and Conway 2010: 18, 187-8, 241) and instances of sponsorship biases regarding the choice of scientific concepts used to interpret experimental results, particularly concerning concepts such as causation, evidence, and uncertainty (for instance, Oreskes and Conway 2010: 30-34, 192-4).

My constructive proposal is to distinguish more clearly between two issues that are not sufficiently disentangled in Oreskes and Conway's work – that is, specifying indicators of epistemic wrongness and clarifying the concept of epistemic wrongness. If one interprets Oreskes and Conway as providing a helpful strategy for identifying *indicators* of epistemic wrongness, then the evidential account can make use of these indicators as reliable guides to detecting instances of sponsorship bias. However, it is important to stress that the evidential account has a different task than to identify indicators of epistemic wrongness. The goal of the evidential account is conceptual clarification – that is, to explicate what epistemic wrongness is. Thus, I take the evidential account to play the role of a fruitful complement to Oreskes and Conway's approach.

## **5. Conclusion**

I have argued for an evidential account of what is epistemically wrong with research affected by sponsorship bias. According to the evidential account, research affected by sponsorship bias is epistemically wrong if and only if the researchers in question endorse false claims about evidential support relationships holding between some hypothesis and some body of empirical data. My focus was on examples of three paradigmatic types of sponsorship bias (regarding the choice of experimental design, of the selection and representation of data, and of the scientific concepts used to interpret the results of a study). I have defended the claim that the evidential account provides a satisfactory analysis of what is epistemically wrong with all three types of sponsorship bias.


I have presented several advantages of the evidential account: it is neutral with respect to two major theories of confirmation (Bayesianism and frequentism), its applicability does not depend on knowing the intentions of biased researchers, and it is neutral with respect to the value-free ideal. I have argued that the evidential account has a broader scope than competing accounts and that it may play, at least under certain conditions, a complementary role for Wilholt’s (2009) view and Oreskes and Conway’s (2010) account.

One fruitful question for future research is to assess whether the evidential account can also be applied to further types of sponsorship bias (for instance, types resting on p-hacking and misleading claims) – and perhaps to other kinds of biases in science (such as publication bias in publicly funded and industry-funded research). In sum, I believe that the evidential account is promising and deserves further attention and discussion.

## 6. Appendix: The Formal Bayesian Argument

In Section 3.1, I provided an informal Bayesian argument for the evidential account in the context of the Bisphenol A case. Here, I will present a more formal Bayesian argument to the same conclusion.

To begin with, we represent the relationship between the variables E, H and K in the following Bayesian network:



That is, we assume that H and K are unconditionally independent, but become correlated if E is given. This is a highly plausible assumption.

Next, we suppose again that  $P(K) = k \approx 1$  and set  $P(H) = h$  to some value between 0 and 1. To fully specify the prior probability distribution over the Bayesian network for the Bisphenol A case, we also assume that  $P(E|K, H) \approx P(E|K, \neg H) \approx 1$ , which follows from the requirement that  $P(E|K) \approx 1$  holds for all possible values of H, and set

$$P(E|\neg K, H) = \beta$$

$$P(E|\neg K, \neg H) = \delta$$

Let us now compute the posterior probability  $P(H|E)$ , i.e.

$$P(H|E) = \frac{P(H, E)}{P(E)}.$$

We can determine  $P(H, E)$  and  $P(E)$  as follows:<sup>30</sup>

$$\begin{aligned} P(H, E) &= P(H, E, K) + P(H, E, \neg K) \\ &\approx h \cdot k + h \cdot \bar{k} \cdot \beta \\ &= h \cdot (k + \bar{k} \cdot \beta) \end{aligned}$$

$$P(E) \approx k + \bar{k} \cdot (h \cdot \beta + \bar{h} \cdot \delta)$$

When substituting for  $P(H, E)$  and  $P(E)$ , we obtain:

$$P(H|E) \approx h \cdot \frac{k + \bar{k} \cdot \beta}{k + \bar{k} \cdot (h \cdot \beta + \bar{h} \cdot \delta)}$$

It is interesting to note that the fraction in the above equation converges to 1 as  $k = P(K) \rightarrow 1$ . Hence,  $P(H|E) \approx P(H)$  in this limit. This result coincides with the conclusion of the informal Bayesian argument: E does not confirm H, or, at most, to a very small degree.

## Acknowledgements

First and foremost, I would like to thank Stephan Hartmann, Maria Kronfeldner, and Anna Leuschner for very constructive and repeated feedback on different drafts. The comments of the editors and two referees helped a lot to improve the quality of the manuscript. Moreover, I did benefit from discussions with Martin Carrier, Mark Colyvan, Michael Cuffaro, Branden Fitelson, Molly Kao, Martin Kusch, Jürgen Landes, Christoph Merdes, Kim Naumann, Barbara Osimani, Lydia Patton, Franziska Reinhard, Jan Sprenger, Daniel Steel, Reuben Stern, Michael Strevens, Gregory Wheeler, and kind audiences in Cologne, Herrsching, and Munich.

---

<sup>30</sup> We use the following notation:  $\bar{x} = 1 - x$ .

## References

- Bero, L.A., and Rennie, D. (1996): “Influences on the Quality of Published Drug Studies”, *International Journal of Technology Assessment in Health Care* 12(2): 209-237.
- Biddle, Justin (2007): “Lessons from the Vioxx Debacle: What the Privatization of Science Can Teach Us About Social Epistemology,” *Social Epistemology* 21: 21-39.
- Biddle, Justin, Kidd, Ian, and Leuschner, Anna (2017): “Epistemic Corruption and Manufactured Doubt: The Case of Climate Science”, *Public Affairs Quarterly* 31: 165-187.
- Biddle, Justin and Leuschner, Anna (2015): “Climate Skepticism and the Manufacture of Doubt: Can Dissent in Science be Epistemically detrimental?”, *European Journal for Philosophy of Science* 5: 261-278.
- Bovens, Luc and Hartmann, Stephan (2003): *Bayesian Epistemology*, New York: Oxford University Press.
- Brown, James Robert (2008): “The Community of Science®”, In M. Carrier, D. Howard and J. Kourany (eds.), *The Challenge of the Social and the Pressure of Practice*, Pittsburgh: University of Pittsburgh Press, 189-216.
- Carnap, Rudolf (1950): *Logical Foundations of Probability*, Chicago: The University of Chicago Press.
- Carrier, Martin (2018): “Identifying Agnotological Ploys: How to Stay Clear of Unjustified Dissent”, In A. Christian et al. (eds.), *Philosophy of Science – Between Natural Science, the Social Sciences, and the Humanities* (European Studies in Philosophy of Science), Dordrecht: Springer, pp. 155-169.
- Carrier, Martin (2017): “Agnotological Challenges: How to Capture the Production of Ignorance in Science”, unpublished manuscript.
- Carrier, Martin (2013): “Values and Objectivity in Science: Value-Ladenness, Pluralism and the Epistemic Attitude”, *Science & Education* 22: 2547-2568.
- Conee, Earl and Feldman, Richard (2004): *Evidentialism. Essays in Epistemology*, New York: Oxford University Press.
- de Melo-Martín, I. and K. Intemann (2018): *The Fight Against Doubt. How to Bridge the Gap Between Scientists and the Public*, New York: Oxford University Press.
- de Melo-Martín, I. and K. Intemann (2014): Who’s Afraid of Dissent? Addressing Concerns about Undermining Scientific Consensus in Public Policy Developments, *Perspectives on Science* 22: 593-615.
- Doucet, M., and Sismondo, S. (2008): “Evaluating Solutions to Sponsorship

- Bias”, *Journal of Medical Ethics* 34: 627-630.
- Douglas, Heather (2006): “Bullshit at the interface of science and policy: global warming, toxic substances and other pesky problems”, In *Bullshit and Philosophy*, Hardcastle, G.L. and Reisch, G.A. (eds.), 213-226. Chicago: Open Court.
- Douglas, Heather (2000): “Inductive Risk and Values”, *Philosophy of Science* 67: 559-579.
- Earman, John (1992): *Bayes or Bust? A Critical Examination of Bayesian Confirmation Theory*, Cambridge, MA: MIT Press.
- Elliott, Kevin (2016): “Standardized study designs, value judgments, and financial conflicts of interest”, *Perspectives on Science* 24: 529-551.
- Fernández Pinto, Manuela (2017): “Agnotology and the New Politicization of Science and Scientization of Politics”, In *The Routledge Handbook of the Political Economy of Science*, D. Tyfield, R. Lave, S. Randalls, and C. Thorpe (eds.), pp. 341-350, London: Routledge.
- Fisher, Ronald Aylmer (1935): *The Design of Experiments*, Edinburgh: Oliver and Boyd.
- Fitelson, Branden and Hawthorne, James (2010): “The Wason Task(s) and the Paradox of Confirmation”, *Philosophical Perspectives* 24: 207-241.
- Glantz, Stanton, John Slade, Lisa Bero, Peter Hanauer, and Deborah Barnes (1996): *The Cigarette Papers*, Berkeley: University of California Press.
- Harker, David (2015): *Creating Scientific Controversies. Uncertainty and Bias in Science and Society*, Cambridge: Cambridge University Press.
- Hitchcock, Christopher (2016): “Probabilistic Causation”, *The Stanford Encyclopedia of Philosophy* (Winter 2016 Edition), Edward Zalta (ed.), URL = <<https://plato.stanford.edu/archives/win2016/entries/causation-probabilistic/>>.
- Holman, Bennett, and Elliott, Kevin (2018): “The Promise and Perils of Industry-funded Science”, *Philosophy Compass* 13.11, DOI: 10.1111/phc3.12544.
- Horwich, Paul (1982): *Probability and Evidence*, Cambridge: Cambridge University Press.
- Howson, Colin and Urbach, Peter (2006): *Scientific Reasoning. The Bayesian Approach*, 3rd edition, La Salle, IL: Open Court.
- Jeffrey, Richard (1956): “Valuation and Acceptance of Scientific Hypotheses”, *Philosophy of Science* 23: 237-246.
- Kitcher, Philip (2011): *Science in a Democratic Society*, Amherst: Prometheus.
- Kourany, Janet (2016): “Should Some Knowledge be Forbidden? The Case of Cognitive Difference Research”, *Philosophy of Science* 83: 1-12.
- Lacey, H. (1999): *Is Science Value Free?*, New York: Routledge.
- Levi, Isaac (1960): “Must the Scientist Make Value Judgments?”, *The Journal of*



- Philosophy* 57: 345-357.
- Lewandowsky, S., J. Cook and E. Lloyd (2018): “The ‘Alice in Wonderland’ Mechanics of the Rejection of (Climate) Science: Simulating Coherence by Conspiracism”, *Synthese* 195: 175–196.
- Lexchin, J., Bero, L. A., Djulbegovic, B., and Clark, O. (2003): “Pharmaceutical Industry Sponsorship and Research Outcome and Quality: Systematic Review”, *BMJ* 326: 1167-1170.
- Longino, H. (1990): *Science As Social Knowledge*, Princeton: Princeton University Press.
- Mayo, Deborah (2010): “Learning from Error, Severe Testing, and the Growth of Theoretical Knowledge”, In D. Mayo and A. Spanos (eds.): *Error and Inference: Recent Exchanges on Experimental Reasoning, Reliability and the Objectivity and Rationality of Science*, Cambridge: Cambridge University Press: 28-57.
- Mayo, Deborah (1996): *Error and the Growth of Experimental Knowledge*, Chicago: The University of Chicago Press.
- Mayo, Deborah and Spanos, Aris (2006): “Severe Testing as a Basic Concept in a Neyman-Pearson Philosophy of Induction”, *The British Journal for the Philosophy of Science* 57: 323–357.
- McGarity, T., and Wagner, W. (2008): *Bending Science: How Special Interests Corrupt Public Health Research*, Cambridge, MA: Harvard University Press.
- Michaels, David (2008a): *Doubt is Their Product. How Industry’s Assault on Science Threatens Your Health*, New York: Oxford University Press.
- Michaels, David (2008b): “Manufactured Uncertainty. Contested Science and the Protection of the Public’s Health and Environment”, In: R. Proctor and L. Schiebinger (2008), 90-107.
- Neyman, Jerzy and Pearson, Egon Sharpe (1967): *Joint Statistical Papers*, Berkeley and Los Angeles: University of California Press.
- Oreskes, Naomi and Conway, Erik (2010): *Merchants of Doubt. How a Handful of Scientists Obscured the Truth on Issues from Tobacco Smoke to Global Warming*, New York: Bloomsbury Press.
- Pearl, Judea (2000): *Causality. Models, Reasoning, and Inference*, Cambridge: Cambridge University Press.
- Proctor, Robert (2011): *Golden Holocaust. Origins of the Cigarette Catastrophe and the Case for Abolition*, Berkeley and Los Angeles: University of California Press.
- Proctor, Robert (2008): “Agnotology. A Missing Term to Describe the Cultural Production of Ignorance (and its Study)”, In *Agnotology. The Making and Unmaking of Ignorance*, R. Proctor and L. Schiebinger (eds.), Stanford: Stanford University Press, 1-33.

- Proctor, Robert and Schiebinger, Londa (eds.) (2008): *Agnology. The Making and Unmaking of Ignorance*, Stanford: Stanford University Press.
- Radder, Hans (ed.) (2010): *The Commodification of Academic Research. Science and the Modern University*, Pittsburgh: The University of Pittsburgh Press.
- Resnik, D. (2000): “Financial Interests and Research Bias”, *Perspectives on Science* 8: 255-285.
- Rudner, Richard (1953): “The Scientist Qua Scientist Makes Value Judgments”, *Philosophy of Science* 20: 1-6.
- Shrader-Frechette, Kristin (1994): *Ethics of Scientific Research*. Boston: Rowman and Littlefield.
- Sismondo, Sergio (2008): “How Pharmaceutical Industry Funding Affects Trial Outcomes: Causal Structures and Responses”, *Social Science and Medicine* 66: 1909–1914.
- Sommers, Sheldon (1985): “Galbraith Deposition Transcript of Sheldon C. Sommers”, 1985 October 31, RJ Reynolds Records, <https://www.industrydocumentslibrary.ucsf.edu/tobacco/docs/mqpd00822018>
- Spirtes, Peter, Clark Glymour, and Richard Scheines (2000): *Causation, Prediction and Search*, Second Edition, Cambridge, MA: M.I.T. Press.
- Sprenger, Jan, and Hartmann, Stephan (2019): *Bayesian Philosophy of Science*, Oxford: Oxford University Press.
- Sprenger, Jan, and Stegenga, Jacob (2017): “Three arguments for absolute outcome measures”, *Philosophy of Science* 84: 840-852.
- Steel, Daniel (2018): “If the Facts were not Untruths, their Implications were: Sponsorship Bias and the Art of Misleading Communication”, *Kennedy Institute of Ethics Journal* 28: 119-144.
- Stegenga, Jacob (2018): *Medical Nihilism*, New York: Oxford University Press.
- Stegenga, Jacob (2015): “Measuring effectiveness”, *Studies in History and Philosophy of Biological and Biomedical Sciences* 54: 62-71.
- Vom Saal, Frederick and Hughes, Claude (2005): “An Extensive New Literature Concerning Low-Dose Effects of Bisphenol A Shows the Need for a New Risk Assessment”, *Environmental Health Perspectives* 113: 926-931.
- Wilholt, Torsten (2013): “Epistemic Trust in Science”, *The British Journal for the Philosophy of Science* 64: 233-253.
- Wilholt, Torsten (2009): “Bias and Value in Scientific Research”, *Studies in the History and Philosophy of Science* 40: 92-101.
- Woodward, James (2003): *Making Things Happen. A Theory of Causal Explanation*, New York: Oxford University Press.