Philosophy of Science (2026), **93**, 675–695
doi:10.1017/psa.2026.10195

PSA PHILOSOPHY OF SCIENCE ASSOCIATION
PHILOSOPHY OF SCIENCE

ARTICLE

Unreasonable Doubt. How Strategic Science Skeptics Exploit the Argument from Disagreement

Alexander Reutlinger

Munich Center for Mathematical Philosophy, Faculty of Philosophy, Philosophy of Science, and Religious Studies, Ludwig-Maximilians-Universität München, Munich, Germany
Email: Alexander.Reutlinger@lmu.de

(Received 01 June 2025; revised 27 December 2025; accepted 02 January 2026)

Abstract

Strategic science skeptics criticize scientific claims solely to promote non-epistemic goals. I will analyze and debunk a philosophically neglected argument exploited by strategic science skeptics: the argument from disagreement. The core of this argument is that one should lower one's confidence in a scientific claim when having learned that there is a scientific disagreement about this claim. I will develop a (Bayesian) *Justificatory Account of Multiple Testimony* to provide a normative characterization of how learning about agreements and disagreements is connected to confirming and disconfirming scientific claims. I will use this account to debunk the argument from disagreement.

1. Introduction: Strategic science skepticism and the argument from disagreement

Suppose you are not a scientific expert, but you know that scientists do research to test a particular scientific claim. Suppose further that you, then, come to know that there is disagreement about that claim in science. How should you react? You should be less confident about the claim than you were before. At least, this advice—to lose confidence in the claim—is the core idea of an argument that strategic science skeptics use to undermine scientific claims: the argument from disagreement. In this paper, my aim is to analyze and to debunk this argument of strategic science skeptics that has been neglected as a subject for thorough philosophical analysis.

Before delving into an analysis of the argument, it is helpful to situate the argument in the context of a larger phenomenon: strategic science skepticism and the claims and arguments of strategic science skeptics. Described as a general phenomenon, strategic science skeptics contest and criticize scientific claims solely

© The Author(s), 2026. Published by Cambridge University Press on behalf of Philosophy of Science Association. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (https://creativecommons.org/licenses/by/4.0/), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

https://doi.org/10.1017/psa.2026.10195 Published online by Cambridge University Press

to promote economic interests and political agendas.[1] Skepticism of this sort targets research in climate science, environmental science and health science—that is, strands of research relevant for the industry and for policy-makers.[2] Consider a well-known example for illustration.[3] According to climate scientists, there is very good evidence that human CO_2 emissions are a major cause of global warming. Yet, strategic science skeptics openly disagree with climate scientists. And strategic science skeptics present arguments for why they disagree, but they disagree for strategic reasons—to achieve certain economic and political goals of their own and/or their sponsors (for instance, the goal of the fossil fuel industry to sell oil and to prevent political regulation of their business). The intended effect of the strategic science skeptics' claims is to create public distrust in climate science.

Strategic science skeptics use a whole range of arguments to contest scientific claims. One pervasive class of skeptical arguments is directly concerned with the *empirical* evidence for scientific claims. Relying on such arguments targeting empirical evidence, strategic science skeptics raise objections to scientific research by *criticizing* that the empirical evidence—produced by scientists' research—does not support the scientists' hypotheses. Science skeptics also actively engage in *doing research* to produce their own (alleged) counterevidence to publicly funded research. Skeptical arguments directed at empirical evidence are implemented in various ways: for instance, by cherry-picking the empirical evidence, by choosing biased research methods to obtain empirical evidence, and by refining concepts used to describe and interpret empirical evidence. Philosophical reflections on strategic science skepticism have been preoccupied with skeptical arguments targeting *empirical evidence*.[4]

The skeptical argument taking center stage in this paper is strikingly different. The argument from disagreement is not directly concerned with empirical evidence. Instead, the argument operates in the domain of the social epistemology of science: it is concerned with a specific form of *testimonial evidence*. The argument rests on assumptions about how justification depends on the awareness of disagreement and, respectively, agreement among the testimonies (or reports) of scientists. While the skeptics' own arguments targeting empirical evidence are often the means to generate disagreements, the skeptical argument in the focus of this paper is concerned with drawing a conclusion from a disagreement *after* it has been generated. The key idea underlying this argument from disagreement is that someone who is not

[1] I will use the label "strategic science skepticism" throughout the paper. Sometimes I will use "skeptics" as shorthand for "strategic science skeptics." For present purposes, I will use this label as being synonymous with alternative established labels such as "merchandizing," "manufacturing," and "creating doubt" (Michaels 2008a; Oreskes and Conway 2010) and "agnogenesis" in the sense of "agnotological ploy," the deliberate and strategic creation of public ignorance about scientific knowledge (Proctor and Schiebinger 2008; Proctor 2011; Kourany and Carrier 2020).

[2] Historians of science produced a vast number of elaborate case studies of strategic science skepticism, launched by the tobacco industry, contesting publicly funded biomedical research that provided evidence for a causal link between smoking cigarettes and various diseases including lung cancer (Glantz et al. 1996; Proctor 1995, 2011; Brandt 2007; Michaels 2008a; Proctor and Schiebinger 2008; Oreskes and Conway 2010, chapters 1 and 5).

[3] Specifically on strategic science skepticism regarding climate science, see, for instance, Michaels (2008a, 2008b); Oreskes and Conway (2008; 2010, chapter 6); Howe (2014); Supran and Oreskes (2017).

[4] For instance, Douglas (2006, 2009); Wilholt (2009); Biddle and Leuschner (2015); Harker (2015); Carrier (2018); Steel (2018).

https://doi.org/10.1017/psa.2026.10195 Published online by Cambridge University Press

a scientific expert, a lay person, should lose confidence in the truth of a scientific claim once that person has come to know that there is a disagreement about the claim. I will take a *normative stance*: I will analyze and debunk this argument. That is, I will show why it is flawed.[5]

More precisely, the plan of the paper is as follows. In section 2, I will propose a reconstruction of the argument from disagreement as a deductively valid inference from two premises to a conclusion. In section 3, I will develop what I call the *Justificatory Account of Multiple Testimony* that will serve as the basis for debunking the argument. This account provides a normative characterization of how learning about agreements and disagreements is connected to confirming and disconfirming scientific claims. I will draw on Bayesian confirmation theory as a framework for articulating a precise version of this account. For this purpose, I treat Bayesian confirmation theory (i) as a *tool* for debunking the skeptics' argument (that is, I will apply, but not attempt to defend or improve, Bayesianism), (ii) as a *normative* theory of how confirmation should work (aligning with the general normative stance of this paper), and (iii) as *one* possible normative framework[6] for articulating the *Justificatory Account of Multiple Testimony* and for debunking the skeptics' argument. In sections 4 and 5, I will use the Bayesian version of the *Justificatory Account of Multiple Testimony* to argue that both premises of the skeptics' argument from disagreement are false. Hence, the skeptics' argument is flawed because it is not sound. Section 6 presents the conclusions of my arguments.

2. Analyzing the argument from disagreement

In its perhaps most iconic form, the argument from disagreement is stated in a confidential, internal document of the tobacco industry, the infamous "Smoking and Health Proposal" (Brown and Williamson 1969). Indeed, historians of science regard this document as *the* paradigmatic case of strategic science skepticism and as *the* model that guided strategic science skeptics in extending skepticism from biomedical research on the health effects of smoking cigarettes to other sciences (including environmental and climate sciences).[7] Moreover, the key phrase of the document—"doubt is our product"—has inspired the analytic terminology of doubt mongering, of producing doubt, and of ignorance that historians of science and subsequently philosophers have begun to use (see footnotes 1–4 for references).

In the document, a public relations strategy is proposed allowing the tobacco industry to respond to biomedical research on the adverse health effects of smoking—research that is the basis of regulation of and lawsuits against the tobacco industry. The strategy is abstractly characterized as follows: "Our consumer I have

[5] The work of some scholars addresses topics in the vicinity of the argument from disagreement, including Oreskes and Conway (2010), Odenbaugh (2012), and Harker (2015, chapter 7). However, what is still missing and what I intend to provide is an in-depth critical analysis of the argument.

[6] It will be a rewarding task for future research to explore alternatives to the Bayesian framework and to assess their (comparative) merit. Such alternatives include other theories of confirmation (such as likelihoodism and frequentism; for a survey, see Sober 2008, chapter 1), and various accounts of the "wisdom of crowds," including approaches in social choice theory (Dietrich and Spiekermann 2020).

[7] This memo is well documented by historians of science, such as Glantz et al. (1996, 171); Brandt (2007, 210); Christensen (2008, 268 and 280–81n2); Michaels (2008a, x and 11; 2008b, 91); Proctor (2008, 1 and 17; 2011, 289 and 617n1); and Oreskes and Conway (2010, 34).

https://doi.org/10.1017/psa.2026.10195 Published online by Cambridge University Press

defined as the mass public, our product as doubt, and our message as truth well stated, and our competition as the body of anti-cigarette fact that exists in the public mind" (Brown and Williamson 1969, 3–4).[8] The document offers advice on how to handle a situation in which the "mind of the general public" (Brown and Williamson 1969, 4) is in the grip of "anti-cigarette facts" and "misinformation" (ibid.) stemming from biomedical research and supporting, for instance, the claim that smoking causes lung cancer. To change the "mind of the general public" (ibid.), it is suggested to create "doubt" about the relevant scientific information:

> Doubt is our product since it is the best means of competing with the 'body of fact' that exists in the mind of the general public. It is also the means of establishing controversy. Within the business we recognize that a controversy exists. However, with the general public the consensus is that cigarettes are in some way harmful to the health. If we are successful in establishing a controversy at the public level, then there is an opportunity to put across the real facts about smoking and health. (Brown and Williamson 1969, 4)

The advice is to establish a "controversy" via raising "doubt" about those scientific claims that are troublesome for the economic interests of the tobacco industry. Questioning and criticizing scientific claims is instrumental to creating the impression in the "mind of the general public" that there is a disagreement or controversy about the scientific claims. In a first approximation, this strategic advice attempts to exploit a connection between (i) the public becoming aware of there being a "controversy" or a disagreement—instead of an agreement—about a scientific claim within science, and (ii) as a response to this awareness, an increased doubtfulness or, equivalently, a decreased confidence regarding that scientific claim in the "mind of the general public."

The second prominent instance of the argument from disagreement is exemplified by the advice of Frank Luntz, a political consultant of the Republican Party in the United States.[9] Luntz writes the following about research in climate science to promote the Republican political agenda (that includes protecting the business interests of the fossil fuel industry and preventing its regulation): "*The scientific debate remains open.* Voters believe that there is *no consensus* about global warming within the scientific community. Should the public come to believe that the scientific issues are settled, their views on global warming will change accordingly" (quoted after Michaels 2008a, 198; original emphases).

Luntz discusses a situation that differs from the problem that the "Smoking and Health Proposal" addresses. As described above, the tobacco industry faced a public opinion according to which there is agreement within biomedical sciences regarding the effect of smoking—and the skeptics' aim was to raise doubt about the existence of such a consensus within science. By contrast, Luntz's strategic advice intends to

[8] Proctor (2011, 617n1) addresses the question of to whom the first personal pronoun in the quoted passage refers.

[9] The Luntz case is referred to, for instance, by Michaels (2008a, xi and 198; 2008b, 92); Oreskes and Conway (2008, 74 and 79); Proctor (2008, 33n53).

https://doi.org/10.1017/psa.2026.10195 Published online by Cambridge University Press

perpetuate an already existing public impression of disagreement within climate science and to prevent a public perception of scientific agreement from emerging.

To prevent that the public "comes to believe that the scientific issues are settled," Luntz suggests creating the impression that the research results of climate science are not settled yet and that, instead, there is a still "open" scientific debate including disagreements and "a lack of scientific certainty" within the climate scientific community (Michaels 2008a, xi, 198; 2008b, 91–93). Luntz draws on the same connection between public awareness of disagreement and decreased confidence of the public that is already implicit in Brown and Williamson (1969). Interestingly, Luntz frames this connection in terms of *belief*. On the one hand, if the public perceives scientists to agree on a scientific claim, then the public will strongly believe it (i.e. be confident about the truth of the claim). On the other hand, if the public thinks that there is scientific disagreement about a claim, then "their views on global warming will change accordingly"—that is, the public will disbelieve the claim at issue or believe it less strongly.

The skeptics do not present their argument from disagreement in a standard premise–conclusion form. I propose the following reconstruction:

1. **Disagreement Premise:** A member of the public learns that there is disagreement about a scientific hypothesis H.[10]
2. **The Doubt Rule:** If a member of the public learns that there is disagreement about a scientific hypothesis H, then this person should be less confident about the truth of H (that is, less confident than before having learned about the disagreement).
3. **Conclusion:** Thus, the member of the public should be less confident about the truth of H.

Let me add two clarificatory remarks about this reconstruction before moving on to describing the challenge the argument poses.

The first remark concerns the expression "member of the public" in my reconstruction. I use an individualistic approach by referring to individuals who are members of the public instead of to a collective subject as suggested by the skeptics' expressions such as "the mind of the general public." I choose this approach, because, for present purposes, it is not necessary to commit to any substantive talk about collective agents and their mental states.

The second remark concerns room for interpretation in reconstructing the second premise. The *Doubt Rule* allows for different possible versions:[11]

(i) **Disbelief version:** If a member of the public learns that there is disagreement about a scientific hypothesis H, then this person should *disbelieve* H (i.e. believe that $\neg H$ is true).

[10] I will use the concept of a hypothesis as a neutral umbrella term covering theories, models, and other theoretical modes of scientific representation with empirical content.

[11] I do not intend this list to exhaust all possible interpretations.

https://doi.org/10.1017/psa.2026.10195 Published online by Cambridge University Press

(ii) **Suspension-of-judgment version:** If a member of the public learns that there is disagreement about a scientific hypothesis *H*, then this person should *suspend judgment* about the truth (and falsity) of *H*.
(iii) **Low-confidence version:** If a member of the public learns that there is disagreement about a scientific hypothesis *H*, then this person should have *low confidence* in the truth of *H*.

The version of the *Doubt Rule* I choose in the reconstruction above is what one might call the *lowering-confidence version*. It demands that a member of the public should—having learned about a disagreement regarding hypothesis *H*—become *less confident* about the truth of *H*. I take the *lowering-confidence version* to be the weakest version of the *Doubt Rule*. In other words, I assume that it is an easier task for strategic science skeptics to defend the *lowering-confidence version* than to defend the three other possible versions of the *Doubt Rule*. I will work with this version of the *Doubt Rule*, because it constitutes the most charitable way of reconstructing the skeptics' argument from disagreement. If one rebuts the weakest version of the *Doubt Rule*, as I intend to do, then its stronger versions become untenable in one sweep.

The challenging character of the skeptics' argument from disagreement arises from the central role of agreement (or consensus)[12] as a marker for credible and justified scientific claims. Indeed, many historians and philosophers of science accept such a central and productive epistemic role of agreement in science and take it as the subject of their analyses.[13] To dramatize the challenge, strategic science skeptics who draw on the argument from disagreement can happily accept the view that agreement is central to good science and to the credibility of scientific claims. Then, the skeptics diagnose that there is disagreement—hence, no agreement—about a specific scientific claim (such as the claim that smoking causes lung cancer). From the skeptics' point of view, this is all that is needed to apply the argument from disagreement to a scientific claim they would like to undermine.

I will explore a way of debunking the argument from disagreement. Indeed, Oreskes and Conway present a glimpse of such an argumentative strategy: "If the scientific community has been asked to judge a matter (as the National Academy of Sciences routinely is)—or if they have self-organized to do so (as in the Ozone Trends Panel or the IPCC), then it makes sense to take the results of their investigations very seriously. [. . .] *It does not make sense to dismiss them just because some person, somewhere, doesn't agree*" (Oreskes and Conway 2010, 272–73; emphasis added). I read this passage as follows: Oreskes and Conway would like to maintain that agreement plays a crucial and productive epistemic role in science. At the same time, they would like to reject the skeptics' argument from disagreement. The problem with Oreskes and Conway's approach is not that it is implausible. Instead, the real problem consists in the lack of a

[12] I will use the notions of consensus and agreement (more precisely, agreement under the conditions described in section 3) interchangeably.

[13] A prominent example is the "new view of science" advocated by Oreskes and Conway (2010, epilogue), but see also, for instance, Shapin and Schaffer (1985, chapter 2), Oreskes (2007), Rehg and Staley (2017), Oppenheimer et al. (2019), among others.

https://doi.org/10.1017/psa.2026.10195 Published online by Cambridge University Press

detailed argument supporting it. For this reason, my main goal is to provide such a detailed argument. I base this argument on the *Justificatory Account of Multiple Testimony*.

3. The (Bayesian) Justificatory Account of Multiple Testimony

In the previous section I have assumed, alongside historians and philosophers of science as well as—perhaps surprisingly—with strategic science skeptics, that agreement plays a central epistemic role in science. But what could this epistemic role be? I will propose an account according to which the epistemic role of agreement consists in *justifying* or *confirming* scientific claims.[14] Since I consider agreement and disagreement as a special sort of testimonial evidence (as I will explain below), I will attempt to capture this justificatory role in terms of the *Justificatory Account of Multiple Testimony*. I will present the general version of the *Justificatory Account of Multiple Testimony* (subsection 3.1). To make the notion of justification or confirmation figuring in this account more precise, I will develop a version of the account that is articulated with the help of Bayesian confirmation theory (subsection 3.2).

3.1. The Justificatory Account of Multiple Testimony—general version

To develop the *Justificatory Account of Multiple Testimony*, I will first lay a terminological foundation: I will propose a definition for the key concepts of agreement and disagreement as forms of testimonial evidence consisting of the testimonies or witness reports of scientists. In a second step, I will make an assumption about which different kinds of reports are relevant in scientific contexts. In a final third step, I will articulate the *Justificatory Account of Multiple Testimony* itself.

First step: Testimonial definitions of agreement and disagreement. I will draw on the concept of testimony taking center stage in social epistemology.[15] In social epistemology, testimonies are also referred to as the reports of witnesses. I will use the notions of testimonies and of witness reports interchangeably and I will assume that in scientific contexts scientists can and occasionally do act as witnesses by reporting or testifying to the truth or falsity of propositions.

I propose to define the notions of scientific agreement and disagreement as consisting of the testimonies or reports of multiple witnesses, where the reports regard the truth of some proposition.[16]

> **Testimonial Definition of Scientific Agreement:** There is a scientific agreement about a proposition *R* if and only if multiple witnesses (typically scientists) all report that proposition *R* is true.

[14] I use "justification" and "confirmation" interchangeably.

[15] Leonard (2023) provides a survey to the debate on testimony in social epistemology. For recent work on scientific testimony, see Gerken (2022).

[16] I adopt truth talk, but the *Justificatory Account of Multiple Testimony* can be reformulated in antirealist terms—for instance, by referring to empirically adequate hypotheses, as constructive empiricists would (van Fraassen 1980, 12 and 64).

https://doi.org/10.1017/psa.2026.10195 Published online by Cambridge University Press

Testimonial Definition of Scientific Disagreement: There is a scientific disagreement about a proposition *R* if and only if multiple witnesses differ in their reports as to whether proposition *R* is true—some say it is, others say it is not (that is, others say that ¬*R* is true).

In what follows, I will adopt these definitions of scientific agreement and disagreement.[17] If defined in this way, both agreement and disagreement can play the role of testimonial evidence consisting of the testimonies or reports of multiple witnesses, or so I will argue below.

Second step: Characterizing reports in science. Are there reports in science? If so, what is the content of scientific reports? The *Justificatory Account of Multiple Testimony* I am about to develop is flexible in allowing for different kinds of reported propositions. In science, reports come in different (often published) forms,[18] including the following:

1. The report of a scientist may consist in summarizing the results of the scientist's research: that is, the empirical evidence she obtained by using a specific method and how this empirical evidence should be interpreted with respect to confirming or disconfirming a hypothesis of interest. Consider a toy example of what such a report might look like: "In my laboratory, I heated water, used a mercury thermometer, and measured that the water boiled at about 100 degrees Celsius. This is empirical evidence confirming the hypothesis that water boils at 100 degrees Celsius." The non-toyish form of such a report is a research paper.
2. Reports sometimes express the expert assessment of a scientist regarding the justificatory status of a hypothesis based on the state of the art in the research of a scientific field. A typical locus for this kind of report is a meta-study or an assessment report. Again, imagine a simple example for illustration: "Given the research in my field of expertise, I can say that the hypothesis that smoking causes lung cancer is empirically well-confirmed."
3. Reports can be expert judgments on a specific subject matter for which there is, at the time, no available empirical evidence, but the judgment is informed by research on directly related matters. For instance, an expert might testify in the following manner: "Although there is no empirical evidence at present as to how the new virus V is transmitted, it is likely that V is transmitted via respiration given that there is a vast amount of research on the transmission of similar viruses and this research provides strong empirical evidence for this mode of transmission."

[17] Defining agreement in terms of unanimity might strike some philosophers of science as granting the skeptics too much, because unanimity is hard to achieve in real scientific life—or so the worry might go. However, this worry is not justified, because what matters for confirming a scientific hypothesis is not necessarily agreement, as I will argue in section 5.

[18] I will work with two assumptions regarding reports. *First*, scientific reports can be (i) asserted by a single scientist or by a group of collaborating scientists (for instance, in the case of a co-authored report), and (ii) directed at different audiences including other scientists, policy-makers, and (non-expert) citizens (Gerken 2022, chapter 1). *Second*, I assume that scientific reports playing the role of testimonial evidence in science are, at least typically, directly about or connected to empirical evidence.

https://doi.org/10.1017/psa.2026.10195 Published online by Cambridge University Press

I do not claim that this list is an exhaustive and systematic typology of scientific reports. I merely assume that, if it makes sense to talk about the reports of scientists at all, the list presents plausible examples of such reports.[19]

Third step: Articulating the Justificatory Account of Multiple Testimony. Having provided testimonial definitions of agreement and disagreement, and having clarified the character of scientific reports, I now proceed to articulating the *Justificatory Account of Multiple Testimony*. I begin with a general version of the account.

> **Justificatory Account of Multiple Testimony (general version):** The reports of multiple witnesses justify some hypothesis H if and only if the total testimonial evidence (that is, the entirety of reports) provides a reason for an epistemic agent to belief that H is (more likely to be) true.

Henceforth, I will adopt the following convention regarding notation: R denotes a positive report justifying or confirming H, while $\neg R$ refers to a negative report that disconfirms H (and, hence, confirms $\neg H$).[20]

Consider a toy example to illustrate the general version of the account (inspired by the first kind of scientific report listed in the second step above). Suppose that the hypothesis H in question is that water boils at 100 degrees Celsius. Imagine further that ten scientists do their experiments individually to test this hypothesis and they all report the following proposition R afterwards: "In my experiment, I measured that water boiled at 100 degrees Celsius." Finally, assume that—surely, counter to fact—the ten reports are all the reports there are regarding this hypothesis. These ten reports constitute the *total testimonial evidence* (the entirety of reports). This is a situation of agreement. According to the *Justificatory Account of Multiple Testimony*, these ten reports justify the hypothesis that water boils at 100 degrees Celsius if and only if the ten reports provide a reason for an epistemic agent to believe that this hypothesis is (more likely to be) true.

Note that, in this toy example, I assume that there is *agreement* in the sense defined above: *all* ten scientists report the same proposition. However, I will argue in section 5 that agreement is not necessary for justifying a hypothesis via testimonial evidence. Even if witnesses disagree, the total testimonial evidence might still justify the hypothesis in question. Hence, to make conceptual room for something less than agreement (relative to a group of witnesses), I refer to "the reports of multiple witnesses" when articulating the *Justificatory Account of Multiple Testimony* in its general version.

3.2. The Justificatory Account of Multiple Testimony—Bayesian version

The general version of the *Justificatory Account of Multiple Testimony* raises a question: What does the vague demand that reports *provide a reason to believe a hypothesis* amount to? Bayesian confirmation theory provides one possible and

[19] The items on the list are inspired by the kinds of reports that Shapin and Schaffer (1985, 55–65), Rehg and Staley (2017), and Gerken (2022, 15–19) distinguish.

[20] I focus on simple cases where the reports consist in asserting a proposition or its negation. For more complex cases, see, for instance, Earman (2000, 56–59), and Bovens and Hartmann (2003, chapters 1–2).

https://doi.org/10.1017/psa.2026.10195 Published online by Cambridge University Press

fruitful way to make the rough idea of *providing a reason to believe a hypothesis* more precise.[21]

> **Justificatory Account of Multiple Testimony (Bayesian version)**: The reports of multiple witnesses confirm hypothesis H if and only if the following condition holds: when an epistemic agent learns about the total testimonial evidence (the entirety of reports), then the agent should have greater confidence in the truth of H.

Bayesians interpret the vague notion of an epistemic agent having *greater confidence in the truth of a hypothesis* in terms of subjective probabilities: the agent assigns a posterior probability to H being true, $\mathbf{P^*}(H)$, that is greater than the prior probability the agent assigned to H, $\mathbf{P}(H)$. In the context of multiple reports, the prior probability is the probability the agent assigned to H before receiving the reports, whereas the posterior probability is the probability the agent assigns to H after having received the reports.

If we use $W_i(R)$ to express that one of n witnesses, witness i, reports that proposition R is true, and K to express background knowledge, then Bayesians articulate the *Justificatory Account of Multiple Testimony* via this inequality (for $i = 1, \ldots, n$):

$$\mathbf{P}^*(H) = \mathbf{P}[H|W_1(R), \ \ldots, W_n(R), \ K] > \mathbf{P}(H|K).$$

This Bayesian articulation rests on three basic assumptions. *First*, I adopt Earman's proposal to state the Bayesian version explicitly as being conditional on the background knowledge K of an epistemic agent (Earman 1992, 33–34; 2000, 26–27 and 55). I take such background knowledge to include previously acquired empirical knowledge (that is, previously obtained research results, including empirical evidence) as well as methodological and conceptual knowledge, shared and accepted in a field of scientific inquiry.[22]

Second, I assume that, if an epistemic agent learns the total testimonial evidence—that is, the reports $W_1(R), \ \ldots, \ W_n(R)$—then it does not make a difference to the posterior probability whether the agent learns the entire testimonial evidence all at once or in chunks over time.

Third, I assume for simplicity's sake that an epistemic agent obtains the posterior probability of H via the updating rule of strict conditionalization, according to which (testimonial) evidence is believed with certainty.[23] In other words, appealing to Bayes' theorem, the posterior probability $\mathbf{P^*}(H)$ is determined as follows:[24]

[21] I am building on seminal statements of the Bayesian approach including Jeffrey (1992, 109–10), Goldman (1999, 223), Bovens and Hartmann (2003, section 3.3), and Sober (2008, 42–43). My main point of reference is Earman (2000, 55–58).

[22] One could say, as a reviewer suggested, that an epistemic agent (partly) demarcates the relevant group of witnesses by ascribing background knowledge.

[23] On the rule of strict conditionalization, see, for instance, Earman (1992, 34) and Sprenger and Hartmann (2019, 22). However, other updating rules are available that allow for uncertain evidence, perhaps most importantly Jeffrey conditionalization (Jeffrey 1983, chapter 11; 2004, chapter 3). It will be a task for future work to explore whether such alternatives hold an advantage over strict conditionalization if applied to the topic at hand.

[24] For informal and formal proofs, see Earman (2000, 55) and Bovens and Hartmann (2003, appendix E.1).

https://doi.org/10.1017/psa.2026.10195 Published online by Cambridge University Press

$$\mathbf{P}^*(H) = \mathbf{P}[H|W_1(R),\ \ldots,\ W_n(R),\ K] = \frac{1}{1 + \left(\frac{\mathbf{P}(\neg H|K)}{\mathbf{P}(H|K)}\right) \times \left(\frac{\mathbf{P}[W_1(R)|\neg H,K]}{\mathbf{P}[W_1(R)|H,K]} \times \cdots \times \frac{\mathbf{P}[W_n(R)|\neg H,K]}{\mathbf{P}[W_n(R)|H,K]}\right)}.$$

To see how the Bayesian version works, it is helpful to assume that two conditions hold: the *Minimal Reliability Condition* and the *Conditional Independence Condition*. The second condition is strictly speaking not a necessary one, if one would like to adopt the *Justificatory Account of Multiple Testimony*. I appeal to it mainly to keep the exposition of the Bayesian version simple. Indeed, there are various proposals of how to phrase and to weaken the (simplifying or idealized) notion of conditional independence to render it plausible in scientific contexts (involving collaboration and exchange between scientists).[25] It is not my goal to assess the merit of such proposals. Here, my focus will be on the minimal reliability of witnesses.

According to the *Minimal Reliability Condition*, it is more probable that a witness W_i reports R if the conjunction of H and, as already indicated above, background knowledge K is true rather than if $\neg H$ and K is true. This condition expresses the intuition that, for a reliable witness, the content of a report (R or $\neg R$) depends sensitively on the truth of H and K. Mathematically, we can express the *Minimal Reliability Condition*[26] regarding W_i's reporting of R as follows in terms of likelihoods:

$$\mathbf{P}[W_i(R)|H,\ K] > \mathbf{P}[W_i(R)|\neg H,\ K].$$

The assumption that a witness is minimally reliable can also usefully and equivalently be expressed as a likelihood ratio that is greater than 1:

$$\frac{\mathbf{P}[W_i(R)|H,\ K]}{\mathbf{P}[W_i(R)|\neg H,\ K]} > 1.$$

Moreover, it holds that a witness W_i is minimally reliable with respect to reporting R if and only if W_i is also minimally reliable in reporting $\neg R$:

$$\frac{\mathbf{P}[\mathrm{W_i(R)|H,K}]}{\mathbf{P}[\mathrm{W_i(R)|\neg H,K}]} > 1 \Leftrightarrow \frac{\mathbf{P}[\mathrm{W_i(\neg R)|\neg H,K}]}{\mathbf{P}[\mathrm{W_i(\neg R)|H,K}]} > 1$$

The *Conditional Independence Condition* requires that the reports of multiple witnesses asserting that R (and, respectively, that $\neg R$) is true be probabilistically independent, *conditional on* whether H and K is true (or whether $\neg H$ and K) is true. The intuition expressed by this condition is that the report of one witness is not affected by what other witnesses report. The *Conditional Independence Condition* with respect to reporting R is stated as follows in formal terms (for $i = 1,\ \ldots,\ n$):

$$\mathbf{P}\big[W_i(R)|W_1(R),\ \ldots,\ W_{i-1}(R),\ W_{i+1}(R),\ldots,W_n(R),\ H,\ K\big] = \mathbf{P}[W_i(R)|H,\ K],$$

and respectively for $\neg R$:

$$\mathbf{P}\big[W_i(\neg R)|\ W_1(\neg R),\ \ldots,\ W_{i-1}(\neg R),\ W_{i+1}(\neg R),\ldots,W_n(\neg R),\ \neg H,\ K\big]$$

[25] Such work on de-idealizing independence includes Bovens and Hartmann (2003, section 3.5), Dietrich and List (2004, 180–84), and Dietrich and Spiekermann (2013, 98 and 104–7).

[26] Note that, according to the *Minimal Reliability Condition*, neither $\mathbf{P}[W_i(R) \mid H, K]$ nor $\mathbf{P}[W_i(\neg R) \mid \neg H, K]$ are required to be greater than 0.5.

https://doi.org/10.1017/psa.2026.10195 Published online by Cambridge University Press

$$= \mathbf{P}[W_i(\neg R)|\neg H,\ K].$$

There is also a second and equivalent way to express conditional independence (for $i = 1,\ \ldots,\ n$):

$$\mathbf{P}[W_1(R), \ldots, W_n(R)|H,\ K] = \mathbf{P}[W_1(R)|H,\ K] \times \cdots \times \mathbf{P}[W_n(R)|H,\ K]$$

and

$$\mathbf{P}[W_1(\neg R), \ldots, W_n(\neg R)|\neg H,\ K] = \mathbf{P}[W_1(\neg R)|\neg H,\ K] \times \cdots \times \mathbf{P}[W_n(\neg R)|\neg H,\ K].$$

In what follows, I will appeal to the second way of stating independence.[27]

To summarize, if the *Minimal Reliability Condition* and the *Conditional Independence Condition* are satisfied, then the Bayesian *Justificatory Account of Multiple Testimony* holds.

To simplify matters for the remainder of the paper, I will talk straightforwardly about likelihood ratios and not about posterior probabilities. In doing so, I will appeal to what I will call the *Confirmation Likelihood Principle*, according to which multiple reports confirm H if and only if the likelihood ratio (regarding all relevant reports) is greater than 1. Formally, we can express this principle as follows (for $i = 1,\ \ldots,\ n$):

$$\mathbf{P}[H|W_1(R),\ \ldots,\ W_n(R),\ K] > \mathbf{P}(H|K) \Leftrightarrow \frac{\mathbf{P}[W_1(R),\ \ldots,\ W_n(R)|H,\ K]}{\mathbf{P}[W_1(R),\ \ldots,\ W_n(R)|\neg H,\ K]} =: x > 1.$$

The *Confirmation Likelihood Principle* is the formal core of the Bayesian version of the *Justificatory Account of Multiple Testimony*.[28]

With this principle in place, let me turn to an illustration of how the Bayesian version of the *Justificatory Account of Multiple Testimony* works. Suppose that two witnesses, W_1 and W_2, both report R, expressed as "$W_1(R)$" and "$W_2(R)$". These reports constitute the total testimonial evidence in the imagined situation. In other words, we imagine a scenario of agreement regarding R (in the sense of agreement defined in subsection 3.1). Given that the two witnesses are minimally reliable and report independently, their individual likelihood ratios can be multiplied (see the second way of expressing conditional independence) and the product of their individual likelihood ratios is greater than 1. This can be stated—regarding the minimally reliable and conditionally independent reporting of R—by the following inequality, in which each fraction expresses the likelihood ratio for a single witness (for $i = 1, 2$):

$$\frac{\mathbf{P}[W_1(R)|H,\ K]}{\mathbf{P}[W_1(R)|\neg H,\ K]} \times \frac{\mathbf{P}[W_2(R)|H,\ K]}{\mathbf{P}[W_2(R)|\neg H,\ K]} > \frac{\mathbf{P}[W_i(R)|H,\ K]}{\mathbf{P}[W_i(R)|\neg H,\ K]} > 1.$$

Now, if the inequality above holds, then, according to the *Confirmation Likelihood Principle*, the reports of the witnesses W_1 and W_2 (the total testimonial evidence) confirm H—that is, the posterior probability is greater than the prior probability $\mathbf{P}(H|K)$.[29]

[27] According to the multiplication rule of probability theory, $\mathbf{P}(A, B) = \mathbf{P}(A) \times \mathbf{P}(B|A) = \mathbf{P}(B) \times \mathbf{P}(A|B)$. Hence, if, given the first definition of independence above, $\mathbf{P}(B|A) = \mathbf{P}(B)$, then $\mathbf{P}(A, B) = \mathbf{P}(A) \times \mathbf{P}(B)$.

[28] In section 4 I will turn to two adaptations of this principle.

[29] Likewise, it is assumed to hold for the minimally reliable and independent reporting of $\neg R$ (for $i = 1, 2$):

$$\frac{\mathbf{P}[W_1(\neg R)|\neg H, K]}{\mathbf{P}[W_1(\neg R)|H, K]} \times \frac{\mathbf{P}[W_2(\neg R)|\neg H, K]}{\mathbf{P}[W_2(\neg R)|H, K]} > \frac{\mathbf{P}[W_i(\neg R)|\neg H, K]}{\mathbf{P}[W_i(\neg R)|H, K]} > 1.$$

https://doi.org/10.1017/psa.2026.10195 Published online by Cambridge University Press

It is worth stressing the increased confirmatory power of multiple reports in comparison to a single report. Although each single report considered *in isolation* constitutes confirming evidence for H, in this illustration the product of the two individual likelihood ratios is greater than the individual likelihood ratio of each witness considered in isolation. For this reason, the reports of the two witnesses—the *total testimonial evidence* in this situation—provide stronger confirming evidence for the truth of H than the report of either W_1 or W_2 in isolation. That is, the posterior probability $\mathbf{P}[H|W_1(R), W_2(R), K]$ is not only greater than the prior probability $\mathbf{P}(H|K)$ but also greater than the posterior probability resulting from conditionalizing on only one of the two witness reports in isolation—that is, $\mathbf{P}(H|W_i(R), K)$. (The distinction between the total testimonial evidence and a single report considered in isolation will resurface in subsection 5.1.)

In what follows (especially in section 5), I will use the notion of *Collective Reliability* as a term of art to refer to the product of the likelihood ratios of a specified group of witnesses. (In the illustration above, the specified group consists of the witnesses W_1 and W_2.) The notion of *Collective Reliability* draws on reliability as defined in terms of likelihood ratios (already introduced for individual witnesses in the *Minimal Reliability Condition* above).

4. The counterargument from unreliable witnesses (against the Disagreement Premise)

To debunk the skeptics' argument, I begin with a counterargument to the *Disagreement Premise*. In a first step, I will propose that Bayesians should revise this premise. Then, I will argue that the revised premise does not hold, at least typically, if the reports of skeptics are part of a disagreement.

First, observe that the *Disagreement Premise* in its original form (section 2) does not require that witnesses (who disagree) be minimally reliable ones. However, on the Bayesian version of the *Justificatory Account of Multiple Testimony*, witnesses must be minimally reliable, if their reports that R (or that $\neg R$) are supposed to confirm (or to disconfirm) H. For this reason, a proponent of the Bayesian version should reject the *Disagreement Premise* as not being sufficiently well formulated and insist on revising it as follows:

> **Revised Disagreement Premise:** A member of the public learns that minimally reliable (and independently reporting) witnesses disagree in their reports (some report R, while others report $\neg R$) regarding a scientific hypothesis H.

It is useful to explain briefly why Bayesians should indeed reject and revise the original form of the premise. For this purpose, it is important to realize that witnesses might fail to be reliable in two different ways.[30]

First, a witness might be *non-reliable* in the sense that it is probabilistically irrelevant for the report R of a witness whether H is true (given K). Mathematically, non-reliability can be expressed as follows:

[30] My main reference point is Earman (2000, 55). Bovens and Hartmann (2003, 14–15) analyze non-reliable witnesses.

https://doi.org/10.1017/psa.2026.10195 Published online by Cambridge University Press

$$\frac{\mathbf{P}[W_i(R)|H,\ K]}{\mathbf{P}[W_i(R)|\neg H,\ K]} = 1.$$

Another way in which a witness might fail to be minimally reliable is even more drastic, as an unreliable witness might be *anti-reliable* in the following sense of the term:

$$\frac{\mathbf{P}[W_i(R)|H,\ K]}{\mathbf{P}[W_i(R)|\neg H,\ K]} < 1.$$

In other words, it is more probable that the witness provides a positive confirming report R if H is false than if H is true (given K). A concrete illustration might consist in a witness reporting R to intentionally mislead someone to believe that H is true.

I take it as a premise that the *Confirmation Likelihood Principle* can be adapted (by drawing on Bayes' theorem)[31] to describe the logical relationship between the likelihood ratio of a witness, on the one hand, and confirmationally neutral and disconfirming reports, on the other hand. For present purposes, this adaptation takes the form of two principles:

(i) If the likelihood ratio of a witness W_i with respect to reporting R is *equal* to 1, then the report of the witness is *confirmationally neutral* with respect to H (relative to K)—that is, $\mathbf{P}[H|W_i(R),\ K] = \mathbf{P}(H|K)$.
(ii) If the likelihood ratio of a witness W_i with respect to reporting R is *smaller* than 1, then the report of the witness *disconfirms* H (relative to K)—that is, $\mathbf{P}[H|W_i(R),\ K] < \mathbf{P}(H|K)$.

Now, according to principle (i), if a *non-reliable* witness reports R, then the report is *confirmationally neutral* with respect to H. Stating that a report is confirmationally neutral does not imply that an epistemic agent ignores the report or "silences" the witness. To the contrary, an epistemic agent learns this report but then realizes that the resulting posterior probability equals the prior probability assigned to the hypothesis. If one imagines multiple non-reliable witnesses, the situation remains unchanged: even the reports of multiple non-reliable witnesses do not increase the posterior probability, as the *Collective Reliability* is equal to 1.

According to principle (ii), the report of an *anti-reliable* witness (that R) should be taken to *disconfirm* H. If multiple anti-reliable witnesses report that R, then the disconfirming impact on H increases, as the *Collective Reliability* of anti-reliable witnesses gets smaller as the number of such witnesses grows.

Moving on to the second step, suppose an epistemic agent learns that there is a disagreement about a hypothesis (partly or entirely) because strategic science skeptics themselves provide disconfirming reports. Does the epistemic agent learn *in this situation* that there is the kind of disagreement required by the *Revised Disagreement Premise*? In other words, is the premise true of such situations? The answer turns on the question of whether the agent should think that the skeptics report in a minimally reliable way.

[31] As described by Earman (2000, 55) and Sober (2008, 13–15).

https://doi.org/10.1017/psa.2026.10195 Published online by Cambridge University Press

To address this question, I will build on various criteria (or indicators) for ascribing unreliability proposed in the research literature on strategic science skepticism.[32] I realize that these criteria might be controversial, and my goal here is not to defend or evaluate these criteria. I only intend to explore the consequences of accepting some of these criteria in a Bayesian framework.[33] I believe a plausible case can be made that skeptics, at least typically, fail to qualify as minimally reliable witnesses if one accepts these criteria.[34]

To start with a prominent proposal for such a criterion, Oreskes and Conway diagnose that strategic science skeptics typically lack scientific expertise[35] in the field of research they target (2010, 8, 270–73).[36] *If* one takes the lack of expertise as a criterion for not being a reliable witness, then skeptics should, at least typically, be deemed unreliable (at least, regarding a specific field).[37] As a consequence, a skeptic's report fails to constitute disconfirming evidence and one should not be less confident about a hypothesis "just because some person, somewhere, doesn't agree" (Oreskes and Conway 2010, 272). The Bayesian version of the *Justificatory Account of Multiple Testimony* can express this consequence in a precise way, if one accepts Oreskes and Conway's criterion for assigning unreliability. Indeed, the Bayesian version provides a nuanced distinction of two failures of minimal reliability and the corresponding rational changes of belief.

If one considers the research literature more broadly, one can rely on further criteria for diagnosing that skeptics are unreliable witnesses (criteria potentially complementing the lack of expertise criterion). For instance, the reports of skeptics frequently contain systematic errors. From a philosophical perspective, these errors are analyzed in different ways; for instance, errors are described as skeptics assessing empirical evidence and how it supports a hypothesis in outright distorting (Douglas 2009; Carrier 2018; Reutlinger 2020) or misleading ways (Steel 2018), as violating methodological standards (Wilholt 2009; Biddle and Leuschner 2015), and as avoiding quality-control mechanisms (such as peer review) accepted in the relevant scientific community (Oreskes and Conway 2010). For current purposes, I take the presence of systematic error to indicate a lack of reliability.

If one accepts these criteria, then one has good reasons for deeming skeptics unreliable witnesses—that is, non-reliable or anti-reliable witnesses.[38] However, as

[32] I take a criterion to increase an agent's confidence that a witness is unreliable.

[33] If there are cases where the criteria do not warrant deeming skeptics unreliable, the second counterargument becomes relevant for debunking the skeptics' argument (section 5).

[34] A proponent of the Bayesian version could also argue that strategic science skeptics fail to report *independently* of one another. I do not want to rule out this additional objection, but in this paper I will stick to the simplifying assumption that the independence condition is satisfied.

[35] Although Oreskes and Conway do not define expertise explicitly, their notion of expertise is tied to making research contributions in a specialized scientific field.

[36] See also Gerken (2022, section 5.5) on a general analysis of "expert trespassing testimony."

[37] Let me add two remarks. First, Oreskes and Conway typically apply this criterion in combination with other criteria (for instance, that skeptics are motivated by non-epistemic interests). Second, although I do not intend to improve this criterion here, a more sophisticated version of the criterion should integrate at least three aspects: (i) that there are different kinds of scientific expertise (Collins and Evans 2007; Anderson 2011); (ii) that there might be kinds of extra-scientific expertise (Epstein 1996; Wynne 1996); and (iii) how the criterion relates to committing systematic errors (see below).

[38] I merely commit to the disjunction here. I leave an analysis of (historical) case studies of strategic science skepticism, in which I could support one of the disjuncts, for another occasion.

https://doi.org/10.1017/psa.2026.10195 Published online by Cambridge University Press

argued above, the *Revised Disagreement Premise* requires that witnesses be reliable. Hence, an epistemic agent does *not* learn the sort of disagreement expressed by the *Revised Disagreement Premise* (that is, a disagreement among reliable witnesses) if the agent learns about the existence of a disagreement (partly or entirely) constituted by the reports of unreliable skeptics. In other words, the *Revised Disagreement Premise* is not true of such situations. This result constitutes a serious limitation to the ability of strategic science skeptics to use the argument from disagreement to undermine specific scientific hypotheses (I will return to this issue of limitation in subsection 5.4).[39]

5. The counterargument from total testimonial evidence (against the Doubt Rule)

Another way to debunk the skeptics' argument from disagreement consists in rejecting the *Doubt Rule*. The skeptics seem to suppose that the *Doubt Rule* is a *general rule* holding in all (or an indefinitely wide range of) situations. I will use the Bayesian *Justificatory Account of Multiple Testimony* to develop the counterargument from total testimonial evidence to reject the *Doubt Rule* as a general rule, since it fails to hold in certain situations. I will do so by arguing that there are situations in which an agent learns about the disagreeing reports of witnesses (in relation to a hypothesis H) and the agent is not forced to lower her confidence in H but to be *more confident* about the truth of H.

The key idea of the counterargument is as follows. The *Justificatory Account of Multiple Testimony*—especially in its Bayesian version—does not require agreement to confirm a hypothesis via testimonial evidence. What really matters for confirming a hypothesis is whether the *total testimonial evidence*—that is, the entirety of reports given by independently and reliably reporting witnesses—confirms the hypothesis. To illustrate what the total testimonial evidence might amount to, it is useful to distinguish two situations. In the first situation, we assume that all the reporting witnesses are equally minimally reliable (subsection 5.1), while in the second situation all witnesses are minimally reliable but not equally so (subsection 5.2). I will argue that the two situations constitute failures of the *Doubt Rule* as a general rule (subsection 5.3) and discuss the merit of a possible, charitable revision of the *Doubt Rule* (subsection 5.4).

5.1. First situation: Equally reliable witnesses

Suppose that the reporting witnesses are all equally minimally reliable—that is, they all have the exact same likelihood ratio—and they report independently of one another. To illustrate, imagine a concrete realization of this kind of situation with a group of three witnesses, W_1, W_2, and W_3. Two of them, W_1 and W_2, report that R is true. The third witness, W_3, reports that $\neg R$ is true. This is a situation of disagreement, according to the testimonial definition of disagreement (subsection 3.1). Given that all witnesses are equally minimally reliable and report independently, W_1 and W_2 exceed W_3 in terms of *Collective Reliability*. This is the case if the following inequalities hold:

[39] Earman indicates but does not elaborate a similar counterargument in the context of "silly tales" (Earman 2000, 61).

https://doi.org/10.1017/psa.2026.10195 Published online by Cambridge University Press

$$\frac{\mathbf{P}[W_1(R)|H,\,K]}{\mathbf{P}[W_1(R)|\neg H,\,K]} \times \frac{\mathbf{P}[W_2(R)|H,\,K]}{\mathbf{P}[W_2(R)|\neg H,\,K]} > \frac{\mathbf{P}[W_3(R)|H,\,K]}{\mathbf{P}[W_3(R)|\neg H,\,K]} > 1,$$

$$\frac{\mathbf{P}[W_1(\neg R)|\neg H,\,K]}{\mathbf{P}[W_1(\neg R)|H,\,K]} \times \frac{\mathbf{P}[W_2(\neg R)|\neg H,\,K]}{\mathbf{P}[W_2(\neg R)|H,\,K]} > \frac{\mathbf{P}[W_3(\neg R)|\neg H,\,K]}{\mathbf{P}[W_3(\neg R)|H,\,K]} > 1.$$

In a situation involving equally minimally reliable witnesses, the total testimonial evidence confirms H, given the *Confirmation Likelihood Principle*, if a *majority* of the witnesses reports that R is true. In the imagined situation, W_1 and W_2 are the majority, W_3 is the minority, and the total testimonial evidence consists of the three reports $W_1(R)$, $W_2(R)$, and $W_3(\neg R)$. If so, the relevant posterior probability $\mathbf{P}^*(H) = \mathbf{P}[H|W_1(R), W_2(R), W_3(\neg R),\ K]$ is greater than the prior probability $\mathbf{P}(H \mid K)$.

Recall the distinction between the total testimonial evidence and one single report (out of multiple reports) considered in isolation (subsection 3.2). It is true that the report $W_3(\neg R)$, taken *in isolation*, disconfirms H. However, the reports of W_1 and W_2 (that R) confirm H more strongly than the report of W_3 (that $\neg R$) disconfirms H. Hence, the total testimonial evidence (all three reports) confirms H in this situation of disagreement.

Abstracting from this specific illustration with three witnesses, Bayesians can rely on the following more general verdict: a majority of equally reliable (and independent) witnesses reporting R provides confirming testimonial evidence for H, because the majority's *Collective Reliability* (that is, the product of the likelihood ratios of the members of the majority) is greater than the minority's *Collective Reliability* (that is, the product of the likelihood ratios of the members of the minority who reliably and independently report $\neg R$).[40]

5.2. Second situation: Witnesses differing in reliability

The Bayesian approach is, however, *not* necessarily a matter of a *majority* of reliably and independently reporting witnesses. Indeed, it is a virtue of the Bayesian version of the *Justificatory Account of Multiple Testimony* to illuminate the limits of such political metaphors in epistemology (here, talk of majority and minority). This becomes clear if one supposes that the reporting witnesses are all minimally reliable, but some witnesses are more reliable than others.[41] Furthermore, suppose that there is a disagreement regarding proposition R within such a group of n unequal witnesses: one subgroup of witnesses reports R, while another subgroup reports $\neg R$. According to the *Confirmation Likelihood Principle*, what matters for confirming H via the reports of multiple witnesses in this kind of situation is that the *Collective Reliability* of the subgroup of witnesses $W_1, \ldots, W_m$ reporting R is greater than the *Collective Reliability* of the subgroup of witnesses $W_{m+1}, \ldots, W_n$ reporting $\neg R$. Mathematically, this claim can be expressed as follows:

$$\frac{\mathbf{P}[W_1(R)|H,\,K]}{\mathbf{P}[W_1(R)|\neg H,\,K]} \times \cdots \times \frac{\mathbf{P}[W_m(R)|H,\,K]}{\mathbf{P}[W_m(R)|\neg H,\,K]} >$$

[40] Jeffrey (1992, 110–14) and Bovens and Hartmann (2003, section 3.6) provide Bayesian approaches to disagreeing reports.

[41] Jeffrey (1992, 110–14) and Earman (2000, 55–56) extend the Bayesian approach to scenarios involving witnesses that differ in minimal reliability.

https://doi.org/10.1017/psa.2026.10195 Published online by Cambridge University Press

$$\frac{\mathbf{P}[W_{m+1}(R)|H,\,K]}{\mathbf{P}\big[W_{m+1}(R)|\neg H,\,K\big]} \times \cdots \times \frac{\mathbf{P}[W_n(R)|H,\,K]}{\mathbf{P}[W_n(R)|\neg H,\,K]} > 1.$$

In this kind of situation, confirming reports can be realized in at least three different ways: (i) the subgroup of witnesses reporting R is the majority (this is similar to the first situation, except that now the witnesses may differ in their individual reliability within the subgroup); (ii) the subgroup of witnesses reporting R and the subgroup of witnesses reporting $\neg R$ are of equal size, but the former is collectively more reliable than the latter; or (iii) the subgroup of witnesses reporting R is a minority of highly reliable witnesses whose *Collective Reliability* exceeds that of a majority of less (but still minimally) reliable witnesses reporting $\neg R$.

5.3. The lesson from both situations

The *Doubt Rule* does not hold in both situations described in subsections 5.1 and 5.2. In both situations, an epistemic agent need not be less confident about H after learning about a disagreement.

In the first situation, the minority's reports (that $\neg R$) do not force an epistemic agent to lower her confidence in H. To the contrary, if one considers the total testimonial evidence—including the majority's reports (that R)—and the majority's greater *Collective Reliability*, one should assign a higher posterior probability to H.

In the second situation, the mere existence of disagreeing reports (whether they are the reports of a minority or not) does not force an epistemic agent to be less confident about H. To the contrary, the agent should be more confident about H, considering the total testimonial evidence, if the witnesses reporting R exceed the witnesses reporting $\neg R$ in their *Collective Reliability*.

In sum, what really matters for confirming a scientific claim via multiple testimonies is not necessarily agreement and not even what the majority of witnesses reports. Confirmation depends on whether the subgroup of witnesses reporting R has a greater *Collective Reliability* than the subgroup reporting $\neg R$. And this is possible even if witnesses disagree. Hence, the Bayesian version of the *Justificatory Account of Multiple Testimony* provides an argument for rejecting the *Doubt Rule* as a general rule and, thereby, one additional way to debunk the argument from disagreement.

5.4. Revising the Doubt Rule and the Argument from Disagreement?

I just provided an argument for rejecting the *Doubt Rule* as a *general rule*. In response, one might charitably ask whether the skeptics could respond to this counterargument by appealing to an appropriately revised, non-general version of the *Doubt Rule* that holds only under specific conditions. Based on the Bayesian *Justificatory Account of Multiple Testimony* and the lessons of subsection 5.3, one could revise this premise as follows:

> **Revised Doubt Rule:** If a member of the public learns that minimally reliable (and independently reporting) witnesses disagree in their reports regarding a scientific hyothesis H and the subgroup of witnesses who reliably (and independently) report $\neg R$ has greater *Collective Reliability* than the subgroup of

https://doi.org/10.1017/psa.2026.10195 Published online by Cambridge University Press

witnesses reliably (and independently) reporting *R*, then this person should be less confident about the truth of *H*.

Suppose one accepts the *Revised Doubt Rule*. If one also accepts the *Revised Disagreement Premise*, one obtains the *Revised Argument from Disagreement*. If so, how should one assess this revised argument?

Being able to state the *Revised Argument from Disagreement* exemplifies the fruitfulness of the Bayesian version of the *Justificatory Account of Multiple Testimony*, as this account is able to articulate what the skeptics *would have* to show to render their argument sound, or at least to render its premises generally more plausible.

Furthermore, proponents of the Bayesian version could accept that the revised premises are generally more plausible than the original ones. At the same time, they could argue that the skeptics have the burden of argument to show that the revised premises are in fact true in crucial kinds of situations: namely, when the skeptics use the argument to criticize specific scientific hypotheses (such as "smoking causes lung cancer"). It remains to be seen whether the skeptics are able to accomplish this task. In this regard, I am, for once, skeptical myself. I am confident that it will be considerably more difficult for strategic science skeptics to exploit the revised version of the argument from disagreement than its original version.

6. Conclusion

I have provided a rational reconstruction of the strategic science skeptics' argument from disagreement. To debunk the argument, I have presented the *Justificatory Account of Multiple Testimony*—particularly in its Bayesian version—as a normative epistemological treatment of the relationship between testimonial evidence and confirmation (and disconfirmation). I have relied on this account to argue against both premises of the argument from disagreement. Finally, I have indicated what it would take to revise the argument and expressed doubts that the revised argument puts the skeptics in a better position to use it.

Acknowledgments. First and foremost, I would like to express my special thanks to Leon Assaad, Stephan Hartmann, and Maria Kronfeldner for their repeated and productive feedback on earlier drafts of the manuscript; philosophical discussions with them encouraged and enabled me to develop my claims and arguments. I also thank two anonymous reviewers for this journal, kind audiences (in Kraków, Munich, and Vienna), and students at Ludwig-Maximilians-Universität München for productive comments that helped to improve the manuscript.

Funding Statement. None to declare.

Declarations. None to declare.

References

Anderson, Elisabeth. 2011. "Democracy, Public Policy, and Lay Assessments of Scientific Testimony." *Episteme* 8 (2):144–64. https://doi.org/10.3366/epi.2011.0013

Biddle, Justin, and Anna Leuschner. 2015. "Climate Skepticism and the Manufacture of Doubt: Can Dissent in Science be Epistemically Detrimental?" *European Journal for Philosophy of Science* 5:261–78. https://doi.org/10.1007/s13194-014-0101-x

Bovens, Luc, and Stephan Hartmann. 2003. *Bayesian Epistemology*. New York: Oxford University Press. https://doi.org/10.1093/0199269750.001.0001

https://doi.org/10.1017/psa.2026.10195 Published online by Cambridge University Press

Brandt, Allan. 2007. *The Cigarette Century. The Rise, Fall, and Deadly Persistence of the Product that Defined America*. New York: Basic Books.
Brown and Williamson. 1969. "Smoking and Health Proposal." UCSF Industry Documents Library, Brown and Williamson Records, Master Settlement Agreement. URL: https://www.industrydocuments.ucsf.edu/docs/jryf0138
Carrier, Martin. 2018. "Identifying Agnotological Ploys: How to Stay Clear of Unjustified Dissent., In *Philosophy of Science – Between Natural Science, the Social Sciences, and the Humanities (European Studies in Philosophy of Science)*, edited by Alexander Christian, David Hommen, Nina Retzlaff, and Gerhard Schurz, 155–69. Dordrecht: Springer. https://doi.org/10.1007/978-3-319-72577-2_9
Christensen, Jon. 2008. "Smoking Out Objectivity: Journalistic Gears in the Agnogenesis Machine." In *Proctor and Schiebinger* (2008), 266–82.
Collins, Harry, and Robert Evans. 2007. *Rethinking Expertise*. Chicago, IL: University of Chicago Press.
Dietrich, Franz, and Christian List. 2004. "A Model of Jury Decisions Where All Jurors Have the Same Evidence." *Synthese* 142:175–202. https://doi.org/10.1007/s11229-004-1276-z
Dietrich, Franz, and Kai Spiekermann. 2013. "Epistemic Democracy with Defensible Premises." *Economics and Philosophy* 29 (1):87–120. https://doi.org/10.1017/S0266267113000096
Dietrich, Franz, and Kai Spiekermann. 2020. "Jury Theorems." In *The Routledge Handbook of Social Epistemology*, edited by Miranda Fricker, Peter Graham, David Henderson, and Nikolaj Pedersen, 386–96. New York: Routledge. https://doi.org/10.4324/9781315717937
Douglas, Heather. 2006. "Bullshit at the Interface of Science and Policy: Global Warming, Toxic Substances and Other Pesky Problems." In *Bullshit and Philosophy*, edited by Gary Hardcastle and George Reisch, 213–26. Chicago, IL: Open Court.
Douglas, Heather. 2009. *Science, Policy and the Value-Free Ideal*. Pittsburgh, PA: University of Pittsburgh Press. https://doi.org/10.2307/j.ctt6wrc78
Earman, John. 1992. *Bayes or Bust? A Critical Examination of Bayesian Confirmation Theory*. Cambridge, MA: The MIT Press.
Earman, John. 2000. *Hume's Abject Failure. The Argument Against Miracles*. New York: Oxford University Press. https://doi.org/10.1093/0195127382.001.0001
Epstein, Steven. 1996. *Impure Science. AIDS, Activism, and the Politics of Knowledge*. Berkeley, CA: University of California Press.
Gerken, Mikkel. 2022. *Scientific Testimony. Its Roles in Science and Society*. Oxford: Oxford University Press. https://doi.org/10.1093/oso/9780198857273.001.0001
Glantz, Stanton, John Slade, Lisa Bero, Peter Hanauer, and Deborah Barnes. 1996. *The Cigarette Papers*. Berkeley, CA: University of California Press.
Goldman, Alvin. 1999. *Knowledge in a Social World*. Oxford: Clarendon Press. https://doi.org/10.1093/0198238207.001.0001
Harker, David. 2015. *Creating Scientific Controversies. Uncertainty and Bias in Science and Society*. Cambridge: Cambridge University Press. https://doi.org/10.1017/CBO9781107706903
Howe, Joshua. 2014. *Behind the Curve. Science and Politics of Global Warming*. Seattle, WA: University of Washington Press.
Jeffrey, Richard. 1983. *The Logic of Decision*, 2nd edn. Chicago, IL: The University of Chicago Press.
Jeffrey, Richard. 1992. *Probability and the Art of Judgment*. Cambridge: Cambridge University Press. https://doi.org/10.1017/CBO9781139172394
Jeffrey, Richard. 2004. *Subjective Probability: The Real Thing*. New York: Cambridge University Press. https://doi.org/10.1017/CBO9780511816161
Kourany, Janet, and Martin Carrier, eds. 2020. *Science and the Production of Ignorance. When the Quest for Knowledge is Thwarted*. Cambridge, MA: The MIT Press. https://doi.org/10.7551/mitpress/12146.001.0001
Leonard, Nick. 2023. "Epistemological Problems of Testimony." In *The Stanford Encyclopedia of Philosophy* (Spring 2023), edited by Edward Zalta and Uri Nodelman. URL: https://plato.stanford.edu/archives/spr2023/entries/testimony-episprob/.
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 *Proctor and Schiebinger* (2008), 90–107.

https://doi.org/10.1017/psa.2026.10195 Published online by Cambridge University Press

Odenbaugh, Jay. 2012. "Climate, Consensus, and Contrarians." In *The Environment. Philosophy, Science, and Ethics*, edited by William Kabasenche, Michael O'Rourke, and Matthew Slater, 137–50. Cambridge, MA: The MIT Press.

Oppenheimer, Michael, Naomi Oreskes, Dale Jamieson, Keynyn Brysse, Jessica O'Reilly, Matthew Shindell, and Milena Wazeck. 2019. *Discerning Expert. The Practices of Scientific Assessment for Environmental Policy*. Chicago, IL: The University of Chicago Press.

Oreskes, Naomi. 2007. "The Scientific Consensus on Climate Change: How Do We Know We Are Not Wrong?" In *Climate Change: What It Means for Us, Our Children, and Our Grandchildren*, edited by Joseph DiMento and Pamela Doughman, 65–99. Cambridge, MA: The MIT Press.

Oreskes, Naomi, and Erik Conway. 2008. "Challenging Knowledge: How Climate Science Became a Victim of the Cold War." In *Proctor and Schiebinger* (2008), 55–89.

Oreskes, Naomi, and Erik Conway. 2010. *Merchants of Doubt. How a Handful of Scientists Obscured the Truth on Issues from Tobacco Smoke to Global Warming*. New York: Bloomsbury Press.

Proctor, Robert. 1995. *Cancer Wars. How Politics Shapes What We Know and Don't Know About Cancer*. New York: Basic Books.

Proctor, Robert. 2008. "Agnotology: A Missing Term to Describe the Cultural Production of Ignorance (and Its Study)." In *Proctor and Schiebinger* (2008), 1–33.

Proctor, Robert. 2011. *Golden Holocaust. Origins of the Cigarette Catastrophe and the Case for Abolition*. Berkeley and Los Angeles, CA: University of California Press.

Proctor, Robert, and Londa Schiebinger, eds. 2008. *Agnotology. The Making and Unmaking of Ignorance*. Stanford, CA: Stanford University Press.

Rehg, William, and Kent Staley. 2017. "'Agreement' in the IPCC Confidence Measure." *Studies in the History and Philosophy of Modern Physics* 57:126–34. https://doi.org/10.1016/j.shpsb.2016.10.008

Reutlinger, Alexander. 2020. "What is Epistemically Wrong with Research Affected by Sponsorship Bias? The Evidential Account." *European Journal for Philosophy of Science* 10: 15. https://doi.org/10.1007/s13194-020-00280-2

Shapin, Steven, and Simon Schaffer. 1985. *Leviathan and the Air-Pump. Hobbes, Boyle, and the Experimental Life*. Princeton, NJ: Princeton University Press.

Sober, Elliott. 2008. *Evidence and Evolution. The Logic Behind the Science*. New York: Cambridge University Press. https://doi.org/10.1017/CBO9780511806285

Sprenger, Jan, and Stephan Hartmann. 2019. *Bayesian Philosophy of Science*. Oxford: Oxford University Press. https://doi.org/10.1093/oso/9780199672110.001.0001

Steel, Daniel. 2018. "If the Facts Were Not Untruths, Their Implications Were: Sponsorship Bias and Misleading Communication." *Kennedy Institute of Ethics Journal* 28:119–44. https://doi.org/10.1353/ken.2018.0011

Supran, Geoffrey, and Naomi Oreskes. 2017. "Assessing ExxonMobil's Climate Change Communications (1977–2014)." *Environmental Research Letters* 12 (8):084019. https://doi.org/10.1088/1748-9326/aa815f

van Fraassen, Bas. 1980. *The Scientific Image*. Oxford: Clarendon Press.

Wilholt, Torsten. 2009. "Bias and Values in Scientific Research." *Studies in the History and Philosophy of Science* 40:92–101. https://doi.org/10.1016/j.shpsa.2008.12.005

Wynne, Brian. (1996): "May the Sheep Safely Graze? A Reflexive View of the Expert–Lay Knowledge Divide." In *Risk, Environment, and Modernity: Towards a New Ecology*, edited by Scott Lash, Bronislaw Szerszynski, and Brian Wynne, 44–83. London: Sage.

Cite this article: Reutlinger, Alexander. 2026. "Unreasonable Doubt. How Strategic Science Skeptics Exploit the Argument from Disagreement." *Philosophy of Science* 93 (3):675–695. https://doi.org/10.1017/psa.2026.10195

https://doi.org/10.1017/psa.2026.10195 Published online by Cambridge University Press