Why we should trust scientific consensus: An extension of Naomi Oreskes's argument

TIMOTHY HIBBINS

Abstract

Since the very beginnings of modern science, findings that have challenged various aspects of the social status quo have been met with extreme resistance. Even after ideas such as heliocentrism, natural selection, vaccination, and global warming achieved consensus support within their respective fields, many opponents have continued to deny them on the grounds that the scientific consensus is corrupt, mistaken, dogmatic, or some combination thereof. It was in response to contemporary manifestations of this concern that the historian of science Naomi Oreskes wrote an article for *Time* magazine titled 'Science isn't always perfect—But we should still trust it' (2019). In the article she presents a nuanced case for why non-scientists can and should trust in scientific consensus. The purpose of this essay is to delve deeper into her argument and ultimately to strengthen it with the addition of new premises. I start by offering a brief reconstruction of the original argument before examining some key objections to it. After assessing the strengths of these objections, I suggest a way of addressing them that involves introducing a new distinction between two different types of consensus. I conclude that consensus only carries epistemic weight when it emerges from the process of scrutiny described by Oreskes, but that when it does, we are obliged to heed it.

Introduction

With regards to the relationship between science and society, there can be no doubt that we live in a rather unsettling historical moment. As we face a deadly global pandemic alongside the growing existential threat of irreversible climate change, the importance of science to the choices of both governments and individuals has perhaps never been greater. At the same time, however, an alarmingly large number of people are distrustful towards science, even on matters of life and death. In September of this year, a full 49 per cent of American adults said they would *not* get a COVID-19 vaccine if it were available, and only 21 per cent said they definitely would (Tyson et al. 2020). It is clear to me that building public trust in science is absolutely essential if we are to overcome these monumental challenges. I imagine it was a similar concern that motivated the distinguished historian of science Naomi Oreskes to pen her article for *Time* magazine titled 'Science isn't always perfect—But we should still trust it' (2019). While she is a historian, this particular article focuses on the present, and possesses a somewhat urgent tone. In it, she laments public distrust in science, its increasing politicisation, and the disinformation campaigns over the past decades that have led up to it. Despite this unfair playing field, Oreskes feels—and I concur—that if we are to get people to trust in science, then communicators in the media need to be able to provide a satisfactory and easily understandable account of *why* scientific consensus is worthy of their trust. Appeals to authority ('you should believe it because this professor in the lab coat said so!') and bandwagon fallacies ('you should believe it because everyone else already believes it!') will not do. Instead, Oreskes presents what I consider to be a novel argument for the reliability of scientific consensus. Before scientific claims can reach the status of consensus, she says, they must withstand an intense process of scrutiny. According to her, only true claims are able to survive this scrutiny—false claims will eventually get thrown out. Hence, the few claims that do become accepted as consensus are almost guaranteed to be correct.

This argument was obviously designed to be convincing to the general public, and in that I hope it succeeded. For the purposes of this essay, though, I am more interested in its philosophical content. I already think that Oreskes's argument is remarkably strong, and I will explain why, but my ultimate purpose in this essay is to suggest a way to make it even stronger. I will begin by offering a reconstruction of her case for trusting scientific consensus, before introducing two possible objections, assessing them, and finally explaining how I think the argument could be amended to address them.

Oreskes's argument

According to Oreskes, the story of the scientific method as written in textbooks will not do as an answer on its own, both because it fails to capture the variety of real scientific practice and because it wouldn't be sufficient grounding for trust regardless. She points out that even if an experiment returns a positive result, this cannot ever fully *confirm* a theory, as there will always be other possible theories that fit the result just as well. Furthermore, she adds, if an experiment returns a negative result, this cannot fully *falsify* a theory either, because the fault could lie with the experiment or with the interpretation rather than the theory itself. By highlighting these issues, she implicitly rejects many traditional accounts of science, at least to some degree. Instead, she argues, it is the way in which claims are evaluated by the *community* of scientists that is of primary importance. Specifically, Oreskes posits that science is defined by the requirement that all claims be subjected to a process of critical scrutiny.

Oreskes's criteria for scientific scrutiny

The concept of scrutiny is central to Oreskes's argument. For her, proper scientific scrutiny possesses three essential characteristics which make it, in her view, extremely reliable at rejecting false claims. Firstly, it is *tough*: scientists are intensely scrupulous when looking for faults in others' research and are not averse to giving harsh criticism. This happens both before publication, when papers are submitted for peer review, and afterwards when other scientists in the field write response papers providing counterevidence and objections. Secondly, scrutiny is *public*: work being exposed to a large and diverse group of other scientists means that claims are assessed from many different angles. Thirdly, and most importantly, scrutiny is *ongoing*: claims always remain fully vulnerable to rejection on the basis of new evidence. To build on the courtroom analogy that Oreskes employs to explain this aspect of scrutiny, not only is there no protection against double jeopardy in science, but there is also no statute of limitations. On her account, any scientific consensus that is subject to this process of scrutiny is thus fundamentally different from, say, religious dogma, in that it is always open to revision.

To describe how this tough, public, and ongoing scrutiny paves the way for scientific understanding, Oreskes borrows Helen Longino's (1990) term 'transformative interrogation'. To both Oreskes and Longino, science is an 'interrogation' because it involves questioning every claim rigorously, and 'transformative' because, in so doing, our understanding is transformed for the better. Oreskes readily concedes that, in the real world, there are often flaws in the process of scrutiny. However, she argues it is extraordinarily unlikely for a claim that hasn't survived *rigorous* scrutiny to become what could reasonably be called scientific consensus—so, scientific consensus is extraordinarily unlikely to be false. Therefore, she says, when scientists are all loudly saying the same thing, be it about the climate, vaccines, or anything else, *we should listen*.

Objections to Oreskes's argument

I believe there are two related but distinct objections one can raise against Oreskes's case for trusting science. I shall refer to these as the *historical objection* and the *systemic objection*.

Historical objection

The historical objection is an application of a famous argument in the philosophy of science known as the *pessimistic meta-induction*. The meta-induction was first postulated by Larry Laudan (1981) as a refutation of scientific realism, or the attitude that scientific theory accurately represents underlying reality rather than merely being useful. It applies induction over the population of scientific theories accepted throughout history and concludes that, since most of them were later proven wrong, currently accepted theories are likely to suffer the same fate. Accordingly, we should thus be *pessimistic* about the truth of currently accepted theories. Proponents of this argument have an abundance of examples at their disposal. In chemistry—or rather, prior to the advent of modern chemistry—it was believed that combustion was caused by the presence of 'phlogiston', which escaped burning objects in the form of fire and smoke (Yudkowsky 2007). In physics, light was long thought to be dependent on an invisible, omnipresent substance called 'luminiferous ether' (Torretti 2007). In biology, phenomena such as the appearance of maggots within animal carcasses were attributed to 'spontaneous generation'; under the right conditions, organisms supposedly materialised out of thin air without connection to any kind of life cycle (Corrington 1961). In all three of these examples, the consensus opinion was completely incorrect. These are not isolated examples either; the history of every field in science is littered with such cases. The pessimistic meta-induction thus directly challenges Oreskes's contention by highlighting the fallibility of consensus. After all, if her argument could be made at any point in history, even when we *know* the consensus was wrong, then why should we trust it today?

Systemic objection

Where the historical objection directly refutes Oreskes's *conclusion*, the systemic objection focuses on undermining the argument's *premises*. Specifically, it sows doubt in Oreskes's narrative of scientific scrutiny by pointing out systemic flaws in modern scientific practice that could lead to the formation of false consensus.

One of the most important principles in scientific research is *statistical significance*, which is typically expressed by what are known as '*p*-values'. Supposing we carry out a study and obtain some result, the *p*-value is the probability that we would see that result by random chance if the hypothesis we were testing was false. The lower the *p*-value, the higher the statistical significance, and the more confidence we can have that our result is not a false positive.

To guard against false positives, most journals will only publish results that meet a maximum *p*-value requirement called an *alpha* value, which is usually set at 5 per cent. In other words, only results that would occur by chance less than 1 time in 20 are allowed. Intuitively, one might infer from this that around 1 in 20 of research findings are false positives, but not more than that. As it happens, that would be a reasonable inference if exactly half of the hypotheses being tested were true, and half false. However, as Pashler and Harris (2012) prove by mathematical means, the expected proportion of false positives is much higher if the majority of hypotheses being investigated in the first place are false, and this is the far more realistic scenario.

The problems don't stop there. Because negative results typically aren't even *submitted* to journals, the expected rate of false positives as a proportion of *published* findings is even higher—this issue is referred to as publication bias. On top of this, since researchers face huge institutional pressures to find surprising and novel results in order to get published, many set up their studies in such a way that they can cherry-pick combinations of variables that happen to yield impressive *p*-values *after* they have all the results—this is known as *p*-hacking. Some go even further and commit fraud by fabricating data.

Of course, scientists aren't ignorant of the problem of false positives. One of the most powerful bulwarks against the proliferation of false ideas in science is also one of the oldest: replication studies. In essence, these involve a group of scientists testing the validity of a published study by conducting a new study following the same procedure and seeing if they come to the same result. If a hypothesis survives multiple replication studies, then it is much more likely to be correct. If it fails replication, then it must be discarded. What is worrying is that many fields are currently experiencing what has been dubbed a 'replication crisis', wherein landmark result after landmark result has been overturned as long overdue replication studies have been conducted (Aschwanden 2015).

In addition to false positives, there is a related problem of hidden cultural biases influencing the way research is actually conducted. Lloyd (1993) offers an illustrative example from primatology. In that field, it was once widely held that female orgasm in various species of primates *must* have evolved in order to fill some kind of reproductive function. There wasn't actually any evidence for this idea—in Lloyd's words it was a *pre-theoretical assumption* rooted in researchers' subconscious beliefs about female sexuality (it was, at the time, a male-dominated field). In fact, there was significant evidence to the contrary: female stumptail macaques almost always had orgasms when they mounted other females, but only occasionally had them during heterosexual coitus, indicating that it likely served a non-reproductive evolutionary function (Goldfoot et al. 1980). Despite this, at least one later study *intentionally* confined itself to the latter category of orgasms because of this unfounded assumption (Lloyd 1993: 142).

Publication bias, *p*-hacking, and even outright fraud have all contributed to many false hypotheses being wrongly accepted. In addition, pre-theoretical assumptions linked to conscious and subconscious bias contribute to many true hypotheses being wrongly denied. Taken together, the systemic objection maintains that these flaws in the process of scientific scrutiny *as it stands* are worrisome enough to cast severe doubt upon even the most seemingly well-established cases of scientific consensus.

Responses to the objections

While the historical objection has strong intuitive appeal, I believe it is the weaker of the two objections I've presented. There are at least three ways it can be significantly mitigated. Firstly, when you investigate the false theories used as evidence in the pessimistic induction, it doesn't always appear that they enjoyed the same level of consensus approval as, say, the theory of natural selection enjoys today. For example, some claim that there was a consensus in the 1970s that the Earth was cooling, but while a 'global cooling' hypothesis did exist, it was never more widely accepted than global warming—that is simply a myth (Peterson et al. 2008). Secondly, I contend that some of the examples offered by those making this objection, such as the humoral theory of medicine (Laudan 1981), cannot be deemed 'scientific' theories, at least in the sense that Oreskes describes, because they didn't undergo a comparable process of scrutiny during their heyday. Including them in an inductive argument against Oreskes's conclusion would thus be an effective straw-man. Thirdly, of the remaining false theories used in the induction, almost all were *less empirically successful* than any theory accepted by scientific consensus today (Park 2011). With these considerations in mind, Park (2011) uses the same inductive argument to reach an *optimistic* view of present scientific theories—if this does not provide independent support for Oreskes's conclusion, then it at least shields it to a meaningful degree against the pessimistic induction.

I contend that the systemic objection is actually much more problematic for Oreskes's argument because, unlike the historical objection, it specifically targets current scientific practice. The field of psychology, for example, is currently facing a crisis because many of its most famous ideas have been thrown into question by replication studies which have failed to find the same results (Aschwanden 2015).

A proposed amendment to Oreskes's argument

To address the aforementioned objections, I propose that Oreskes's argument be amended by introducing an additional condition for trust in consensus. Part of what makes her argument so strong already is its relatively narrow scope, and it can be made even stronger by being narrowed further. Oreskes doesn't try to explain exactly how scientists come up with their ideas or even the statistical methods they use to test them. Instead, she outlines a process of scrutiny that prevents false claims from reaching the status of consensus *regardless* of the methods used in a particular field. She then argues that this rigorous scrutiny justifies public trust in scientific consensus—nothing more, nothing less. But while she does a good job of defining 'scrutiny', and 'trust' is fairly self-explanatory, her precise notion of 'consensus' remains somewhat vague. The only thing she clarifies is that scientific consensus can only exist among experts *within* a specific field (that is to say, the opinion of neurobiologists cannot affect consensus amongst particle physicists, and vice versa). However, I believe that pinning down precisely what we mean by consensus is key to strengthening Oreskes's argument and neutralising these two objections. I believe this because, in my view, scientific consensus can actually take on two very different forms, and Oreskes's argument only applies to one of them. My amendment is therefore to add the following important caveat to the original argument: consensus may be either *active* or *passive*, and only the former kind warrants trust.

Active and passive: The two kinds of consensus

To count as *active* consensus, a claim must survive rigorous scrutiny to be actively endorsed by an overwhelming majority in a given field. An example of *active* consensus would be the widespread agreement among climate scientists—on the basis of incredibly robust evidence—that global warming exists and is caused by human activity. I contrast this with *passive* consensus, which occurs when a shared assumption is deemed scientific because it happens to be *believed* by most scientists in a field, but hasn't actually ever faced the kind of tough, public, and ongoing scrutiny that Oreskes describes. An example of *passive* consensus would the myth of 'spontaneous generation' in biology which I mentioned earlier. While it was certainly widely believed, it wasn't actually based on anything but supposition—unsurprisingly, it quickly fell apart when it was finally subjected to proper scrutiny (Corrington 1961). I would also label as *passive* consensus the previously cited example discussed by Lloyd (1993) of faulty pre-theoretical assumptions about female orgasm in primates.

By explicitly restricting itself to *active* consensus, this modified argument resists both the historical and systemic objection to a greater degree. That is because the examples of consensus used by both the historical and systemic objections almost all fall under the *passive* category, while the examples of consensus that Oreskes has in mind are *active* because they only came about after being subjected to rigorous scrutiny.

There does remain a significant problem, which is that laypeople will naturally have difficulty differentiating between *passive* consensus and *active* consensus. After all, almost by definition, laypeople are not able to directly observe whether the process of scrutiny was adequate when deciding whether to trust the scientific consensus on a given issue. This problem seems somewhat unavoidable, though I argue it is not a weakness in the argument I've presented so much as an indication that we need better reporting on science. It's one thing to talk about trusting consensus when it exists, but polling shows that a large number of Americans do not believe in the *existence of consensus* on climate change, and that appears to be largely due to the way the issue is covered in the media (Sanders 2019). I do think this problem is mitigated by the fact that *passive* consensus will likely fall apart if it becomes central to an issue of public importance. If a particular scientific consensus has survived mainstream news coverage and political debate for a significant period of time, then it is more likely than not to be *active* consensus.

Conclusion

In her article, Oreskes succeeds in presenting a novel argument that is not only convincing, but philosophically robust. She outlines in simple but effective terms the nature of scientific scrutiny with the three fundamental criteria of being *tough*, *public*, and *ongoing*. She demonstrates how this kind of scrutiny inevitably works to remove false claims from scientific discourse while allowing true claims to build on one another in accordance with Helen Longino's (1990) notion of transformative interrogation. Ultimately, Oreskes uses this to show why scientific consensus is so reliable, and hence trustworthy, even for laypeople.

However, in its original form Oreskes's argument was vulnerable to two major objections. The historical objection points to a myriad of examples of past scientific consensus being overturned to infer that even current scientific consensus is likely to be wrong, and hence not trustworthy. The systemic objection goes one step further and highlights the ways in which even contemporary science fails to live up to Oreskes's standards for scrutiny. Both objections can be responded to, but in order to properly neutralise them I propose we need to slightly amend the original argument. Specifically, we must distinguish between *active* consensus that was born from a process of tough, public, and ongoing scrutiny, and *passive* consensus that would be better described as pre-theoretical assumption than theory. If we discard the latter and specify that our argument applies only to the former, then I believe both the historical and systemic objections are satisfied. In the end, I thus concur with Naomi Oreskes that even modern science is far from perfect, but that we laypeople should still put our trust in *active* scientific consensus when it exists—especially when it is a matter of life and death.

References

Aschwanden, C (2015) 'Science isn't broken', *FiveThirtyEight*, 19 August. Available at: fivethirtyeight.com/features/science-isnt-broken/ (accessed 20 May 2020).

Corrington, JD (1961) 'Spontaneous generation', *Bios*, 32(2): 62–76.

Goldfoot, D, Westerborg-van Loon, H, Groeneveld, W and Slob, AK (1980) 'Behavioral and physiological evidence of sexual climax in the female stump-tailed macaque (Macaca arctoides)', *Science*, 208(4451): 1477–1479. doi.org/10.1126/science.7384791.

Laudan, L (1981) 'A confutation of convergent realism', *Philosophy of Science*, 48(1): 19–49. doi.org/10.1086/288975.

Lloyd, EA (1993) 'Pre-theoretical assumptions in evolutionary explanations of female sexuality', *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, 69(2/3): 139–153. doi.org/10.1007/BF00990080.

Longino, HE (1990) *Science as social knowledge: Values and objectivity in scientific inquiry*. Princeton University Press, Princeton, NJ. doi.org/10.1515/9780691209753.

Oreskes, N (2019) 'Science isn't always perfect—But we should still trust it', *Time*, 24 October. Available at: time.com/5709691/why-trust-science/ (accessed 13 May 2020).

Park, S (2011) 'A confutation of the pessimistic induction', *Journal for General Philosophy of Science*, 42(1): 75–84. doi.org/10.1007/s10838-010-9130-0.

Pashler, H and Harris, CR (2012) 'Is the replicability crisis overblown? Three arguments examined', *Perspectives on Psychological Science*, 7(6): 531–536. doi.org/10.1177/1745691612463401.

Peterson, TC, Connolley, WM and Fleck, J (2008) 'The myth of the 1970s global cooling scientific consensus', *Bulletin of the American Meteorological Society*, 89(9): 1325–1337. doi.org/10.1175/2008BAMS2370.1.

Sanders, R (2019) 'In media coverage of climate change, where are the facts?', *Berkeley News*, 19 September. Available at: news.berkeley.edu/2019/09/19/in-media-coverage-of-climate-change-where-are-the-facts/ (accessed 4 December 2020).

Torretti, R (2007) 'Getting rid of the ether: Could physics have achieved it sooner, with better assistance from philosophy?', *Theoria: An International Journal for Theory, History and Foundations of Science*, 22(3(60)): 353–374.

Tyson, A, Johnson, C and Funk, C (2020) 'U.S. public now divided over whether to get COVID-19 vaccine', *Pew Research Center*, 17 September. Available at: www.pewresearch.org/science/2020/09/17/u-s-public-now-divided-over-whether-to-get-covid-19-vaccine/ (accessed 4 December 2020).

Yudkowsky, E (2007) 'Fake causality', *LessWrong*. Available at: www.lesswrong.com/posts/RgkqLqkg8vLhsYpfh/fake-causality (accessed 2 December 2020).