



Science, assertion, and the common ground

Corey Dethier¹

Received: 19 July 2021 / Accepted: 20 December 2021
© The Author(s) 2022

Abstract

I argue that the appropriateness of an assertion is sensitive to context—or, really, the “common ground”—in a way that hasn’t previously been emphasized by philosophers. This kind of context-sensitivity explains why some scientific conclusions seem to be appropriately asserted even though they are not known, believed, or justified on the available evidence. I then consider other recent attempts to account for this phenomenon and argue that if they are to be successful, they need to recognize the kind of context-sensitivity that I argue for.

Keywords Assertion · Common ground · Context · Norm of assertion · Science communication · Social epistemology of science

1 Introduction

There is a growing literature on the relationship between individual norms of inquiry, norms of proper scientific practice, and the reward structures of science.¹ One theme that has emerged in this literature is the idea that the declarative sentences found in published scientific research are not (and should not be) judged by the same standards that we use to evaluate assertions in ordinary contexts.

¹ See, e.g., Mayo-Wilson et al. (2011), Zollman (2018), Heesen (2015), Bright et al. (2018), O’Connor (2019), Holman and Bruner (2017), Weatherall and O’Connor (2021), and Fleisher (2018, 2021).

✉ Corey Dethier
corey.dethier@gmail.com

¹ Leibniz Universität Hannover, Hannover, Germany

I think the opposite conclusion is correct: everyday and scientific assertions are to be judged according to the same standards.² Importantly, however, those standards are context-sensitive. More precisely, whether an assertion is appropriate depends on the “common ground” presupposed by the conversants in a way that hasn’t previously been emphasized by philosophers. Because scientists tend to assume that their conclusions will only be read by other scientists, they generally make very specific assumptions about the propositions in the common ground in a way that affects what is appropriately asserted.

An outline of the paper is as follows. Section 2 discusses the phenomenon that has motivated philosophers to posit a difference in how assertions are evaluated in scientific and ordinary contexts. Section 3 shows that in general, the appropriateness of an assertion depends on what’s in the “common ground” of the conversation. For example, if we’ve explicitly assumed that P is false, it’s not appropriate (within the context of that conversation) for me to represent myself as knowing (believing, justifiably believing, etc.) P , and thus not appropriate for me to assert it. Section 4 argues that scientists generally make relatively strong assumptions about the common ground when publishing their papers because they assume their audience consists solely of other experts. Together with the context-sensitivity of assertion, these assumptions explain the phenomenon laid out in Sect. 2.

In Sect. 5, I identify a class of cases in which scientists inappropriately assert: sometimes scientists assume that their audience consists solely of other experts when in fact they are addressing non-experts as well. My account offers a nice explanation of what’s gone wrong in these cases: the scientists have made a mistake about the character of the common ground. Finally, in Sect. 6, I argue that the cases described in Sects. 2 and 5 cannot both be accommodated without introducing a similar context-sensitivity mechanism to the one that I propose.

2 When scientists defend what they don’t believe

Though there’s debate about whether there are any special norms of assertion, it’s usually assumed that there are at least some necessary conditions on warranted or proper assertion.³ That is, there is some norm of the form “One may assert P only if ...”; various philosophers fill in the ellipses in different ways. The most common view, following Williamson (2000), is that one may assert P only if one knows that P . Other views focus on justification, reasons, or evidence (see, e.g. Lackey, 2007; McKinnon, 2015; Maitra and Weatherson, 2010); on truth (Weiner, 2005); on certainty (Stanley, 2008); or on mere belief (Bach, 2008).

Recently, however, philosophers such as Fleisher (2018, 2021) and Dang and Bright (2021) have argued that scientists frequently make claims that satisfy none of these accounts—that is, in which P is asserted but none of the proposed necessary conditions

² In what follows, I’m largely going to put aside the question of whether the declarative sentences found in scientific publications are “really” assertions. It seems to me that this question either boils down to the normative question—such sentences constitute assertions iff they are governed by the same standards as assertions—or is largely uninteresting.

³ For dissent, see Azzouni (2020, Chap. 4) as well as the citations therein.

are met. Dang and Bright (2021) offer the following schematic thought experiment to illustrate how this might go:

Zahra is a scientist working at the cutting edge of her field. Based on her research, she comes up with a new hypothesis. She diligently pursues inquiry according to the best practices of her field for many months. Her new hypothesis would be considered an important breakthrough discovery. Zahra knows that many more studies will have to be done in the future in order to confirm her hypothesis. Further, she has read the current literature and realizes that the existing research in her field does not, on net, support her hypothesis. She does not believe that she has conclusively proven the new hypothesis. Nonetheless, Zahra sends a paper reporting her hypothesis to the leading journal in her subdiscipline. In the abstract of the paper, the conclusion, and talks she gives on her work, she advocates for her hypothesis. Peer reviewers, while also sceptical of the new hypothesis, believed that her research had been carried out according to best known practices and her paper would be a valuable contribution to the field. Her paper, which purports to have advanced a new hypothesis, is published and widely read by members of her community. In subsequent years, additional research in her field conclusively demonstrates that Zahra's hypothesis was false (Dang & Bright, 2021, 8191).

Like Dang and Bright and Fleisher, I take it that this case is fairly representative: scientists commonly defend hypotheses that are not known, believed, or justified by the evidence.

We need to careful here, however. As both Dang and Bright (2021) and Fleisher (2021) explicitly recognize, the proposed norms of assertion are not violated if scientists merely "advocate for" or "advance" P when P is not known, believed, or justified by the evidence. So a scientist might "advance" P , for example, by asserting that " P has not yet been ruled out by the evidence, and should be treated as a live hypothesis." This is a way of advancing the hypothesis that P is true; it is something that it would be reasonable for Zahra to assert even if we suppose that she is constrained by the most restrictive of the proposed norms of assertion. Equally, Zahra might advance P by saying that "The evidence provided by the present study supports P ." As before, this is in some sense a way of advancing P , but it is one that Zahra could assert without believing or knowing P and even if the evidence on the whole favors $\neg P$.

The kinds of hedging just surveyed allows scientists to present their claims without (systematically) violating any of the proposed norms of assertion: I can know that P is supported by the study that I have conducted, for instance, without knowing, believing, or being justified in believing the proposition that P is true. Dang and Bright make precisely this point in clarifying their position:

We do not doubt that it would be possible to reform scientific communication behaviour such that one sticks to scientific public avowals that are proper according to one of the surveyed norms by insisting on appropriate hedging. ... But our concern is not really with the precise linguistic form such claims would take so much as the social uptake amongst scientists. However results are conveyed, scientists must decide what claims are worthy of further tests. ... Our point is

that it would be inappropriate for scientists to insist that (in the absence of fraud or mistake or misfortune) these pursuit worthy claims must be true, or justified, or believed to be as much by their proponents (Dang & Bright, 2021, 8198–99).

I'm more interested in the "precise linguistic form" than Dang and Bright are. As we'll see, there are important links between the precise linguistic form that scientific assertions should take and their social uptake by the audience. Or, in other words, I'll be arguing that even if Dang and Bright are right that neither truth, justification, nor belief is necessary for pursuit-worthiness (and I think they are right on this point), that doesn't mean that the statements of scientists shouldn't be assertions that are held to the relevant norm.

As the literature illustrates, however, there are cases in which scientists make assertions where the precise linguistic form of the assertion violates our everyday assertoric standards. For instance, Dang and Bright (2021) quote William Henry Bragg as asserting that "the experimental proof of the material nature of the γ rays carries with it, almost surely, a corresponding proof as regards the X-rays" (Bragg 1908, 270; see also the discussion of the evolution of dogs in Fleisher 2021). Bragg's assertion involves two different plausible violations of any standard assertion norm. First, it presupposes not only that γ rays are material, but that there is an "experimental proof" of such; Dang and Bright argue that Bragg himself did not believe this at the time and that he would not have been warranted in doing so.

The second plausible violation is more interesting. Bragg's assertion commits him to the claim that there is "almost certainly" a proof that X-rays are material, and thus, given that proof is factive, to the claim that it is "almost certain" that X-rays are material. Given Dang and Bright's analysis of the situation, however, it seems implausible that Bragg was in a position where this probabilistic assertion was warranted. While it is an open question how to extend standard accounts of assertion to probabilistic or hedged assertions (Benton & vanElswyk, 2020), it seems plausible that the assertion that P is "almost certain" will require either that P is in some sense objectively much more likely than not on the total evidence or that the testifier believes that P is much more likely than not on the total evidence, if not both. If, as Dang and Bright argue, Bragg neither believed that X-rays were material nor was in a position to justify this belief, then it seems unlikely that he met this minimal condition.

In what follows, I'm going to follow Dang and Bright in holding that Bragg did nothing wrong in these cases—or at least that there are nearby cases like this in which the relevant assertions were appropriate.⁴ The existence of such cases is thus a phenomenon in need of explanation. That is, I take that there are cases in which there is a face-value violation of any of the proposed norms of assertion yet we nevertheless want to say that the relevant declarative sentence were appropriate. In the literature to date, the dominant explanation of this fact is that the declarative sentences of science aren't governed by the same norms that govern assertions in ordinary life. I'm going

⁴ It's not clear to me that Bragg's assertions are appropriate; I'm tempted to say that he should have more carefully hedged his conclusions. For the sake of investigating the general claim, however, I'm going to assume that Dang and Bright's analysis is apt.

to argue for the opposite: contrary to appearances, the cases in question (largely) meet everyday standards for assertion, and where they don't, they are in fact inappropriate.⁵

3 Appropriate-in-a-context

It's often thought that the appropriateness of assertions can vary between contexts, though there is disagreement on how exactly this context-sensitivity works.⁶ In this section, I propose that whether an assertion is appropriate depends on the "common ground" in a way that has not previously been emphasized. To argue for this conclusion, I'm going to put aside questions concerning the proper formulation of *the* norm (or norms) of assertion; I take it that the proposed form of context-sensitivity is compatible with most (and possibly all) of the views in the literature.

Here's the basic picture, which I borrow largely from Stalnaker (2002, 2014). When we engage in conversation, there are various shared presuppositions about what the members of the conversation know (the "common ground").⁷ These presuppositions can be modified either explicitly or implicitly; one way that this can be done is by utterances such as "assume that *P*." My claim is that whether it is appropriate to assert *P* in a context depends not on what one actually knows (believes, justifiably believes, etc.), but instead on whether one can legitimately represent oneself as knowing (etc.) *given* the common ground of the conversation—or, more minimally those elements of the common ground that are either true or recognized as assumptions in the relevant context. My reasoning is simple and easily illustrated by assuming that knowledge is the relevant assertion norm. When we assume *P* or add it to the common ground, it's treated as common knowledge. But if *P* is common knowledge then I can't without contradiction represent myself as knowing (or even believing) $\neg P$. Similarly, when *P* is common knowledge, I can legitimately claim to know any proposition that I know to be a consequence of it. Insofar as the appropriateness of assertion tracks the the appropriateness of my claim to know *P*, the appropriateness of an assertion depends on what's in the common ground.

Most of the time, this kind of sensitivity isn't important, because most of the propositions in the common ground are known (or at least assumed) to be by the participants of the conversation). As noted above, however, there are mechanisms for introducing

⁵ Note that similar cases exist in philosophy. Indeed, a number of philosophers have run similar arguments to motivate the view that philosophical assertions need not be known, believed, or justified (see Barnett, 2019; Goldberg, 2015; Plakias, 2019). Importantly, the philosophy-focused discussion relies more heavily on substantive theses about disagreement and higher-order evidence that are outside the scope of the present paper. In what follows, therefore, I'll continue to focus on the scientific case. My arguments will mostly apply to philosophical assertions as well, however, and so if I'm right about the scientific cases, then there are at least some cases where philosophers are warranted in asserting *P* even though they do not know, believe, or justifiably believe it.

⁶ Various accounts are presented in DeRose (2002), Stanley (2008), Goldberg (2015), McKinnon (2015); and Turri (2010). None of these are obviously equivalent to the account offered here. It's likely, however, that the appropriateness of an assertion is sensitive to context in multiple ways and possible that there's some way of subsuming the kind of context-sensitivity I argue for under one of the other accounts

⁷ The notion of "presupposition" here should not be understood in terms of belief, but a broader kind of "acceptance": so we might presuppose *P* for the sake of argument, as a conjecture, or even as part of the background for a fictional story (Stalnaker, 2014, 39).

propositions to the common ground even when said proposition is not known, and even when it is known to be false. In these cases, common ground-sensitivity makes a difference to what it is appropriate to assert. Consider:

Kim and Mira are discussing whether P is true. They both think that Q , which seems relevant to P , is probably true, and explicitly agree that—for the sake of argument—they’re going to assume Q . After some thought, Mira realizes that $Q \rightarrow P$. She exclaims “Aha!” and then follows up with an explanation that begins with “So, Q , right?” and ends with “Therefore, P !”

There isn’t anything wrong with Mira’s assertions in this case, *even though* she doesn’t know either Q or P . Why not? The obvious explanation is that Kim and Mira have (explicitly!) presupposed that Q is true.⁸ (This feeling is reinforced if we alter the case so that Kim and Mira decide not to assume Q .) Accordingly, while Mira doesn’t in fact know P , in the context of the conversation she can legitimately represent herself as knowing P given that Q has been explicitly presupposed. After all, if we assume that her thinking is correct, then it’s true that if Mira did know Q , she would know P . Since both Mira and Kim have assumed that Q is common knowledge, their presuppositions, combined with what they know, entails that Mira knows P . So therefore there’s nothing epistemically inappropriate about Mira asserting P .

Of course, the legitimacy of Mira’s assertion of P is fragile in an important sense: as soon as Q is no longer part of the common ground, it becomes illegitimate for her to represent herself as knowing P . Given how easily the common ground can shift in a conversation—particularly in the context of explicit assumptions that can be withdrawn at any time—Mira’s authority to represent herself as knowing P can be eroded much more easily than her actual knowledge can be. But this fragility doesn’t undermine the point. Consider the following extension of the case:

“Of course,” Mira adds with a sigh, “we don’t actually know if Q is true.” Kim turns to her in confusion: “Well then why did you assert P ?!”

Kim is clearly making a mistake here in that she’s treating an assertion that was made in one context as though it were made in another. Kim isn’t in a position to criticize Mira’s assertion of P given that she explicitly agreed to assume a proposition that make it appropriate for Mira to assert P ; when that assumption is later canceled, it’s not appropriate to go back and reinterpret what was said given the new common ground.⁹

⁸ This is not the only possible explanation: we could also explain intuitions in this case by claiming that the proposition that Mira “really” communicates by uttering “ P ” is $Q \rightarrow P$. I’m skeptical that this is a plausible reading of this case, but there is an interesting question here: do standards of assertion govern what the speaker actually utters (P) or the proposition that is communicated (allegedly $Q \rightarrow P$)? The difference here won’t matter for my purposes: if one wants to insist on the latter, it’s easy enough to reinterpret what I have to say in what follows as claims about what’s communicated by scientific utterances that are ostensibly in violation of assertion norms.

⁹ A related question here is whether we can criticize Mira for asserting P even if Kim can’t. I don’t think so. To see why, complicate the case a bit. Suppose that Anne enters the room in between when Kim and Mira explicitly suppose Q and when Mira asserts P . Anne might reasonably exhibit some confusion: “Mira, why did you assert P ? It seems like what you’ve shown is $Q \rightarrow P$.” If we assume that Anne is not illegitimately butting into the conversation and has standing (as it were), this is a legitimate criticism: Q is not in the common ground of the Kim-Mira-Anne conversation. But Mira has an excuse, namely that they

There's a clear explanation for why assertion should be sensitive to the common ground in this manner. As Williamson (2000, 252, fn 4) notes, assertion is one of a class of speech acts that require some sort of authority to be appropriate—other examples might include commands, orders, apologies, promises, and witnessings. Williamson argues that in the context of assertion, the relevant authority is simply knowing P . When we, like Mira and Kim, suppose propositions, we change the authority calculus—most obviously, if we suppose P , then we're simultaneously supposing that no one knows $\neg P$ and thus that no one has the relevant authority to appropriately assert it. The same phenomenon occurs with other authority-dependent speech acts: if we've supposed that you're in charge, I can't appropriately give you a command; if we've supposed that I've done nothing wrong, I can't appropriately apologize for my actions, etc. In order to appropriately carry out these speech acts, the presupposition must first be canceled.

Two comments on the argument for context-sensitivity that I've given in this section. First, we might worry that there could be contradictory information in the common ground, or (more broadly) that the common ground encodes assumptions that are in some sense impossible. I must admit that I am not terribly worried about this possibility: counterpossible reasoning seems to be required for all sorts of purposes, and I see no reason to balk at this particular application. Perhaps we want to deal with such cases by invoking impossible worlds (Berto, 2017); perhaps we want to deal with them by invoking less powerful logics (Bueno & da Costa, 2007). Perhaps there are simply informal and pragmatic norms about when you're allowed to use explosion. To model how we reason, we need some way of accounting for reasoning under impossible scenarios, and it seems plausible that whatever tools we invoke in those contexts will do just as well for complicating the present analysis. I'll be putting these technical questions aside in what follows because I take it they aren't necessary for making the present case.¹⁰

Second, and more importantly, notice that changes to the common ground can both shrink and expand what it is appropriate to assert. It can shrink it because we can explicitly remove from consideration propositions that are in fact known; it can expand it by adding to the common ground propositions that we do not know. Moreover, recall the point in the last section that many of the plausible violations of our common assertoric practices involve strictly unwarranted probabilistic assertions such as Bragg's claim that X-rays are "almost certainly" material. If we suppose that Bragg was justified in making *some* hedged assertion regarding the materiality of X-rays, then it seems

Footnote 9 continued

assumed that Q before Anne entered the room. At this point, it seems that all Anne has left to criticize is the assumption of Q itself; *given* that assumption, she's not in a position to criticize the assertion of P . Insofar as we see ourselves as having grounds for criticizing Mira, it's because we're confusing our third-person perspective with that of Anne's; we're working from the perspective of someone who is in the conversation but not governed by the (explicitly laid out) common ground. And this doesn't look like a legitimate position to take.

¹⁰ There are other interesting technical considerations that may be more relevant. So, for instance, the analysis given above might be in need of complication if there are important differences between cases where we add propositions to the common ground in an indicative mode and those in which we add them to the common ground in a subjunctive mode (see Eva et al., forthcoming). This is a possibility to be explored in future work.

highly plausible that changing the common ground can alter the degree of hedging that is appropriate in *either* direction. So, for example, if we explicitly remove from the common ground the evidence that undermines Bragg's hypothesis, he will be warranted in making a hedged assertion with a higher degree of confidence than would otherwise be justified. By contrast, there are propositions that we might assume that would lower the probability that he could legitimately assign to the hypothesis.

The crucial upshot of the discussion in this section is that the appropriateness of one and the same assertion can vary depending on the common ground, and thus the audience, of said assertion. In what follows, I'll be making this kind of variation do a lot of work; as we'll see, the apparent violations of the norm of assertion identified by Dang and Bright and Fleisher can and should be explained by recognizing that scientists are (or really, take themselves to be) speaking to a very specific audience.

4 Scientific publishing and the common ground

In this section, I'll argue that the cases identified as potential violations of individual norms of assertion can be explained by the fact that scientists *tend* to assume that the audience for their cutting-edge research consists solely of other experts—and, in so doing, they make very specific assumptions about the common ground in their published papers.

It is a truism that the best methods of communication are sensitive to the composition of the audience. If you are reporting new results to a room full of non-experts, the best approach is to carefully lay out all of the prior results relating to the hypothesis under test before turning to a discussion of your new results and how they fit into the already-established picture. On the other hand, if your audience consists solely of experts, it's more efficient to put aside earlier results as much as possible and only outline the newest results available. Because the experts are (by definition) well-informed on the subject already, they are in a position to evaluate your new results without an extensive presentation of the relevant background. The resulting efficiency isn't a purely pragmatic good. If every paper or presentation had to thoroughly review all of the prior evidence, the dissemination of knowledge would be substantially hindered.

The "putting aside" described in the second strategy means treating previous results in one of two ways. Some results—those that are sufficiently well-known and that have agreed-upon implications for the hypotheses being tested—can simply be treated as elements in the common ground. Not all prior results have this status, however. Insofar as new results address an open question in the field, there's liable to be differences in how the experts evaluate at least some subset of the prior results (see Solomon, 2001). If there's no uncontroversial evaluation of these prior results to be had, however, then there is (by definition) no common ground regarding the open question. There's no—or, more realistically, limited—agreement about what we ought to believe going into the evaluation of the new evidence. In these cases, it thus makes sense for scientists to "bracket" the relevant prior results—that is, treat them as unknown—rather than building them into the common ground. This "bracketing" move allows the disagreeing parties to all evaluate the new evidence for themselves given their different views on what is warranted by the prior results.

Of course, the two strategies that I've outlined are idealized extremes; in real life, actual presentations of new results are likely to fit somewhere in between. The thesis I defend in the rest of this section is that, in general, scientific articles tend to be written for other experts and thus tend to approximate the second strategy—that is, when presenting cutting-edge research, scientists tend to presume that some uncontroversial prior results are known while bracketing evidence or results with controversial implications for the study at issue.

I take it that the part of this thesis that asserts that scientists aim their cutting-edge research at other experts is obvious. One can tell from the language used in cutting-edge scientific publications that the target is other scientists. And while your average scientific paper includes a short lit review at the beginning of the paper, these lit reviews are usually designed more to motivate the question that the paper asks than to genuinely evaluate the strength of *all* the prior evidence related to the hypothesis under study.

The claim that scientists often bracket contentious prior evidence is more interesting, and I'll offer two additional pieces of evidence for it. The first piece of evidence comes from the use of Bayesian methods. Since these methods explicitly involve the use of priors, we would expect that if background knowledge about the probability of the hypothesis under test was to show up anywhere, it would show up here. But it doesn't, at least not consistently: scientists who employ Bayesian methods are often not concerned with selecting priors that match either their own prior beliefs or what they take to be the beliefs that are most warranted on the prior evidence. Instead, they're concerned with either selecting a prior distribution that is as close to mathematically unbiased as possible or with demonstrating that the results of the given study hold for a variety of potential prior distributions.¹¹ This is exactly the behavior that we would predict if scientists are bracketing the old evidence that's immediately relevant to whether or not their hypothesis is true.

Further evidence comes from the presence of “review articles” that summarize the previous results relating to some hypothesis.¹² Such articles are quite common—there are entire journals dedicated to them—and they are often treated as *de facto* authorities: citations to a review article ends up standing-in for citations to an entire literature (McMahan & McFarland, 2021). Importantly, these review articles do not generally seem to be directed entirely or even primarily at the public. Instead, they're directed at other scientists working in the field.

The existence of review articles directed at other experts would be surprising if it was expected that every paper would itself review the prior evidence, or even all of the contested evidence, in a thorough way. In that case, there would be no need to

¹¹ Some examples from climate science: Schnur and Hasselmann (2005) survey a range of priors in their discussion of attribution climate change to humans; the common “regularization” method for estimating internal climate variation is based on what is essentially a (nearly) uninformative prior (Ribes et al., 2013); and even where climate scientists use priors that are in some sense subjective (as in the estimation of “climate sensitivity” found in Sherwood et al. 2020), they generate said priors by polling various experts in the field, which I take to indicate a desire to survey all of the plausible interpretations of the prior evidence. For one theoretical defense of these kinds of practices, see Gelman and Shalizi (2013).

¹² Review articles have unfortunately received little attention from philosophers of science, though see George E. Smith's historical work on Newtonian astronomy (Smith, 2014) and Avogadro's number (Smith and Seth, 2020).

survey all of the different studies, because each study would do so itself. But both the existence and the centrality of review articles is to be expected if the standard practice is to bracket a subset of the prior evidence. One of the main worries raised by bracketing is that it encourages a piecemeal and partial examination of the evidence, and that it thus can inhibit recognition that some hypothesis has been conclusively settled or blind us to worrying unexamined alternatives. Review articles act as checks on these concerns.

As we saw in the prior section, the kind of bracketing that I'm claiming is common in the sciences influences which assertions are appropriate. In particular, it will often be appropriate for scientists to assert that P is more or less probable than they in fact believe it is on the grounds that some of the relevant total evidence has been (explicitly or implicitly) bracketed for the purposes of presenting the new evidence and evaluating its implications. So, for instance, suppose that the evidence collected in the study supports P ; it will then be appropriate for the scientist to assert that P is probable *regardless* of what the bracketed evidence says.¹³ On the view presented, such assertions don't violate our everyday assertoric standards insofar as scientists are correct that their audience is composed of fellow experts who (a) are themselves aware of the bracketed evidence and (b) understand that this evidence has been bracketed because they understand the publication norms of the discipline. Such an audience is in a position to uptake the newly presented results in the way that is in fact warranted by the total evidence.

The practice of bracketing thus explains the phenomenon of scientists seeming to offer unwarranted assertions without requiring the systematic discrepancy between our everyday assertoric practices and the assertoric practices we find in the sciences. On the picture I offer, science publishing is a conversational "game" in which there are tacit norms about what you can and cannot assume. These norms influence what you can represent yourself as knowing in a given context, and thus what it is appropriate to assert within that context. It's not that these assertions aren't governed by standards norms, but rather that the tacit norms that influence the common ground figure into our evaluation of an assertion in a predictable way: what you can legitimately represent yourself as knowing in the context of a scientific paper is (sometimes) different from what you in fact know.

It's worth worrying whether this picture really diverges from that offered by others. After all, here's how Dang and Bright characterize their positive view:

We think that what is required here is some form of contextualist epistemic norm. ... Context, provided by the history and present consensus in a field, specifies some amount of the previous literature one must survey to check for coherence, and which methodological procedures one must carry out to reach a conclusion that is worth reporting to others. One's avowal must be such that if one's total evidence were what one had gathered in the methodologically proper way for one's latest study, combined with whatever one has taken from the mandated subset of the previous literature contains, then one would be justified in believing one's scientific public avowal (Dang & Bright 2021, 8199–20).

¹³ This claim should be qualified with a "within reason" clause; as I'll motivate in the next section, many of the normative distinctions should be treated as relatively blurry.

Both Dang and Bright and I therefore offer contextualist readings of the norms operative in science, and indeed, we agree that what's appropriate in science is largely a function of the study in question and not a function of all of the other evidence that has previously been collected.¹⁴ So while there are minor differences in how we spell out the relevant context-sensitivity, it seems plausible that the two positions will largely agree on individual cases.

This agreement on individual cases shouldn't blind us to a central—and important—theoretical difference. The previous literature on this subject, including both Dang and Bright (2021) and Fleisher (2018, 2021), has claimed that the assertions found in scientific publications can be appropriate without living up to our normal assertoric standards.¹⁵ I claim that scientific assertions must meet these standards if they are to be appropriate. This isn't a verbal dispute. It's a dispute between a view on which scientific assertions are governed by fundamentally different norms than regular assertions and one on which there are contextual factors that can be found throughout our assertoric practices that show up in a systematic way in scientific discourse. On my view, the nature of legitimate assertion doesn't fundamentally change when in the context of scientific publishing. On the contrary, the legitimacy of assertion is always sensitive to context, and scientific assertions occur in a context that is structured in very specific ways.

I think it tells in favor of my view that it better unifies our assertoric practices. More importantly, though, I think that there are cases that motivate thinking that any successful account of our assertoric practices must build in something like the kind of context-sensitivity that I've argued for in the last two sections. The next section spells out these cases and shows that my account gives the right result; the final section argues that it's impossible to get the right result without something akin to the sort of common ground-sensitivity that I'm arguing for.

5 False assumptions about the common ground

Dang and Bright explicitly limit their discussion of scientific avowals to what they call "inter-scientific" discussions:

Public avowals in science should not be confused with extra-scientific testimony or "public scientific testimony." Statements from scientists aimed at policy-makers are not the type of utterances we are interested in. It is important to distinguish between claims aimed at the scientific community and claims aimed at the general public or policy-makers. The IPCC assessment report on climate change, for example, while made publicly, is primarily aimed at political bodies, and such testimony is properly held to a different standard. Extra-scientific testimony is hence not the target of our paper (Dang & Bright, 2021, 8190–91).

Fleisher (2021), similarly, seems to be concerned primarily with how scientific assertions are interpreted by other researchers. That is, these alternative accounts make a

¹⁴ On my reading, Fleisher (2021) agrees on the first of these points but not the latter one.

¹⁵ Perhaps this is too strong: given Dang and Bright's disinterest in the "precise linguistic form," it might be better to say that Dang and Bright offer a picture of scientific assertions in which obedience to norm of assertion doesn't even come up.

similar assumption to the one I've claimed is operative in the sciences, namely that the audience of the relevant assertion consists of other experts in the field, and *not* the general public.

This assumption is not always justifiable, either for scientists or for philosophers of science theorizing about the norms governing science. The public nature of journal articles means that anyone with the appropriate resources or know-how can access any published result. While it might once have been reasonable to suppose that no one without the relevant background would pick up a print copy of the journal and read through the conclusions in its pages, that assumption is no longer warranted: online repositories and open access publishing have made it increasingly easy to access scientific publications and thus increased the probability that the actual audience of a paper will be composed partly of non-experts. This possibility matters: it's at least possible for a scientific assertion to be inappropriate because it was made under the assumption that the audience would be composed entirely of experts when in fact the audience includes non-experts.

Let me push the point on the level of the theory first. Speakers (or in this case, authors) don't have infallible access to the common ground; they can be wrong both about what their audience knows or presupposes and about who is in their audience. And, as such, while they might think that their assertions are appropriate given their mistaken assumptions about the context, they can nevertheless fail to make appropriate assertions because the common ground is not as they believe it to be. In many cases, of course, these gaps between intended and actual audience are either harmless or excusable, and so even if we think that there is something bad about the assertion itself in the context in which it is uttered, intuitions will differ in these cases as to whether or not it was genuinely inappropriate.¹⁶ In at least some cases, however, a speaker's assertion will be genuinely inappropriate because they failed to account for the fact that their actual audience was sufficiently likely to be different from their intended audience.

As an example where a speaker makes this kind of mistake in a situation where it isn't easily excusable, consider the following case:

Edward is a climate scientist working on the cutting-edge of climate change attribution. He's found a new way of estimating the human contributions to climate change, and this new line of evidence—taken by itself—would strongly indicate that humans are less responsible than previously thought. Edward knows that if one were to thoroughly review his results in the context of the prior three decades worth of research, the most likely conclusion would be that there's something wrong with his new method. But he can't find any errors in his method and he's done nothing that violates standard practice in the field. He also knows that denialist news outlets trawl new publications for results that go against the consensus in climate attribution research—we can suppose that he's completely certain that if he reports his research in a way that easily allows for a denialist reading, his research will be picked up by these outlets and misinterpreted. Nevertheless, he publishes a paper in which he claims that the new evidence

¹⁶ I suspect that intuitions here simply track those at the ethical level—i.e., the question is simply whether I'm wrong for committing an action that caused unforeseen harms.

“strongly indicates that humans have contributed less to global warming than previously thought.” His colleagues, who are aware of the thirty years of research that Edward bracketed, reasonably interpret these results as primarily calling for further investigation into this new method. Predictably, however, his research is picked up by denialists in the media and used to justify continuing inaction on climate change.

Edward has done something wrong. What he’s done wrong, most obviously, is that in presenting his results the way he does, he’s failed to account for the fact that his research will be disseminated to those outside of his intended audience, and thus for the fact that his actual audience includes people who don’t have the scientific background necessary to understand that his assertion concerning the state of the evidence doesn’t take into account the previous thirty years of research on the subject. And while this insensitivity to one’s actual audience might be harmless in many cases, Edward has failed to account for his audience in a case in which this gap between intended and actual audience is extremely predictable and where the ethical stakes are extremely high.

My view therefore both says that Edward has asserted inappropriately in this case and provides a natural explanation of what’s gone wrong: Edward’s case is akin to that of an actor who shouts “Fire!” in a situation where she knows that her audience is unaware that she’s acting.¹⁷ In both cases, the error is one of failing to take account of the audience that one actually has.

In the next section, I’ll argue that accommodating this case requires common ground-sensitivity. But first a few comments on cases like Edward’s. First, on the picture I’ve offered, a difference between the intended and actual audience of a paper affects whether the assertions found within that paper are appropriate. I’ve motivated this view by focusing on a clear cut case: Edward *knows* that his audience is going to include non-experts. Most cases won’t be like Edward’s. It’s often extremely difficult to predict who is going to read a given piece of writing. As such, the view I offer requires scientists to evaluate tradeoffs between presenting their research in a way that is most useful to other experts and presenting it in a way that will lead to the fewest misinterpretations. I suspect that in most cases, whether we think a scientist has asserted appropriately will depend on whether they had sufficient reason to *expect* that their assertion—or the misinterpretations of their assertion—would cause harm in virtue of reaching unintended audience members. So our evaluations are likely to be much murkier in practice than Edward’s case might make it seem.

Second, I’ve offered a relatively simple (and I think highly intuitive) view of how scientists should present their research, namely that it depends on the audience. In cases where there’s a relatively low chance of the results being read by non-scientists, the shared conventions and assumptions of science are such that it’s sometimes permissible to assert what you don’t strictly speaking know (believe, justifiably believe); the audience will understand that this assertion should be interpreted in a particular way. In cases where there’s a sufficiently high chance of the results leaking out into the

¹⁷ Perhaps not quite: maybe there’s an important difference between asserting-under-a-pretense and the kind of context-sensitivity that I’ve claimed is at work in Edward’s case. The case can be modified to make the analogy tighter, however.

public sphere, by contrast, scientists should explicitly *hedge*—that is, explicitly qualify their avowals to make it clear how they should be interpreted. They often do hedge in these cases: IPCC reports such as IPCC Working Group 1 (2013), for example, are an extraordinary source of hedged assertions.

Third, and in some tension with the last point, I think there's a case to be made that scientists should generally hedge more (perhaps much more) than they in fact do—as indicated in note 4, I'm tempted to say this about Dang and Bright's Bragg example. As anyone who has taught introductory philosophy of science can tell you, misconceptions of the methods of science are relatively common among non-scientists. In the popular imagination, a single study is usually definitive; scientists go out, they investigate the world, and they come back with the truth. I think this view is harmful, both because those who hold it are too trusting of individual scientific results and because it so easily lends itself to crass conspiracy theories. If there's truth-revealing science and biased science and nothing else, then any area in which results are messier must involve bad faith actors.

Of course, the general populace aren't usually reading scientific papers directly—they get information about scientific discoveries through second- or third-hand summaries. Nevertheless, we should prefer norms of science that discourage this kind of view, all other things being equal. Plausibly, that means that in general scientists should habitually hedge their claims, because doing so presents a more accurate picture of how science actually works. Notice: if I'm right on this front, the context-sensitive account of assertion that I've defended explains why: even if the immediate audience for any given scientific conclusion is largely going to be other experts, there's consistent enough divergence between the intended and actual audiences of scientific conclusions that misrepresentations that are harmless individually can add up to cause harm over time by reinforcing false views about how science works.

6 Escaping Edward's case

In the last section, I identified a class of cases involving inappropriate assertions, and showed that my account could explain what was inappropriate in these cases. This section discusses various ways of reconciling the view that it is sometimes appropriate for scientists to assert what they do not know (believe, justifiably believe) with the cases outlined in the last section. I argue that any successful account must build in common ground-sensitivity.

The most obvious means of responding to Edward's case is simply to deny that Edward has done anything wrong. But rejecting the view that Edward has done *something* wrong here just doesn't look plausible. Inductive risk arguments in philosophy of science (e.g. Douglas 2000) and encroachment arguments in epistemology (e.g. Fantl and McGrath 2002) both motivate *at least* the view that whether one should assert or publicly commit oneself to *P* depends on moral and pragmatic considerations of the “stakes” involved in the assertion.¹⁸ This kind of sensitivity to stakes

¹⁸ Indeed, many of these arguments claim to go much further, arguing that one's evidence and epistemic reasons for believing cannot be separated from moral and pragmatic reasons (see, e.g. Basu 2019; Basu and Schroeder 2019).

has clear theoretical justification. Borrowing phrasing from Maitra and Weatherston (2010), whether one should assert that P is just a specific instance of the question of whether ϕ is “the thing to do” in a given situation—and obviously, whether ϕ is “the thing to do” depends on the moral and pragmatic implications of ϕ -ing. Indeed, while there’s substantial disagreement about *how* to explain assertion’s sensitivity to stakes (see note 6), I think it’s fair to say that the consensus is that stakes play some sort of role in determining whether an assertion is appropriate. It’s hard to argue that Edward hasn’t asserted inappropriately.

If we accept that Edward’s assertion is inappropriate, there are two species of response for distinguishing Edward’s case from presumably appropriate cases like Bragg’s. The first is to explain why Edward’s assertion is inappropriate in a different way. In particular, Dang and Bright (2021) appeal to the standards of the scientific community in question in their positive view of what it is appropriate for a scientist to assert. One way of accommodating Edward’s case is thus to say not that Edward has done something wrong, but that there’s something wrong with the community standards.

So, for instance, we might say that climate scientists ought to (and perhaps do) hold themselves to a higher standard than Edward does given the ethical import of their research. Though natural, this move seems to get nearby variations on the case wrong. We can imagine that Edward’s case differs in various ways that alter the nature of his audience but not anything else. So, for instance, imagine that Edward gets the opposite result: his new findings reinforce the prior thirty years of research and so won’t be picked up by denialist news outlets. Or imagine that he lives in a world in which climate science is not the political issue that it is in our world, and so there’s no worry about anyone other than experts reading Edward’s paper. Or, finally, imagine that Edward is in a situation where for structural reasons his audience is limited to only experts—perhaps he’s at a conference where the working papers are not accessible to the public. In all of these cases, Edward’s assertion appears to be appropriate; after all, he’s lived up to the standards of field, and the audience is in a good position to understand his assertion and the background evidence against which it occurs and thus to evaluate it appropriately. The only difference between these cases and the original case is how likely it is that his assertion will reach an audience of non-experts.

These variations on the case make it clear that what makes Edward’s assertion inappropriate is precisely the composition of his audience. The only way to account for these variations, then, is to build common ground-sensitivity (or something effectively akin to it) into the standards of the science itself. It doesn’t seem likely to me that common ground-sensitivity is part of the standards of science, but (more importantly) even if it is (or should be) the reason why seems to be that common ground-sensitivity is part of what makes an assertion appropriate. So even if this move works, it only provides a partial explanation of why Edward has gone wrong; the ultimate explanation seems to be that he’s violated the norms of everyday assertoric practice.

The final line of response is to argue that while Edward’s assertion is inappropriate, it’s not inappropriate *qua* scientific assertion—it’s inappropriate for some other reason. Perhaps, for instance, an opponent could claim that Edward has done something that’s scientifically legitimate but morally wrong, and argue that I’ve conflated the scientific

dimensions of evaluation with the moral one in the analysis of the case.¹⁹ To my eyes, this move looks ad hoc and undermotivated. While there are real arguments to be had as to whether Edward's *evaluation* of the evidence should depend on moral considerations, in this case we only require the much less radical view that his actions should. And we're happy to freely mix together scientific and ethical considerations in other respects: failure to secure ethics board approval is a sufficient reason to reject or retract a paper; the paper's scientific qualifications are undermined by its questionable ethical status. Further, it seems unlikely that this line of argument is compatible with holding that scientists can sometimes appropriately assert what they do not know (believe, justifiably believe): if we're going to distinguish between a "purely" scientific evaluation of assertion and the pragmatic ethical evaluation of that same assertion to account for what Edward has done wrong, why not use the same distinction to account for the cases that motivate the literature in the first place?

A more promising way of running this response involves distinguishing between two different kinds of assertion ("scientific" and "everyday," maybe) or between two different roles that Edward plays ("colleague" and "expert"). Based on this distinction, one might argue that Edward's behavior is appropriate in one of these two respects but not the other. So, for instance, we might hold that Edward acted appropriately *qua* colleague but not *qua* expert.

In broad strokes, this is essentially the move that Fleisher (2018, 2021) makes. For present purposes, Fleisher's account has three commitments. First, a kind of context-sensitivity in which the standards of appropriateness of an assertion depends in part on the purpose or goal of the conversation. Second, a view according to which researchers sometimes put propositions forward for *advocacy* purposes and sometimes put them forward as *evidence*. Third, the claim that it is appropriate for researchers to assert *P* in an advocacy context so long as they "endorse" it, where endorsing is a cognitive attitude that is sensitive to both epistemic and pragmatic reasons.

So Fleisher can respond to Edward's case by arguing that Edward's assertion is appropriate *qua* advocacy role but not evidence-updating role. The question is then simply which of the two roles he is playing at that time. If Fleisher accepts my interpretation of Edward's case, he could say that Edward asserts as though he's (clearly?) in an advocacy context despite knowing full well that his statement will be taken to be evidential. In this respect, he can simultaneously uphold a principled distinction between scientific assertions (or at least assertions *qua* advocacy) and our everyday assertoric practices while giving the correct analysis of Edward's case.

I'm not enthusiastic about the resulting picture, as I think that the lesson that we should take from cases like Edward's is that this kind of principled distinction doesn't exist. Scientists, like people more generally, are almost always performing many different roles at once—colleague, advocate, expert, critic, friend, etc.—that place different obligations and responsibilities on them. Our assertoric standards should be sensitive to these differences, but (it seems to me) the messy hybrid case is the normal one,

¹⁹ Note that I think this is a misdescription of the case. What Edward has done wrong concerns assertions *qua* assertions. In most cases we forgive the kind of mistake he makes because they are harmless. In Edward's case, however, the mistake isn't harmless. (Compare shouting "Fire!" when there's some evidence of a fire, but not anything close to sufficient evidence of such.) This description begs the question by presupposing the context-sensitivity that I argued for above, however, and so my response won't rely on it.

meaning that it's not usually appropriate to say that a particular assertion should be judged as an advocacy-role assertion. Instead, the question is how the role of advocate affects the scientist's overall collection of responsibilities, which is what matters for the question of whether they asserted appropriately. I find it unnatural to say that a scientist acted appropriately qua advocate but not qua presenter of evidence; at least in a healthy scientific community, failing in the second respect seems to entail failing in the first.²⁰ The account that I offer thus rejects these distinctions entirely in favor of a view in which our assertoric practices are continuous but context-sensitive. I think it ought to be preferred.

I haven't offered an argument for that conclusion, however, and Fleisher gives his own arguments for the principled distinction between advocacy and evidence-updating; addressing these would take us too far afield from the main takeaway of the present section.²¹ That takeaway is this: in order to account for both cases like Bragg's and cases like Edward's, we need the appropriateness of an assertion to be sensitive to the common ground in a way that allows for the possibility that a scientist can assert inappropriately due to mistaken assumptions about said common ground. Because the barriers between scientific and non-scientific communities are porous at best, these mistakes are going to happen. Regardless of whether or not we ultimately want to make sharp distinctions between our everyday assertoric practices and those found in science, therefore, our standards need to be able to account for the interactions that happen at the border between the two domains. And so regardless of whether we prefer continuous or discrete accounts of our assertoric practices, we're going to need some mechanism that accounts for how mistakes concerning the common ground can shift the standards of appropriate assertion.

7 Conclusion

I've argued that the appropriateness of an assertion is sensitive to the common ground of a conversation in a way that hasn't previously been emphasized by philosophers. This kind of context-sensitivity explains why some scientific (and philosophical) conclusions seem to be appropriately asserted even though they are not known, believed, or justified on the available evidence: scientists who assume that their audience contains only scientists make very specific assumptions about the common ground, assumptions that are rational in that they allow for efficient scientific communication. I then showed these assumptions can go wrong when scientists end up with an audience that is composed partly of non-experts, with the result that the relevant assertions are no longer appropriate. This class of cases provides us reason to think that any successful account of our assertoric practices is going to have to build in some form of common ground-sensitivity, because there's no other good way to accommodate both cases like Edward's and the assumption that it is sometimes appropriate for scientists to assert what they do not know, believe, or justifiably believe.

²⁰ Compare: we *could* say of an actor who shouts "Fire!" to an audience who is unaware that she is acting that her shout is appropriate qua acting but inappropriate qua assertion, but this feels like unnatural hair-splitting. Imagine trying to use "I acted appropriately qua acting" as an *excuse*.

²¹ See in particular Fleisher (2021, §5).

In the interest of living up to my own earlier statements about hedging, I note that this starting assumption is exactly that. It seems to me to be plausible that scientists can sometimes appropriately assert what they do not know (etc.), and I've argued for a couple of (even more) plausible claims that would explain why this is the case. Nevertheless, I think the arguments for this assumption are nowhere close to definitive: it remains relatively likely that this element of scientific practice ultimately can't be justified on epistemic grounds and is instead a kind of unwarranted imprecision in presentation that is widespread because scientists are prone to the same epistemic mistakes as the rest of us.

Acknowledgements I would like to thank Haixin Dang, Liam Kofi Bright, and two anonymous reviewers for constructive comments on an earlier version of this paper. The paper also benefited from an audience at Leibniz Universität Hannover as well as conversations with Will Fleisher, Samia Hensi, and Mack Sullivan.

Funding Open Access funding enabled and organized by Projekt DEAL. Funding for this paper was provided by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) - Project 254954344/GRK2073.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Azzouni, J. (2020). *Attributing knowledge: What it means to know something*. Oxford University Press.
- Bach, K. (2008). Applying pragmatics to epistemology. *Philosophical Issues*, 18(1), 68–88.
- Barnett, Z. (2019). Philosophy without belief. *Mind*, 128(509), 109–38.
- Basu, R. (2019). The Wrongs of racist beliefs. *Philosophical Studies*, 176(9), 2497–515.
- Basu, R., & Schroeder, M. (2019). Doxastic wronging. In B. Kim & M. McGrath (Eds.), *Pragmatic Encroachment in epistemology* (pp. 181–205). Routledge.
- Benton, M. A., & van Elswyk, P. (2020). Hedged assertion. In S. Goldberg (Ed.), *The Oxford handbook of assertion* (pp. 245–63). Oxford University Press.
- Berto, F. (2017). Impossible worlds and the logic of imagination. *Erkenntnis*, 82(6), 1277–97.
- Bragg, W. H. (1908). The nature of and X-rays. *Nature*, 77(1), 270–71.
- Bright, L. K., Dang, H., & Heesen, R. (2018). A role for judgment aggregation in coauthoring scientific papers. *Erkenntnis*, 83(2), 231–52.
- Bueno, O., & da Costa, N. C. A. (2007). Quasi-truth, paraconsistency, and the foundations of science. *Synthese*, 154(3), 383–99.
- Dang, H., & Bright, L. K. (2021). Scientific conclusions need not be accurate, justified, or believed by their authors. *Synthese*, 199(3–4), 8187–203.
- DeRose, K. (2002). Assertion, knowledge, and context. *The Philosophical Review*, 111(2), 167–203.
- Douglas, H. (2000). Inductive risk and values in science. *Philosophy of Science*, 67(4), 559–79.
- Eva, B., Shear, T., & Fitelson, B. (forthcoming). Four approaches to supposition. *Ergo*.
- Fantl, J., & McGrath, M. (2002). Evidence, pragmatics, and justification. *The Philosophical Review*, 111(1), 67–94.
- Fleisher, W. (2018). Rational endorsement. *Philosophical Studies*, 175(10), 2649–75.
- Fleisher, W. (2021). Endorsement and assertion. *Noûs*, 55(2), 363–84.

- Gelman, A., & Shalizi, C. R. (2013). Philosophy and the practice of Bayesian statistics. *British Journal of Mathematical and Statistical Psychology*, 66, 8–38.
- Goldberg, S. (2015). *Assertion: On the philosophical significance of assertoric speech*. Oxford University Press.
- Heesen, R. (2015). How much evidence should one collect? *Philosophical Studies*, 172(9), 2299–313.
- Holman, B., & Bruner, J. (2017). Experimentation by industrial selection. *Philosophy of Science*, 84(5), 1008–19.
- IPCC Working Group 1. (2013). *Climate Change 2013: The Physical Science Basis*. In T. F. Stocker (Ed.), *Fifth Assessment Report of the Intergovernmental Panel on Climate Change*. Cambridge: Cambridge University Press.
- Lackey, J. (2007). Norms of assertion. *Noûs*, 41(4), 594–626.
- Maitra, I., & Weatherson, B. (2010). Assertion, knowledge, and action. *Philosophical Studies*, 149(1), 99–118.
- Mayo-Wilson, C., Zollman, K. J. S., & Danks, D. (2011). The independence thesis: When individual and social epistemology diverge. *Philosophy of Science*, 78(4), 653–77.
- McKinnon, R. (2015). *The norms of assertion: Truth, lies, and warrant*. Palgrave Macmillan.
- McMahan, P., & McFarland, D. A. (2021). Creative destruction: The structural consequences of scientific curation. *American Sociological Review*, 86(2), 341–76.
- O'Connor, C. (2019). The natural selection of conservative science. *Studies in the History and Philosophy of Science Part A*, 76, 24–29.
- Plakias, A. (2019). Publishing without belief. *Analysis*, 79(4), 638–46.
- Ribes, A., Planton, S., & Terray, L. (2013). Application of regularised optimal fingerprinting to attribution. Part I: Method properties and idealised analysis. *Climate Dynamics*, 41(11–12), 2817–36.
- Schnur, R., & Hasselmann, K. (2005). Optimal filtering for Bayesian detection and attribution of climate change. *Climate Dynamics*, 24(1), 45–55.
- Sherwood, S. C., et al. (2020). An assessment of earth's climate sensitivity using multiple lines of evidence. *Review of Geophysics*, 58(4), e2019RG000678.
- Smith, G. E. (2014). Closing the loop: Testing Newtonian gravity, then and now. In Z. Beiner & E. Schliesser (Eds.), *Newton and empiricism* (pp. 262–351). Oxford University Press.
- Smith, G. E., & Seth, R. (2020). *Brownian motion and molecular reality: A study in theory-mediated measurement*. Oxford University Press.
- Solomon, M. (2001). *Social empiricism*. The MIT Press.
- Stalnaker, R. (2002). Common ground. *Linguistics and Philosophy*, 25(5/6), 701–21.
- Stalnaker, R. (2014). *Context*. Oxford University Press.
- Stanley, J. (2008). Knowledge and certainty. *Philosophical Issues*, 18(1), 35–57.
- Turri, J. (2010). Epistemic invariantism and speech act contextualism. *The Philosophical Review*, 119, 1.
- Weatherall, J. O., & O'Connor, C. (2021). Conformity in scientific networks. *Synthese*, 198(8), 7257–78.
- Weiner, M. (2005). Must we know what we say? *Philosophical Issues*, 14(2), 227–51.
- Williamson, T. (2000). *Knowledge and its limits*. Oxford University Press.
- Zollman, K. J. S. (2018). The credit economy and the economic rationality of science. *The Journal of Philosophy*, 115(1), 1–33.