

THE BELIEF NORM OF ACADEMIC PUBLISHING

WESLEY BUCKWALTER

*Department of Philosophy, Institute for Philosophy and Public Policy,
George Mason University*

The belief norm of academic publishing states that researchers should believe certain claims they publish. The purpose of this paper is to defend the belief norm of academic publishing. In its defense, the advantages and disadvantages of the belief norm are evaluated for academic research and for the publication system. It is concluded that while the norm does not come without costs, academic research systematically benefits from the belief norm and that it should be counted among those that sustain the practice of academic publishing.

Keywords: belief; norms; research; truth; knowledge; publishing; methods; inquiry

A belief of the primary sort is a map of neighbouring space by which we steer. It remains such a map however much we complicate it or fill in details.

– Frank Ramsey

PUBLISHING in academic journals is a heavily norm-governed activity. It is reasonable to suppose that the practice is sustained by many norms. One norm that may sustain the practice is that authors believe what they publish. According to this norm, when researchers publish certain claims during scholarly inquiry, they should think that those claims are true. Like many norms, this is not an exceptionless rule, but rather, a general expectation regarding what authors should do. According to this expectation, authors need not believe every single word they publish. Inquiry sometimes proceeds by contesting prior data or claims, labouring under conditionals, or by making reductio style arguments which posit false assumptions. Rather, the expectation under the belief norm is that many published statements should be

Contact: Wesley Buckwalter <wesleybuckwalter@gmail.com>

believed. For the purposes of the present investigation, a tentative list of such statements includes the central thesis statement and the core contribution being made to the research record, statements concerning the central pieces of evidence presented in support of the thesis, statements made about how evidence was collected and analyzed, statements concerning the instruments and procedures used, descriptions of the views of others, and that these things otherwise meet field specific standards for publication. While it is tempting to conclude that a belief norm of some form is accepted in several fields, the matter has not been systematically investigated.

Several researchers have questioned the value of belief in various aspects of publishing, research, inquiry, and methodology (Barnett 2019; Beebee 2018; Dang & Bright 2021; Fleisher 2020; Millson 2020; Palmira 2020; Plakias 2019). With respect to academic publishing, researchers have argued that the belief norm should be limited or rejected (Dang & Bright 2021; Fleisher 2020; Plakias 2019). Just as accepting the norm does not mean that every published sentence should be believed, rejecting it does not mean every disbelieved sentence should be permitted. For example, it would not be permissible to publish discriminatory comments just because one happens not to believe them, or to represent data as reliable that one strongly believes are fraudulent (Fleisher 2020). Though the details of these arguments vary, the unifying idea is that scholarly inquiry would be improved if there was no expectation that many published statements, and perhaps even statements concerning core evidence, contributions, and conclusions are believed. On such a view, for example, it could be permissible for researchers to publish the conclusion that the universe is a computer simulation, even though they do not actually think we are living in a simulation.

The purpose of this paper is to offer a systematic defense of the view that authors should believe what they publish. To do so, the paper will proceed by broadly evaluating whether academic research would be better off with or without the belief norm of academic publishing. In the following section, several arguments in favor of the belief norm are presented which suggest that academic research and publishing are better off with the belief norm. The section after that one considers several objections to the norm and argues that several negative consequences are outweighed by the benefits of the norm. It is concluded that while the norm does not come without costs, academic research has systemically benefited from the belief norm of publishing and that it should be counted among those that sustain the practice. The final section summarizes these arguments, discusses their limitations, and suggests avenues for future research on the norms that govern academic research.

Before proceeding, several clarifications are needed regarding the shape and scope of the argument. Though the scope of the belief norm is admittedly vague, the present focus will be on statements concerning the central research contributions of published papers as described above. The argument for this will be in some ways broader and in some ways narrower than arguments that have come before. While some arguments against the belief norm reject belief for all public avows that

researchers make (Dang & Bright 2021), the present arguments will be restricted to central claims published in peer reviewed academic journal articles. And while some arguments against the belief norm are constrained to certain fields of study such as philosophy (Plakias 2019), the present arguments for the belief norm will not be restricted by discipline. Instead, the paper will argue that the belief norm of academic publishing should apply to those publications reporting truth-conducive inquiry. Many researchers publishing in philosophy journals as well as researchers in several other fields across the humanities, natural, and social sciences take themselves to be reporting inquiry of this sort and therefore belong to this category. And while it is widely agreed that a major goal of research published in scholarly journals is to reach the truth, no argument in the paper depends on it being the only goal a researcher could have or on this fully representing the activity of individuals. Rather than focusing on any one area or discipline, the argumentative strategy will be to canvas a broad range of considerations that many researchers face when conducting inquiry and present them collectively as evidence in favor of the norm.

1. Advantages of the Belief Norm

There are many ways in which the belief norm benefits academic publishing. These advantages should be counted in favor of the norm.

1.1. Inquiry Aims at Truth and Belief is a Good Guide to Truth

One goal of inquiry is to discover truth. A long and venerable tradition in epistemology and philosophy of mind holds that there is a reliable connection between belief and truth (Shah & Velleman 2005; Velleman 2000). According to this tradition, beliefs constitutively aim at truth. To believe a proposition carries with it a commitment to the truth of that proposition. The fact that belief is truth-directed in this way is often taken to reveal several further properties of belief and these properties make it a good candidate for a norm of academic publishing.

One property is that belief is involuntary (Alston 1989; Bennett 1990; cf. Turri, Rose, & Buckwalter 2017; Williams 1973). Thomas Hobbes, for example, writes that we have no choice but to “believe whatever is proved by valid arguments or is related in a credible manner, whether we want to or not” (Hobbes 1641/2006: Obj. 13). The reason, it is often thought, is because forming beliefs based on our inclinations or desires would be to form beliefs irrespective of their truth, which would violate the constitutive aim of belief. Instead, beliefs are formed respective of truth-conducive considerations, such as evidence. Thus, if belief aims at truth in this way, and belief is the norm of academic publishing, then published claims will be

more likely informed by truth-conducive considerations than those that are published but not believed. Of course, this does not mean that published claims will always be true or even probably true. What it does mean is that the fact that one believes a proposition is always at least some evidence for its truth. If this widely held view of belief is correct, then all else equal, that proposition is more likely to be true, even if ever so slightly, than a proposition that one does not believe to be true. In short, the norm increases the likelihood of fewer false statements.

It is also thought that the truth-directedness of belief distinguishes belief from other attitudes (Engel 2004). Consider other attitudes researchers might have toward their published statements if not belief. For example, they might be curious, ambivalent, or indecisive about the claims they publish. Or they might claim that what they publish is true based on a larger rhetorical strategy to shift the literature in a certain direction or persuade other researchers to study it (Fleisher 2020). Research based on such attitudes may indeed have positive outcomes. However, these attitudes can be distinguished from belief in that they lack constitutive connections to truth. If so, then all else equal, we can expect that claims based on attitudes that lack constitutive connections to truth will be less likely to be true than those based on attitudes with a constitutive connection to truth. And while other closely related attitudes sometimes associated with publishing such as suspecting, wondering, exploring, or worrying may sometimes have similar epistemic benefits to believing, they will probably not lead to as many true statements as believing.

1.2. The Norm Encourages Us to Qualify Our Claims and to Modulate Their Strength to Our Evidence

Consider the cases of the “Repentant (non) Realist” and the “Doubtful Graduate” introduced by researchers as evidence against the belief norm of academic publishing (Plakias 2019: 639). In the former, a researcher publishes a response that they find compelling to an objection offered against non-naturalism, thereby defending non-naturalism. But they do not believe that non-naturalism is true. In the latter, a researcher publishes an argument against a popular position p even though they suspect that p might be correct. Furthermore, their argument against p relies on empirical evidence they suspect is false. In both cases, it might be thought, the belief norm would lead to worse outcomes by preventing publication of these research articles.

However, the norm does allow publication of these projects in at least two forms. First, the core research claims being made in these examples could be publishable, according to the norm, after they are qualified. The norm does not prohibit these researchers from claiming to defeat an objection to non-naturalism or that empirical evidence supports $\sim p$, so long as those more specific things are what gets published and believed. Publishing arguments one thinks are right for a

view one thinks is wrong for other reasons is not incompatible with the norm. Second, the papers could be publishable, according to the norm, without the stronger claim that the arguments in question entail that the disbelieved views are true. Specifically, the norm would discourage the claim that a defeated objection to non-naturalism demonstrates non-naturalism or that questionable empirical evidence demonstrates $\sim p$ is true. This is a good thing because they likely don't do that.

Given that the papers are publishable in these forms, the norm promotes two epistemic benefits. For one, the belief norm encourages clarity and caution about the arguments being made. In this hypothetical, it was a condition of publishing that researchers clarify their arguments involving the defeat of an objection or the presentation of empirical evidence in support of their argument. For another, the belief norm helps modulate the strength of the claim published to the strength of the evidence supporting it. A defeated objection to a theory does not entail that the theory is true. And claiming that a theory is true when one has good evidence to the contrary can be intellectually irresponsible and harmful to the research record. The norm makes these important errors salient and helps us to avoid them.

The case involving empirical evidence is more complicated, but the result is the same. It would be wrong to claim that data were reliably obtained when they are known or strongly believed to be fraudulent. And perhaps one should hesitate to publish an argument based on data that one believes are poor. Drawing on data one has some reason to doubt is less clear. Based on the nature of this doubt, the norm might lead researchers to make weaker claims with respect to what they believe those data show. Or the norm might lead researchers to share their doubts and make more circumscribed claims about the way those data support some theoretical position. In both cases, the norm promotes good epistemic practices that increase the likelihood that published claims are well supported. It may potentially even help alert researchers to the possibility of mistakes or fraud and slow the spread of unreliable evidence. These practices amount to researchers making more careful claims in ways that better match the strength of their evidence.

One example raised by prior researchers in the history of science involves physics and claims made about neutral material particle theory (Dang & Bright 2021). In this example, Sir William Henry Bragg published a series of articles claiming that γ and X rays may be or were even "almost surely" material (Bragg 1907; 1908). However, researchers have pointed out that copies of Bragg's personal correspondence around that time call into question whether Bragg believed that the rays were material (Dang & Bright 2021: 8194). Instead, it appears that Bragg only believed that the theory seems "to be the best model to be devised at present" and that he had "no right to claim more" (Wheaton 1981). With this latter statement, Bragg seems to acknowledge the belief norm. And, as researchers note, Bragg is careful (in at least some publications) to make conditional statements or to hedge his claims according to what he thinks is only likely to be true (Dang & Bright

2021: 8193). It is unclear whether Bragg's beliefs fully match his hedged assertion above or if a more thorough analysis of the historical record would reveal other publications that violate the belief norm. Perhaps an argument could be made that a stronger hedge would have better aligned with his beliefs or improved inquiry. In the present quotations however, it appears that there is evidence that the norm was being acknowledged and that it did encourage norm-concordant behavior in published statements. If this is correct, then the example is useful for the present argument because it illustrates how acknowledging the belief norm and encouraging statements in accordance with it may have led to progress in science.

1.3. The Norm Promotes Honest Engagement and Good-faith Argumentation for Functional Research Communities

Recent years have seen high-profile hoaxes published in academic journals (Mounk 2018; see Fleisher 2020, for discussion) and some have offered examples involving hoaxing as evidence against the belief norm (Plakias 2019: 639). Such hoaxes are often perpetrated by researchers who believe that the claims they publish are false, foolish, or absurd. In some cases, they may even know this. Hoaxes are not necessarily perpetrated with bad intentions. To the contrary, they might also be done with the intention of revealing weaknesses in the peer review system. Audits can be good. Every so often they can be, at least. Would inquiry be improved if there was a general expectation that academic publications were not hoaxes?

It is difficult to see how a system of academic publishing could be sustained for long with hoxing as anything other than a rare exception. An incredible amount of academic research depends on sincerity and trust. We trust that researchers share the goal of inquiry to discover the truth and pursue this goal in good faith. We trust that researchers are motivated to apply terms, describe their evidence, and represent others as accurately as possible. And we trust that journals are committed to these norms and uphold them judiciously when they select articles for publication. We trust researchers because disseminating findings would be difficult or impossible without this foundation. Researchers trust journals because the process of peer review and editorial decisions are often opaque or anonymous. And members of the general public trust the process when they support public funding, attend universities, and accept discoveries that might one day radically impact their lives.

A publication system without an expectation against hoaxing erodes this trust. On the practical level, hoaxing wastes resources with respect to refereeing, space, and editorial handling that delays or prevents the publication of deserving papers. On the social level, researchers will be less likely to productively engage with research they suspect is insincere or might make them look foolish. And there is probably no quicker way to undermine public trust in research than to normalize the publishing of papers hoaxing it. To be clear, these are empirical hypotheses about what might

happen to research communities if there was no rule against authors publishing certain kinds of claims they believe are false. But ask yourself: would you want to submit or review for journals where this was permissible? Would you want your work published alongside hoax papers or trust research that has? And given the risks involved, would you be less likely to research such topics, in the first place?

It may be objected that even if the practical and social consequences of normalizing hoaxes are a disaster, good faith argumentation is not required for evaluating research contributions. Instead, it might be thought, what matters for evaluating research contributions is ultimately a matter of assessing the quality of the arguments printed on the page for all to see. Additionally, it might be thought, it is difficult or impossible for peer reviewers or other evaluators to determine what the mental states of authors are anyway. Together, these considerations point to a powerful reason to overlook belief, and perhaps mental states entirely, in our epistemic evaluations of published journal articles.

It may be true that we sometimes do not care or cannot determine what the beliefs of individual authors are when we read their papers. Nonetheless, the expectation that they believe still benefits the system of academic publishing. The reason is because the norm licenses an assumption of sincerity that helps us to evaluate the quality of research contributions more accurately and efficiently. One reason for this is that conducting and evaluating research are difficult enough activities to do well without someone actively trying to disrupt them. Another reason is that evaluating contributions sometimes depends on factors that go beyond the printed word. In the natural and social sciences, for example, much of the research process goes unreported. We often have no choice but to trust that researchers did what they say they did and did not do what they failed to disclose. And it can be inefficient to track down every citation or supplementary detail. In the humanities, we often trust that researchers are representing the views of those cited accurately and that they aren't leaving out important information discovered during their scholarship. Sometimes, the lived experiences and social situations of authors are also thought to be important for evaluating research on matters of social and ethical concern. Other times, community and contextual values could be important to conducting and evaluating research. Perhaps it is possible that extensive academic detective work could reveal that researchers did not do, value, or experience what they claim to. But this often takes considerable time and resources to uncover, at which point there may have been significant damage to the research record. The norm can help minimize these negative outcomes by limiting the degree to which this happens.

1.4. The Norm Curates the Research Record

Even in many of the smallest fields, thousands of research papers are published each year. There are simply far too many papers that can be accepted for publica-

tion, let alone read, cited, or meaningfully engaged with. One positive aspect of the norm is that it limits the research record to papers with core claims that are believed. This is a positive outcome because given the constitutive aim of belief, the research record is more likely to include claims that are true, clear, epistemically humble, and better modulated to the strength of evidence with the belief norm in place. Again, this is not to say that publications that are not believed cannot have these features. For example, such a paper might usefully draw scholarly attention to a new area or finding (Dang & Bright 2021; Fleisher 2020, though see below). But given that the number of research papers published far exceeds what can be attended to anyway, we might as well have papers that direct attention to new areas and have the positive epistemic features that come with belief. Research programs which build on a firm foundation are those most likely to thrive.

1.5. The Norm Stems From Aspects of Our Psychology That Would Be Difficult or Unwise to Ignore

It is possible that researchers could publish a paper with almost any attitude about it. But typically, this is not the way human beings approach inquiry in the long run. Instead, inquiry usually proceeds in steps as one discovers, reasons, and infers things across several months or years establishing research programs. Spurred on by intellectual curiosity, researchers build from their evidence to advance new claims and discoveries. Neither do researchers typically select topics, spend their time collecting evidence, or make arguments with no cognitive pro attitudes toward what they are studying. Rather, researchers typically gravitate to topics that have some personal significance, given their training and backgrounds. This will result in an immensely complex system of beliefs that researchers draw on to characterize the world. Some have even likened belief to a map that helps us navigate more smoothly or efficiently (Ramsey 1931). In other words, beliefs play important inferential and motivational roles that guide and structure future activity.

It is unclear if it would be possible to interrupt these organizational and motivational aspects of belief just in case researchers are publishing an academic journal article, or if it would be wise to do so. It is possible that when researchers publish core contributions in academic journal articles, they have any attitude towards them. But we would typically expect a strong cognitive pro-attitude. That is, given the costs, risks, personal demands, and sheer hard work required for pursuing scholarly research, it would be surprising if publication was cordoned off from people typically thinking the things that they were pursuing were right. The reason this is surprising is because it seems unlikely to sustain the practice of publishing in the months and years required to pursue research questions over a scholarly career.

Of course, it is possible that researchers are sometimes motivated by epistemic states that are related to but fall short of belief. For example, researchers might be motivated by feeling a worry about an argument or having a doubt about a conclusion. The present suggestion isn't that these states cannot be valuable sparks to inquiry. Rather, the suggestion is that doubts or worries often lead to thinking that something about them is true when we publish claims they motivate, and these latter sorts of beliefs are what tend to ground continued and systematic scholarship in the long run.

1.6. Ordinary Norms Are Much Stronger, but Academics Can't Even Believe?

It is initially unclear whether claims in published research articles constitute assertions like those that we make in everyday life or if assertions made in scholarly contexts should be judged by the same standards (see Dethier 2022 for an argument that they should). However, a comparison between published claims and ordinary assertion may still be informative insofar as both represent important ways that we communicate information with one another. When it comes to ordinary assertion, at least, a strong case has been made that knowledge is the norm of assertion (Turri 2016; Williamson 2000). One powerful argument for this view comes from the observation that the knowledge norm can be traced throughout our phylogenetic and ontogenetic pasts (Turri 2017). Specifically, a strong case has been made that factive norms govern communication of very young children and many animal species. This suggests that the norms that sustain the practice of assertion have played significant roles in human development and have benefited many species across the animal kingdom. Of course, the fact that strong norms have been selected for so heavily across various systems of animal communication does not necessarily mean that it is optimal in other contexts. However, this fact is one reason why we should take this norm under serious consideration in scholarly communication too. At the very least, it should be considered against other norms philosophers have defended for related communicative practices.

As it happens, the belief norm under consideration for academic publishing is among the weakest of those defended in the context of ordinary assertion (Gerken & Petersen 2020). Most of the research on ordinary assertion, by contrast, favors much stronger norms in terms of justification, truth, knowledge, or even certainty. Many epistemologists think these stronger norms are constitutive of ordinary assertion and that there is evidence that they sustain the practice even in cases where rather mundane information is exchanged. This fact is striking because the subjects of scholarly communication are often far from mundane. Researchers often communicate highly impactful information that has the

potential to affect health outcomes, expand our understanding of the universe, reveal and respond to injustice, or have other important societal consequences. We must then ask ourselves: why should professional academics charged with discovery and learning at the highest levels of society be held to a far lower standard than those that govern the way we communicate information in everyday life, and perhaps even the standards of toddlers, vervet monkeys, and sparrows?

2. Disadvantages of the Belief Norm

There are some ways in which the belief norm might not benefit academic publishing. These disadvantages should be weighed against the benefits of the norm.

2.1. *The Norm Punishes the Blessed*

If publishing requires belief and belief is involuntary, then researchers cannot fully control what they believe. Continuing this line of thought, involuntarists often say that we can only believe in accordance with our evidence. Of course, others throughout the history of philosophy such as Augustine, Aquinas, Pascal, or James may have disagreed. But the possibility leaves open the following concern. If there are large individual differences in what researchers regard as credible or consider to be good evidence, and evidence limits what it is possible to believe, then this could potentially disadvantage researchers who have higher standards of evidence than others (Plakias 2019: 640). Researchers with higher standards will believe less and thus publish less than those researchers who have lower standards. While this is an important challenge, it is unclear that it creates a unique problem for the belief norm over other norms of academic publishing, that it constitutes a limitation, or that the limitation it constitutes is a disadvantage.

First, the involuntarist objection claims that some researchers may not be able to believe, and thus may not be able to publish claims because their standards of evidence required for belief are higher than others. But the fact that some researchers could profit from having lower evidential standards than others is a problem for many norms of publishing. For instance, it is also a problem for evidence norms. An evidence norm which says that publishing claims requires having strong evidence might also disadvantage conscientious researchers. Such a norm would also be susceptible to the worry given that some researchers have different opinions about what constitutes good evidence than others do, or are more conscientious at collecting it, paying attention to details, and so on, than others are. For example, suppose one researcher does not put much stock in certain fMRI evidence given widespread technological errors and low replication rates in neuroscience

(Eklund, Nichols, & Knutsson 2016), while another continues to generate that fMRI evidence uncritically. Because they are unconscientious, they do not double check calculations or notice software errors that conscientious researchers would. In this case, the former researcher will have a harder time publishing claims than the latter researcher because they will ultimately have less evidence. There is a sense in which it is regrettable that the conscientious researcher will have a more difficult time publishing work in their field than the unconscientious researcher will. But it is tempting to think that the right response is not to relax evidence norms so that the conscientious can publish poor work too. Instead, it seems that the right response is to uphold stronger norms to improve research and publication.

To take another example, suppose two researchers have different views about the role of peer disagreement in scholarly activity. One treats peer disagreement in an area of research as evidence that claims are misguided, while the other dismisses the idea that disagreement constitutes evidence. The latter researcher will have a much easier time publishing claims based on good evidence in fields with high disagreement, such as philosophy. It is currently unknown what the best response is to peer disagreement of this sort and whether disagreement should be counted as evidence or by how much. The present point is just that rejecting the belief norm would not solve this problem. The problem will continue to persist if two researchers accept an evidence norm of academic publishing and have different views about what constitutes evidence.

A second response is that the norm does not substantively limit the subject matter that researchers can publish about. Suppose a conscientious researcher with high evidential standards and an unconscientious researcher with low evidential standards both had access to the same strong but imperfect evidence that our universe is a computer simulation. The unconscientious researcher believes based on this evidence and publishes the conclusion that “we live in a simulation”. The conscientious researcher does not believe this. They want more evidence. But they are not barred from publishing their work. Instead, they publish the claim that “evidence strongly suggests we live in a simulation”. If they’re even more cautious, as was the case in at least one recent article on the subject, they might publish that the “Chances Are about 50–50” (Ananthaswamy 2020). As this example illustrates, conscientious researchers can publish almost anything that less conscientious researchers can publish about so long as their claims are proceeded with qualifications that bring them into alignment with their beliefs and evidence. This does not seem like much of a limitation. Careful qualification and hedging are common features of successful academic writing.

A third response is that it is unclear that if qualification is a significant limitation, it is one that significantly disadvantages researchers. Insofar as the norm modulates claims to the strength of evidence, belief improves the likelihood that the claims a researcher publishes are true. For this reason, having high evidential stan-

dards may ultimately help a researcher make more true discoveries over the course of their careers. And the greater connection between belief and truth than between truth and other mental states may also help researchers make discoveries more efficiently by minimizing wasted time and resources on dead ends and false starts. On the merits, these are strong advantages that researchers are bound to welcome.

What about practical disadvantages relating to status and promotion, does hedging damage career prospects? It might be thought that stronger claims are rewarded over weaker claims. One implication of this is that researchers who publish stronger claims may be more likely to get more papers published in better venues and have more successful academic careers. There are certainly many examples where brash publications have been rewarded. But it is also unclear that this is the rule, and the matter requires further study. In the long run, publishing strong unqualified claims can also have high social costs that harm career prospects. In philosophy, for example, the “Reluctant (non) Realist” and “Doubtful Graduate” will likely be viewed negatively for making unsupported inferences and uncaredful appeals to empirical science. Similarly, in the physics example, the researcher making the hedged claim is bound to be viewed more favorably than another flat out asserting the claim that we live in a simulation. And the former was still published in *Scientific American*. In any event, researchers will need to decide for themselves what status and promotion is worth.

Recent initiatives in social science and the open science movement in response to the replication crisis might also indicate that belief, hedging, and qualification are becoming more highly regarded. In social psychology, for example, the “loss of confidence” project is a publication initiative designed to be “an academic safe space for researchers to declare for all to see that they no longer believe in the accuracy of one of their previous findings” (Resnick 2019). The purpose of the project is to improve the self-correcting nature of science by making error correction more efficient and to normalize individual self-correction as a routine part of science (Rohrer et al. 2021). The success of this project is itself a testament to the association between publishing and belief. Belief encourages and motivates researchers to correct the research record in a way that they might not have otherwise felt without the norm. And learning what researchers believe helps the community to better assess and respond to areas of scholarship.

Moving forward as a research community, belief sharing also promotes more responsible research practices. With respect to future publications, perhaps the broader goal of the loss of confidence project is to prevent errors by fostering a culture of understanding and intellectual humility in scientific communities. While fostering intellectual humility might involve doing many things, one aspect of it is encouraging scientists to monitor their confidence in findings and to better modulate this confidence to the strength of their evidence given the possibility that they might be wrong. This suggests that professional incentives

have shifted toward rewarding more modest claims. This is a promising start toward improving the quality of research publications and it is guided by belief.

2.2. The Norm Precludes Collaboration, Especially Large Ones

Collaborations sometimes involve hundreds or even thousands of authors (see, for example, G. Aad et al. ATLAS Collaboration 2015). Collaboration rates differ between fields. Large collaborative publications are common across the social and natural sciences. They are rarer, by comparison, in certain humanities fields. Nevertheless, researchers in most fields collaborate with one another from time to time. And when this happens, it might be suggested, it is unlikely that all authors believe the claims in the paper they put their names to. Call a published paper where at least one coauthor does not believe its core claims an “incredulous collaboration”. To the extent that rates of incredulous collaboration are high among coauthored publications, accepting the belief norm may have significant costs. For example, it might suggest that many collaborations should not have been published, or that some contributors to them should not have been authors. This represents a significant challenge to the belief norm worthy of serious consideration.

One thing that an advocate of the belief norm might say in response to this challenge is that the presence of conflicting norms is an unavoidable outcome of any sufficiently complex norm-governed activity. It is likely that there will always be at least some tension between norms pertaining to individual researchers and those that involve promoting a diversity of perspectives when researchers are working collectively. That said, however, critics of the belief norm may be unsatisfied by this response. What is needed, it might be thought, is an accounting of how such norms interact to benefit the activity when they conflict. To that end, the following considerations are offered that, while not eliminating the tension between collaborative publication and the belief norm, do attempt to clarify that conflict and address its negative consequences.

A first response is that incredulous collaboration is relatively rare and that many genuine instances of it probably should be avoided. At first glance, this may seem implausible. After all, researchers can be extremely fastidious (picky) people. This makes it unlikely that researchers in collaborations would ever fully agree about a manuscript. Concessions and compromises must be made for collaborations to succeed. But recall that the belief norm says that researchers do not have to agree about every sentence of a paper. The version of the norm presently under consideration holds that at least core claims and substantive contributions to the published research record should be believed. And it also allows that almost anything can be published with the right qualification. These qualifications make incredulous collaboration rarer than it may first appear. It would

be surprising if researchers continue to collaborate with each other up until the point of publication if they did not even agree with even the central contribution of the paper (they can be picky, after all!). And if the research team cannot find a way to phrase even the central contribution of their paper, then maybe they should think twice about publishing it.

A second response is that collaborations benefit from the belief norm when there is disagreement among members of a research team. The existence of the norm encourages collaborators to publish core claims they think are true. So, when team members disagree about a core research claim being made in the paper, the belief norm offers guidance and suggests a practical strategy for how to resolve such disagreements. Namely, the researchers should modulate the strength of core claims until team members agree that they are supported. Given a connection between belief and truth, the resulting publication will be more likely to include claims that are better supported by evidence than those in which there is no norm encouraging authors to do this.

One demonstration of this is the practice of adversarial collaboration (Clark & Tetlock in press; Mellers, Hertwig, & Kahneman 2001; Tetlock & Mitchell 2009). Adversarial collaboration is a technique for resolving epistemic impasses between researchers. It occurs when researchers with opposing views come together in a single paper to specify empirical tests of hypotheses, conduct research, and to publish their findings in a way that satisfies all authors. It is often suggested that this method of publishing is more effective than publishing individual reply articles by researchers with different beliefs. Because researchers disagree about many of the core claims before the research begins, adversarial collaboration forces researchers to carefully specify their predictions, experimental designs and procedures, and limitations of theories in response to findings in a way that all authors can live with. Sadly, adversarial collaborations are rare, and are perhaps rarer still when they are needed the most (Tetlock & Mitchell 2009). This supports the point that incredulous collaboration is also relatively rare. The practice also demonstrates the epistemic benefits that can occur for individuals and communities of researchers when belief is the salient organizing principle for guiding collaborative publication.

A third response is to deny that incredulous collaboration is prohibited by the norm by appealing to group beliefs. There are many views that one might have about the metaphysics of groups and the relationships between attitudes and authors (Bright, Dang, & Heesen 2018; Wray 2017; Fagan 2011; Gilbert 1992; 1987). One view of collaborative publication is that groups of authors hold attitudes that reduce to the attitudes of each individual author. But another view is that collaborations can express group attitudes. On some of these views, social groups can hold certain states that are distinct from those of individual members, such as acceptance (Wray 2017) or belief (Gilbert 1992; 1987). If this is true, then it might also be

possible that a group of researchers has attitudes toward core research claims of a collaborative publication that are irreducibly the views of the group. Putting these pieces together, then, a possible response to the objection is that it might also be possible that the group thinks the core research claims of a paper are true when some individual authors do not. In short, incredulous collaborations may not violate the belief norm because authors are groups and groups can believe.

A fourth response is that some incredulous collaborations may be exceptions to the rule. On this response, certain incredulous collaborations might belong to a category that researchers have labeled “blameless transgression” (Turri & Blouw 2015). To explain, for nearly any rule, it is possible to imagine a way to blamelessly break that rule. For example, it might be thought that the fact that someone secretly tampered with your speedometer is a good excuse for why you were driving over the speed limit. Likewise, one possibility is that very large collaborations with thousands of authors can sometimes be a good excuse for violating the belief norm, since it is impractical to check the beliefs of all authors. The absence of blame for breaking these rules does not entail that a rule was not broken. Moreover, the fact that there can be blameless instances of incredulous collaboration does not entail that there should not be a belief norm, just as the desire to excuse the speeding driver does not entail that there should be no speed limits.

A fifth response is that some amount of unexcused rule breaking is optimal for norm-governed systems. Take, for example the question of whether we should accept a factive norm for everyday assertion. It might be thought that accepting this standard will increase the amount of objectively true statements uttered and increase the true beliefs that individual community members are likely to have. This is a good outcome. But, what is good for the individual is not necessarily good for communities (Mayo-Wilson, Zollman, & Danks 2011). For instance, too much endorsement of a rigid standard might disincentivize individuals from communicating entirely, which would negatively impact communicative networks and ultimately decrease benefits for the typical community member. Thus, there is likely to be some sort of trade-off between the benefits of a high objectivity rate as a result of the factive rule and the negatives that come from complete adherence to that rule for the community. The same might be said for belief and publishing.

To better understand this trade-off in ordinary assertion, researchers ran several computational agent-based computer simulations to model the gains and losses of different objectivity rates in simple question-and-answer exchanges (Turri 2023). Based on these simulations, researchers were able to determine the point at which too much adherence to the truth norm became counterproductive. In doing so, researchers identified a low but non-trivial proportion of non-objective communication that was optimal for ordinary communicative systems. This suggests that certain communication networks are optimal when they include a factive norm, high adherence rates to that norm, and a small number of holdouts.

Although comparison between domains is fraught, this research may begin to shed light on the community benefits of the norm in scholarly contexts. It is reasonable to expect that there will be a similar sort of trade-off in scholarly publishing. And although the belief norm is significantly less demanding than a truth norm, it is possible that there will also be a point at which the community loses more than it gains from all adhering to it. An exciting opportunity for future research would be to conduct similar computer simulations to determine this point and if it is also reached after high rates of adherence to the norm. However, the fact that some incredulous collaborations lead to progress is not a reason to reject the belief norm, this response hypothesizes, because some incredulous collaboration could be among the low but non-trivial proportion of rule breaking that is optimal for publishing systems.

A final response is that it is possible that the arguments above are wrong, and that the belief norm should not apply to certain collaborative publications. In this case, the belief norm should apply to single authored publications, and become slowly outweighed as more coauthors are added. This would be a significant restriction of the thesis, since a central tendency of many scientific fields is collaboration. In such an eventuality though, this result would still tell us something important about the norms governing optimal inquiry, for three reasons. First, thousands of single-authored papers are published each year, and currently comprise somewhere between 70% to 83% of papers in philosophy alone (Bourget & Weinberg 2021). Second, the fact that belief is and should be the norm until replaced by other norms would increase our understanding of both optimal scholarly activity and advance a more sophisticated discussion of conflicts between norms. Third, analyses suggest that collaborative publishing rates in philosophy have rapidly increased in a relatively short amount of time, which hastens the need to understand the shift between norms as the field evolves. If belief is no longer the mental state that authors should have toward their conclusion once they start working together, then what is?

2.3. *The Norm Makes Us Jerks*

One potential challenge to the belief norm is that it perpetuates scholarly disagreements that lead to philosophical scepticism about the aims of philosophical activity (Beebe 2018). If belief does perpetuate disagreements of this sort, it is interesting to speculate about the social effects that this may have on research communities. For example, disagreement might have negative social consequences for research communities by making research overly personal. Doing so might blur the lines between the questions of whether, say, non-naturalism is silly and whether an author's view about non-naturalism is silly (and for believing it, perhaps that

author too). When things become overly personal, this can encourage dogmatic adherence to claims, unprofessional behavior, and ad hominem attack. Given the climate of some research communities, the worry that the norm contributes to unprofessional behavior or dogmatism is an important one to address.

Depersonalizing research is an important goal that may improve inquiry. However, it's also unclear that rejecting the belief norm will prevent these negative consequences. It is likely that some researchers will judge authors or dogmatically defend claims so long as there are authors of papers to judge. Alternatively, in well-functioning scholarly communities, personal connections can sometimes increase collegiality. For example, they might remind us that we work in a community rather than a vacuum, that we are accountable for what we say, or are required to respond professionally when others disagree with us. This can lead us to be more respectful and to consider the arguments and evidence of others more seriously than we might otherwise if it was published without belief.

2.4. Researchers Aim to Advocate, Not Just Educate

Prior researchers have argued that the belief norm should apply to some published research claims and not others, depending on the intentions of the authors who publish them and the roles those claims end up playing in the resulting research community (Fleisher 2020). According to this theory, "Evidential Role Claims" are those that are published when an author intends to "increase the common stock of evidence available to inquirers" or by adding "what is mutually accepted by inquirers" to the research record (2020: 243). For example, this might describe many of the claims published in a literature review or encyclopedia article. By contrast, "Advocacy Role Claims" are claims that function as "entries into debate" or ones that an author uses to "promote productive debate and disagreement" (2020: 242). When advocating for a claim, authors provide new arguments for views and position themselves as defenders of them for the purposes of participating in scholarly exchanges. Researchers have argued that the belief norm of academic publishing should apply to evidentiary role claims but not to advocacy role claims.

Why shouldn't the belief norm apply to advocacy role claims? Or put another way, why should we think that the debate and disagreement following advocacy is more productive without belief? Presumably, the idea is that publishing the claim that something is true that one does not believe can still have good effects for scholarly research that would be lost if publication was limited by belief. It is possible to imagine that such publications could sometimes still provoke productive discussion, promote valuable discourse, and advance considerable arguments that motivate new areas of scholarship. Therefore, it might be thought, the norm should be rejected because it would deprive the research community

of these positive effects. But it is unclear whether these effects would not be better with belief. To evaluate this question, we must also consider the negatives of advocacy without belief on the publication system.

While advocating for claims one does not think are true can lead to good outcomes, this also has downsides for the publication system. These downsides make advocacy susceptible to considerable bias and abuse. After all, claims that elicit disagreements, draw attention to findings, or constitute entries to debates form a very broad category. Not all disagreements are good and there are many unproductive outcomes that eliciting disagreements between researchers can and often does have in academia. The power to influence the direction of research discussions within a field might be motivated by good intentions, but it may also be motivated by many personal or self-serving interests beyond discovering truths or promoting productive inquiry. For example, disagreements are often associated with and perpetuated by credit seeking, career advancement, reputation management, and scholarly competition. Given the risks that advocacy is often motivated by these things, one might be tempted to think that advocacy role claims stand to benefit even more from the belief norm than evidential role claims do. That is, when one wishes to start debates and influence others by publishing research articles, the quality of that debate is likely to be better under a norm that grounds publication in belief and the search for truth. Similarly, when one launches a new area of research, this is more likely to lead to productive scholarly activity in the long run when sustained by a norm with a constitutive connection to truth (for extended discussion of several ways this might improve research in philosophy, see Sarihan 2022). If this is true, then it suggests that advocacy role claims will fulfill their function of promoting productive debate and disagreement even better under a belief norm.

2.5. Researchers Change Their Minds

As inquiry proceeds, researchers sometimes reject claims made in prior publications. And updating beliefs when considering new evidence can often be a sign of healthy inquiry. Does the norm say that authors should always continue to believe what they publish? In response, the norm is perhaps best interpreted as holding that authors should believe claims at some point during the advanced publication process. As a corollary to this interpretation, the scholarly community should recognize that people are more than authors of journal articles and that the norm applies to what authors publish when they publish it. Distinguishing between person and author in this way further reduces the risk of overpersonalization, directs our focus towards the content of distinct contributions, and encourages more intellectually humble claims lest we ourselves change

our minds. Incidentally, this recognition is reflected in the way that academic research articles are primarily cited and organized, according to author and date.

2.6. Scientists Rely on False Approximations That Are Not Believed

Given inherent limitations in human faculties, technologies, and instruments, scientists must often use approximation to represent values, sizes, distances, weights, volumes, and temperatures in the papers they publish. Scientists also publish idealized models and diagrams to help explain phenomenon and underlying mechanisms to their audience, even though they may include assumptions that are not actually true or that do not represent anything in the actual world. In these senses, researchers regularly publish things they know are false, but which may still have immense practical and pedagogical value. In this way, publishing strictly false propositions improves scientific understanding.

There are large and sophisticated literatures regarding measurement, approximation, idealization, and “useful fictions” in science (Elgin 2004; Potochnik 2017; Trout 2003). But there are at least two initial responses that can be made to this objection in the present context. First, researchers do not believe that rounded values are actual values or that idealized models are actual, but nor do they typically publish claims that they are. It is usually clear that this is not what is being asserted in published research articles, either through the explicit communication of model assumptions or implicit to the field specific conventions governing rounding and reporting. If this is true, then authors do not violate the belief norm when publishing approximated values or claims that utilize idealized models. Arguably, the degree to which idealized models effectively improve scientific understanding depends on clarifying that they are not actual. Second, and perhaps more interestingly, this observation seems to capture something important about the epistemic practices of science, more broadly. For instance, it is often speculated that the sort of propositions believed by scientists ‘is not one expressed by a formula “F,” but one with the content “F, to some approximation”’ (Goldman 1999: 246). It could be that approximation and idealization are instances of a more general tendency to believe approximate truths (for discussion, see Buckwalter & Turri 2020). From this observation, a further possibility worthy of future research is that the belief norm of publishing can be formulated in terms of beliefs that are approximately true.

2.7. Belief is Too Weak

Some philosophers have argued that ordinary notions of belief are weak (Hawthorne, Rothschild, & Spectre 2016). According to these arguments, believing something

does not require too much confidence in it. In many cases, believing that something is true may only require thinking that something is likely to be true. As speculated above, this may often be the attitude that scientists have toward discoveries. Researchers have further argued that the standards for everyday assertion are much higher than this, which suggests that proper assertion requires more than just belief. If belief only requires thinking something is likely to be true or approximately true, as it sometimes seems to in science, this may also suggest that proper assertion in published journal articles should require much more than just belief.

In response, it is likely that belief is just one norm among many stronger norms that sustain the practice of academic publishing. When it comes to reaching the truth or curating the research record, for instance, it might be suggested that knowledge is an even better candidate than belief. Some have even argued that knowledge is the “the prime pedagogical principle” for “transmitting information or skills” during instruction and demonstration (Buckwalter & Turri 2014). Additionally, a factive norm like knowledge would also be less susceptible to many of the objections raised above. For example, it would not be susceptible to the objection that researchers could get ahead by having low standards or that they sometimes reject prior claims they publish. What is true doesn’t change because people are lazy or ambitious or change their minds. These considerations point to the merits of adopting a considerably stronger norm, namely knowledge. Future research might profitably explore the benefits of strengthening the belief norm by adopting a knowledge norm of publishing.

3. Conclusion

This paper examines the strengths and weaknesses of the belief norm for scholarly publication in academic journals. Several arguments are presented in favor of the belief norm, including the claims that it increases the veracity of the research record, modulates the strength of published claims, motivates evidence responsiveness, encourages good faith engagement necessary for healthy research communities and public trust, helps organize effective collaborations, coheres with broader communicative norms and the standards for assertion more generally, and is associated with recent methodological reforms in science. Considering these strengths, the paper concludes that the published research record has systematically benefited from the belief norm and that it should continue to be counted among those that sustain the practice.

The present research has several noteworthy limitations that also suggest avenues for future research. The first regards the scope of the norm and the force of the argument. It has been argued that the norm should apply to certain claims that researchers make, such as core research contributions to the research record.

It has also been argued that the norm should apply to these claims within certain contexts, such as when they appear in published academic journal articles. Future research might characterize the types of claims regulated by the norm in greater detail and explore other types of scholarly contributions in which adopting doxastic norms could be beneficial. Second, the argument given is that the norm should apply to publications reporting truth conducive inquiry. While it is typically assumed that the aim of research across many fields is to conduct inquiry, it might not be the only goal that researchers have when they publish academic articles. Thus, future research might profitably explore the extent to which researchers in different fields publish with that aim in mind and whether the norm should continue to apply as other motivations take precedence. Lastly, a final area for improvement would be to collect more evidence for the norm through computational modeling of scholarly communication networks. Doing so will further improve our understanding of the advantages and disadvantages of the belief norm of academic publishing on research and academic communities.

Acknowledgements

For helpful feedback, discussion, and comments on previous drafts, I thank Carolyn Buckwalter, Alexandra Plakias, Joel Smith, John Turri, an editor, and anonymous reviewers at *Ergo* and *Analysis*.

References

- Aad, G., B. Abbott, J. Abdallah, O. Abidinov, R. Aben, M. Abolins, O. S. Abouzeid . . . ATLAS Collaboration (2015). Combined Measurement of the Higgs Boson Mass in pp Collisions at $\sqrt{s}=7$ and 8 TeV with the ATLAS and CMS Experiments. *Physical Review Letters*, 114(19), 191803. <https://doi.org/10.1103/PhysRevLett.114.191803>
- Alston, William P. (1989). *Epistemic Justification : Essays in the Theory of Knowledge*. Cornell University Press.
- Ananthaswamy, Anil (2020, October 13). Do We Live in a Simulation? Chances Are about 50–50. *Scientific American*. Retrieved from <https://www.scientificamerican.com/article/do-we-live-in-a-simulation-chances-are-about-50-50/>
- Barnett, Zach (2019). Philosophy without Belief. *Mind*, 128(509), 109–38.
- Beebe, Helen (2018). Philosophical Scepticism and the Aims of Philosophy. *Proceedings of the Aristotelian Society*, 118(1), 1–24. <https://doi.org/10.1093/arisoc/aox017>
- Bennett, Jonathan (1990). Why Is Belief Involuntary? *Analysis*, 50(2), 87–107.
- Bourget, David and Justin Weinberg (2021, October 14). Co-Authorship in Philosophy over the Past 120 Years. *Daily Nous*. Retrieved from <https://dailynous.com/2021/10/14/co-authorship-in-philosophy-over-the-past-120-years-bourget-weinberg/>

- Bragg, William H. (1907). On the Properties and Natures of Various Electric Radiations. *The London, Edinburgh, and Dublin Philosophical Magazine*, 14(82), 429–49.
- Bragg, William H. (1908). The Nature of γ and X-Rays. *Nature*, 77, 270–71.
- Bright, Liam Kofi, Haixin Dang, and Remco Heesen (2018). A Role for Judgment Aggregation in Coauthoring Scientific Papers. *Erkenntnis*, 83(2), 231–52. <https://doi.org/10.1007/s10670-017-9887-1>
- Buckwalter, Wesley and John Turri (2014). Telling, Showing and Knowing: A Unified Theory of Pedagogical Norms. *Analysis*, 74(1), 16–20.
- Buckwalter, Wesley and John Turri (2020). Knowledge and Truth: A Skeptical Challenge. 101(1), 93–101. <https://doi.org/10.1111/papq.12298>
- Clark, Cory and Philip Tetlock (in press). Adversarial Collaboration: The Next Science Reform. In C. L. Frisby, R. E. Redding, W. T. O'Donohue, and S. O. Lilienfeld (Eds.), *Political Bias in Psychology: Nature, Scope, and Solutions*. Springer.
- Dang, Haixin and Liam Kofi Bright (2021). Scientific Conclusions Need Not Be Accurate, Justified, or Believed by Their Authors. *Synthese*, 199, 8187–203. <https://doi.org/10.1007/s11229-021-03158-9>
- Dethier, Corey (2022). Science, Assertion, and the Common Ground. *Synthese*, 200(1), 1–19. <https://doi.org/10.1007/s11229-022-03580-7>
- Eklund, Anders, Thomas E. Nichols, and Hans Knutsson (2016). Cluster Failure: Why fMRI Inferences for Spatial Extent Have Inflated False-Positive Rates. *Proceedings of the National Academy of Sciences*, 113(28), 7900–905. <https://doi.org/10.1073/pnas.1602413113>
- Elgin, Catherine Z. (2004). True Enough. *Philosophical Issues*, 14(1), 113–31.
- Engel, Pascal (2004). Truth and the Aim of Belief. In D. Gillies (Ed.), *Laws and Models in Science* (77–97). King's College Publications.
- Fagan, Melinda Bonnie (2011). Is There Collective Scientific Knowledge? Arguments from Explanation. *The Philosophical Quarterly*, 61(243), 247–69. <https://doi.org/10.1111/j.1467-9213.2010.676.x>
- Fleisher, Will (2020). Publishing Without (Some) Belief. *Thought: A Journal of Philosophy*, 9(4), 237–46. <https://doi.org/10.1002/tht3.466>
- Gerken, Mikkel and Esben Nedenskov Petersen (2020). Epistemic Norms of Assertion and Action. In Sanford Goldberg (Ed.), *The Oxford Handbook of Assertion* (683–706). Oxford University Press.
- Gilbert, Margaret (1987). Modelling Collective Belief. *Synthese*, 73(1), 185–204.
- Gilbert, Margaret (1992). *On Social Facts*. Princeton University Press.
- Goldman, Alvin I. (1999). *Knowledge in a Social World*. Clarendon Press.
- Hawthorne, John, Daniel Rothschild, and Levi Spectre (2016). Belief Is Weak. *Philosophical Studies*, 173(5), 1393–1404. <https://doi.org/10.1007/s11098-015-0553-7>
- Hobbes, Thomas (2006). Third Set of Objections, by a Famous English philosopher. In R. Ariew and D. Cress (Eds.), *Meditations, Objections, and Replies* (100–114). Hackett. (Original work published 1641)
- Mayo-Wilson, Conor, Kevin J. S. Zollman, and David Danks (2011). The Independence Thesis: When Individual and Social Epistemology Diverge. *Philosophy of Science*, 78(4), 653–77. doi:10.1086/661777
- Mellers, Barbara, Ralph Hertwig, and Daniel Kahneman (2001). Do Frequency Representations Eliminate Conjunction Effects? An Exercise in Adversarial Collaboration. 12(4), 269–75. <https://doi.org/10.1111/1467-9280.00350>
- Millson, Jared A. (2020). Seeking Confirmation: A Puzzle for Norms of Inquiry. *Analysis*, 80(4), 683–93. <https://doi.org/10.1093/analys/anaa017>

- Mouk, Yascha (2018). What an Audacious Hoax Reveals about Academia. *The Atlantic*. Retrieved from <https://www.theatlantic.com/ideas/archive/2018/10/new-sokal-hoax/572212/>
- Palmira, Michele (2020). Inquiry and the Doxastic Attitudes. *Synthese*, 197(11), 4947–73.
- Plakias, Alexandra (2019). Publishing without Belief. *Analysis*, 79(4), 638–46. <https://doi.org/10.1093/analysis/anz040>
- Potochnik, Angela (2017). *Idealization and the Aims of Science*. University of Chicago Press.
- Ramsey, F. P. (1931). *Foundations of Mathematics and Other Logical Essays*. Routledge.
- Resnick, Brian (2019, January 4). Intellectual Humility: The Importance of Knowing You Might Be Wrong. *Vox*. Retrieved from <https://www.vox.com/science-and-health/2019/1/4/17989224/intellectual-humility-explained-psychology-replication>
- Rohrer, Julia M., Warren Tierney, Eric L. Uhlmann, Lisa M. DeBruine, Tom Heyman, Benedict Jones, . . . Tal Yarkoni (2021). Putting the Self in Self-Correction: Findings from the Loss-of-Confidence Project. *Perspectives on Psychological Science*, 16(6), 1255–69. <https://doi.org/10.1177/1745691620964106>
- Sarıhan, Işık (2022). Problems with Publishing Philosophical Claims We Don't Believe. *Episteme*. Advance online publication. <https://doi.org/10.1017/epi.2021.56>
- Shah, Nishi and J. David Velleman (2005). Doxastic Deliberation. *Philosophical Review*, 114(4), 497–534.
- Tetlock, Philip E. and Gregory Mitchell (2009). Implicit Bias and Accountability Systems: What Must Organizations Do to Prevent Discrimination? *Research in Organizational Behavior*, 29, 3–38. <https://doi.org/10.1016/j.riob.2009.10.002>
- Trout, J. D. (2003). *Measuring the Intentional World: Realism, Naturalism, and Quantitative Methods in the Behavioral Sciences*. Oxford University Press.
- Turri, John (2016). *Knowledge and the Norm of Assertion: An Essay in Philosophical Science*. Open Book Publishers.
- Turri, John (2017). Experimental Work on the Norms of Assertion. *Philosophy Compass*, 12(7), e12425. <https://doi.org/10.1111/phc3.12425>
- Turri, John (2023). Simulating Norms of Assertion: Loss, Benefit, Inequality, and Objectivity. Unpublished manuscript.
- Turri, John and Peter Blouw (2015). Excuse Validation: A Study in Rule-Breaking. *Philosophical Studies*, 172(3), 615–34.
- Turri, John, David Rose, and Wesley Buckwalter (2017). Choosing and Refusing: Doxastic Voluntarism and Folk Psychology. *Philosophical Studies*, 175(10), 2507–37. <https://doi.org/10.1007/s11098-017-0970-x>
- Velleman, David (2000). On the Aim of Belief. In David J. Velleman (Ed.), *The Possibility of Practical Reason* (244–81). Oxford University Press.
- Wheaton, Bruce R. (1981). Impulse X-Rays and Radiant Intensity: The Double Edge of Analogy. *Historical Studies in the Physical Sciences*, 11(2), 367–90.
- Williams, Bernard (1973). Deciding to Believe *Problems of the Self* (136–51). Cambridge University Press.
- Williamson, Timothy (2000). *Knowledge and Its Limits*. Oxford University Press.
- Wray, K. Brad (2017). The Impact of Collaboration on the Epistemic Cultures of Science. In Thomas Boyer-Kassem, Conor Mayo-Wilson, and Michael Weisberg (Ed.), *Scientific Collaboration and Collective Knowledge: New Essays* (117–34). Oxford University Press.