

* Forthcoming in *Analysis*

Searching for Epistemic Norms that Matter

Epistemologists are engaged, among other things, in the business of formulating epistemic norms. That is, they formulate principles that tell us what we should believe and to what degree of confidence, or how to evaluate such epistemic states. In *The End of Epistemology As We Know It*, Brian Talbot argues that thus far, most of the theories resulting from these efforts are flawed, because:

The epistemic norms should matter. The ones philosophers typically focus on do not matter enough. So we should replace them... [T]he replacement norms... will permit us to form ... beliefs that violate all standard norms by going against our evidence, being incoherent, or even being clearly false. (p. 1)

According to Talbot, the sin that most normative theories commit is that they universally disallow and devalue believing what is clearly contrary to one's evidence, or incoherent, or clearly false. Talbot suggests that once we ask ourselves when and why our beliefs matter, we should conclude that sometimes, there is nothing wrong with such beliefs, and that they can even be the right beliefs to have. According to Talbot, standard epistemic norms are on the wrong track. The title 'the end of epistemology as we know it' is a bit of an exaggeration, since standard epistemic norms don't comprise the whole of epistemology. However, if the arguments succeed, some widely accepted theories should be discarded, so the stakes are high.

Talbot has a clear and easy style of presentation, and in the course of his argument, walks the reader through a large swath of recent literature in epistemology and ethics. I learned a lot from reading and thinking through the book, you probably will too. The structure of the book's eight chapters is as follows. Chapter 1 argues that epistemic norms need vindication, and it defines concepts used in the chapters that follow. Chapters 2, 4, 5, 6, and 7 discuss various approaches to vindicating epistemic norms and argue that these approaches are either implausible or else they vindicate non-standard norms. Chapter 3 discusses some objections to the project, and Chapter 8 speculates as to where all this leads.

My initial reaction to the book was that it can't be right. One of his primary examples of a theory that posits norms that supposedly don't matter is Bayesianism. But Bayesian norms do seem to matter. We use Bayesian principles to guide our beliefs whenever we rely on statistical reasoning: to test medical treatments, to predict the weather and to develop artificial intelligence. They have proven so useful that it is difficult to deny that they matter, and yet Talbot follows his arguments to this conclusion. So where do those arguments go wrong? I cannot do justice to all of them in this short notice, so I will comment on only three.

1 Does it matter that the norms often don't matter?

The rules of chess don't always matter so much. As they are standardly taught, for instance, the rules say that a pawn can *never* move sideways. Taken simplistically, this would mean that you shouldn't move a pawn sideways even when you are not playing chess. However, when not playing chess, it doesn't matter at all how you move a pawn. Most of us don't spend most of our time playing chess, so most of the time, the rules of chess don't matter to us. Even when we *are* playing chess, the rules don't always matter to the same degree because some games are more important than others. If the winner will receive a significant prize, for example, the game and its rules matter more than they do for players who are playing merely to relieve boredom. The fact that the rules of chess don't always matter—and that when they matter, they don't always matter to the same degree—is not a reason to replace the rules of chess, nor does it seem especially perplexing.

Let's move to epistemology. One of Talbot's first arguments starts from the claim that not all epistemic states matter to the same degree. It is much more important, for example, to have correct beliefs about the moral value of human beings than to have correct beliefs about the precise number of sugar grains in a particular jar. The former is an example of what Talbot calls *interesting beliefs*, beliefs that matter in their own right; the latter is an example of *pointless beliefs*, beliefs that don't matter. But the epistemic norms, as they are normally formulated, don't make any such differentiation, and contrary to what I suggested in relation to chess, Talbot argues that this fact requires attention. Why? Here's one of his lines of argument:

What if we say that the epistemic norms don't matter to any significant degree when they govern pointless beliefs? This would mean that we need an explanation for why the

epistemic norms matter when they govern non-pointless beliefs. Why? The overwhelming majority of things we can possibly believe are pointless. (Sec. 1.4.1)

The term “need an explanation” can mean more than one thing (Baras, 2022, sec. 1.3). I take it that here Talbot means that we should expect them to have an explanation, rather than being brute facts, and probably also that it is a worthwhile project for theoreticians to seek such an explanation.

My complaint about this argument is not about its soundness—I accept each of its premises and its conclusion. Rather, I think the argument is superfluous. There is a simple explanation for why the norms matter when applied to non-pointless beliefs, whereas they don’t matter when applied to pointless beliefs. The explanation is, simply, that the pointless beliefs are pointless. It’s just like the rules of chess not mattering when we’re not playing chess: True, the vast majority of possible chess-piece moves can be made while we aren’t playing chess, but that fact doesn’t create any special need to explain why the rules matter when we *do* play chess. If we’d like to be nitpicky, we can add to the rules of chess that they only apply when playing chess. Similarly, if we want to formulate epistemic norms that track what matters, we can say that the norms that are aimed at accuracy only matter when and to the extent that the accuracy of our beliefs matters. But it seems to me that we get along fine without making that domain restriction explicit.

Does that mean that theoreticians shouldn’t spend their time seeking an explanation for why epistemic norms matter when they do? Not necessarily. This question—what makes the truth or accuracy of some beliefs important in its own right, and how various norms can promote truth or accuracy—is interesting. So I agree with Talbot’s conclusion: there is interesting work here for theoreticians. I just don’t think the fact that epistemic norms don’t matter when applied to pointless beliefs adds much to the significance of this project.

2 How much does maximizing matter?

I move now to Chapter 2, which discusses versions of epistemic consequentialism. According to epistemic consequentialism, epistemic norms are vindicated by considering their consequences. A primary example is justifying Bayesian norms by proving that following them will increase the accuracy of one’s credences (Pettigrew, 2016). Such arguments rely on an assumption that we

should try to make our credences as accurate as possible. That is, they rely on an assumption that accuracy should be *maximized*. Talbot argues against this assumption and thus undermines such arguments for Bayesian norms. Mirroring a similar debate among consequentialists in ethics, Talbot argues that if some version of epistemic consequentialism is correct, it is not a maximizing version. Instead, we should aim at *satisficing* norms—norms that permit credences that are not maximally accurate, so long as they are close enough (up to some distance ϵ from the maximal option).

Maximizing norms are, by definition, norms that yield the best consequences. How could any norms, satisficing or otherwise, beat that? The answer is, of course, that sometimes there is a price to maximizing. For instance, applying maximizing norms may come at a cost to our cognitive resources, and it might in some cases make sense to use less costly heuristics instead. Talbot is explicit (p. 56), however, that this is not his argument. Rather, he argues that even if there were no cost associated with maximizing, the norms we should follow should still be laxer than maximizing norms. Why? One of his arguments on this topic appears in the following passage:

The consequentialist explains how much some norm matters by comparing the consequences of conformity with that norm to the consequences of conforming to another norm instead. The norms vindicated by maximizing and satisficing consequentialism differ in cases where there are suboptimal options that are just a bit worse than the best option. In such cases, there will either be a much worse option as well, or there won't be. If there is a much worse option, then norms vindicated by maximizing consequentialism matter much less than norms vindicated by satisficing. That's because, to explain how much a norm matters, we look at the benefits of conforming to it rather than to another, and conforming with maximizing rather than satisficing norms matters much less than conforming with satisficing norms rather than even more lax norms. (Section 2.3.1)

The argument relies on ranking norms not simply by their consequences but rather by the distance between their consequences and the consequences of some other set of norms. How much it matters whether we use maximizing norms is measured by the difference between following maximizing norms and following satisficing norms. By definition, satisficing norms

are not so far apart from maximizing, so the result is that maximizing has an insignificant advantage. However, when measuring how much satisficing norms matter, Talbot compares them not to maximizing but rather to some unspecified norm that, he assumes, yields significantly worse results. So, in comparison to the much laxer norms, satisficing norms matter a lot.

The argument, as presented in the quoted passage, might be summarized as follows:

- (1) A norm matters to the extent that it yields better consequences than another norm.
- (2) In comparison to a satisficing norm, a maximizing norm makes an insignificant improvement.
- (3) In comparison to an even laxer norm, a satisficing norm makes a significant improvement.

Conclusion: Satisficing norms matter more than maximizing norms.

However, this is not a valid argument: The problem is that its first premise, which states a method for ranking norms, does not specify which norm to compare to. The choice (encompassed in premises 2–3) of which comparisons to consider—comparing maximizing to satisficing and satisficing to laxer norms—is unmotivated. As will soon be explained, other comparisons yield different results, and therefore we could just as easily derive the opposing conclusion that maximizing norms matter more.

True, in comparison to norms that recommend completely inaccurate credence (i.e., accuracy-minimizing-norms), satisficing norms make a big difference. However, when compared to the same accuracy-minimizing-norms, maximizing norms make at least as much a difference, and often slightly more. Similarly, when compared to satisficing norms, maximizing norms make only a small difference for the better. But the same is true in the other direction: compared to maximizing norms, satisficing norms make only a small difference for the worse. And these are only three possible sets of norms—there are numerous, perhaps infinite, possible sets of norms one can follow and possible credal states one can be in. For example, we might compare satisficing norms to norms that are only slightly laxer than satisficing norms. In comparison to those norms, it only matters a bit that we follow satisficing norms. When comparing maximizing

norms to these slightly-laxer-than-satisficing norms, it matters a lot more that we opt for maximizing norms.

As illustrated here, depending on which norms one compares to, any other norm can appear to be as significant or insignificant as one wishes. The method for ranking norms assumed in Talbot's argument (premise 1) thus provides no ranking at all and should be rejected. And the comparisons between norms that the argument relies on (premises 2–3) do not show there to be any advantage to choosing satisficing norms over maximizing norms.

3 The argument from tradeoffs

One of the main arguments of the book comes from more the familiar territory of considering potential tradeoffs. An example of a tradeoff would be a case in which believing something that goes against one's evidence, or is incoherent, or is clearly false, would allow one's beliefs to be better overall. For instance, a good tradeoff might sacrifice a belief that doesn't matter so much in order to gain more accurate beliefs that do matter. In such a case, Talbot thinks that we should accept the tradeoff. He reasons that because standard norms tell us never to maintain such beliefs, standard norms can't be the ones that matter. His argument can be summarized as follows:

- (1) Standard epistemic norms imply that all tradeoffs are impermissible.
- (2) Some tradeoffs are permissible.
- (3) If (1) and (2), then standard epistemic norms don't matter.

Conclusion: Standard epistemic norms don't matter.

A similar argument can be applied to attempted vindications of epistemic norms. Suppose the vindication is premised on some conjunction of proposition P. If P forbids tradeoffs, then P (the theory that is supposed to provide vindication) should be rejected, and therefore the vindication fails. If, on the other hand, P permits tradeoffs, then P cannot vindicate standard norms that forbid tradeoffs. Either way, we'll conclude that all vindications of standard norms fail. (Something like this is presented in section 7.2 as the Master Challenge.)

Before discussing some ways of resisting the argument, I'd like to point out one way of accepting its conclusion while nevertheless preserving much of epistemology as we know it. In one way or another, a tradeoff involves believing contrary to what is supported by the evidence. Thus, the concept of a tradeoff presupposes a concept of evidential support. Perhaps, in turn, evidential support belongs to a broader class of epistemic support, which also includes a priori forms of support. We may want to interpret some of the principles put forward by epistemologists as explications of epistemic support rather than as norms of belief. Indeed, some theorists are explicitly interested in the former rather than the latter. Theories of evidential support are outside the scope of Talbot's arguments, which focus on what we should believe. Perhaps sometimes we shouldn't believe what our evidence supports—but often we should, and that's enough to make theories of evidential support matter.

3.1 Do we actually face tradeoffs?

My initial response to this argument was to consider whether I ever felt like I should be making a tradeoff in my beliefs, and I was then and am still unable to think of any case in which I thought I should have a clearly false belief in order to obtain some good. If my experience mirrors that of most people, and we never *actually* face circumstances in which we should make tradeoffs, then perhaps it doesn't matter that standard norms imply that tradeoffs are impermissible. So long as they provide good guidance in all of the actual cases we confront, then they clearly matter for that reason. (In this way, I'm rejecting premise 3.)

That said, things may look different when thinking about how epistemic norms might guide us in assessing and reacting to other people's beliefs. In some ways, I am less inclined to judge beliefs that I see as clearly false if they don't seem to matter, or if they seem beneficial overall. So perhaps there's something wrong about standard norms if we think of them as norms for assessing others. Even if that is correct, if standard norms provide good first-personal guidance in actual cases, that's enough for them to matter.

But maybe I'm wrong—*do* we actually face tradeoffs in the actual world? With this question in mind, I returned to the book looking for examples. Talbot's initial examples of tradeoffs appear in the opening chapter (section 1.4.5). Two of them, Talbot acknowledges, are unrealistic. They are not cases of a kind that readers of this essay are likely to encounter—one is about a

researcher who will only receive funding if she sincerely believes in God on insufficient evidence, and the other is a researcher who will only receive funding from a cookie company if she sincerely believes something about the cookies that is false and contrary to her evidence. The third example is of a self-fulfilling belief: Someone is sick, and they will get better if they believe that they will. This last example might be realistic, but, as Talbot acknowledges, it is not truly a trade-off example because, if a belief is evidently self-fulfilling, then it is not contrary to the evidence to maintain that belief.

Only one of his examples is a candidate for being both realistic and a genuine case of a tradeoff:

Overconfidence: Researchers who are overconfident about their abilities are more likely to make discoveries.

Talbot continues to suggest that in such circumstances, researchers are justified in being overconfident. Shortly, I will question Talbot's verdict in this case. Beforehand, I want to strengthen his argument by noting that this example points to a whole cluster of examples that could potentially strengthen his argument and that are worth considering.

Believing that one has greater abilities than the evidence suggests is but one example of a well-documented optimism bias (Sharot, 2012).¹ This optimism bias is but one example of a long list of well-documented human biases (Kahneman, 2011), some of which are the result of subconscious heuristics that in many ordinary circumstances lead to true beliefs but sometimes lead to false ones. One example is our use of stereotypes to form beliefs about people around us; another is our belief in falsehoods that lead to beneficial actions. We might think our children are better than other people's children or that our group (religion, nation, ethnicity) is better than other groups, and we might be overconfident in our abilities and beliefs. These versions of the optimism bias have positive effects—they might increase our motivation and ability to care for our children, to cooperate with our social groups, and to overcome worries about risk when engaging in valuable projects. Some of the benefits of these false beliefs can be epistemic, like allowing us to conduct risky research that can end up generating valuable new knowledge. In all

¹ In section 5.2, Talbot provides some more examples for how the optimism bias can be beneficial. If it were possible to revise the book, I would suggest incorporating this section in chapter 1. Since tradeoffs play such a crucial role in Talbot's argument, the book should start with the best attempt at providing realistic examples and not save them for an arbitrary location in the middle.

such cases, we can ask—once we know that we have these biases and that they are generally beneficial—should we embrace them or try to resist them? According to Talbot’s view, at least sometimes the resulting false beliefs should be embraced.

There are other contexts in which some people think that it’s worth having false beliefs because such beliefs can have various benefits. Consider religious fictionalism—the view that certain religious claims are false but that we should, in some ways, live our lives as if they were true (Scott & Malcolm, 2018). Religion is not the only domain in which some have opted for fictionalism—moral fictionalism exists as well (Joyce, 2001, Chapter 8). Not every form of fictionalism implies that one should have false beliefs. A fictionalist might be content with false discourse and pretense, without belief. In fact, it would probably be incoherent for anybody to be a fictionalist and at the same time really believe the fiction—that is, if you really believe the fiction, you can’t genuinely believe that it is just a fiction. However, some fictionalists go as far as justifying false beliefs, and their views would fall on Talbot’s side rather than on side of the standard norms that he criticizes.

Contrary to such views, I do not feel compelled to embrace clearly false beliefs in any of these examples. First, it should be clear—and Talbot should agree—that we have no reason to embrace false beliefs that result from a heuristic. Because a heuristic is a mechanism that saves us cognitive resources and generally serves us well enough, we have reason to embrace the *mechanism*—but we have no reason to embrace individual beliefs produced by such mechanisms, once we realize that they are probably false. That is because in the case of heuristics, there are no benefits to believing falsely; rather, the false beliefs are merely byproducts of *mechanisms* that are generally beneficial.

Talbot’s argument is more compelling when a false belief can be beneficial, such as with the optimism bias or, as some argue, a religious fiction. However, I do not feel compelled to embrace clearly false beliefs even in such cases. I tend to aspire to be truthful even with regard to my own abilities and prospects, regardless of possible benefits of false beliefs, and I feel no attraction to believing false religious or other fictions. Is this irrational of me?

The argument for embracing our false beliefs, in cases of optimism bias, would be that they will have beneficial practical implications. Generally, the claim is that without such false beliefs, we

would engage less in risky projects because we would worry about failure. This claim presupposes that engaging in such projects is overall expected to be beneficial. But if they are beneficial overall, then why would their riskiness deter us from engaging in them? Probably because we are irrational in our assessment of risk or in our decision-making when faced with risk. Thus, the claim must be that the optimism bias, when beneficial, uses irrationality to counter irrationality.

Consider the earlier example of overconfidence in a researcher. A feature of a significant discovery is that it diverges significantly from what was previously believed. Because of this divergence, such discoveries are typically unpredictable. To make unpredictable discoveries, researchers will often have to engage in unpredictable projects. More importantly, they'll have to engage in projects to which they should initially attribute low probability of success. (Or, at least, all of this seems plausible from the armchair.) Once we realize this, perhaps the rational thing to do is engage in projects with a low probability of success. The reason we need the optimism bias to help us do so, so the reasoning must go, is that we have trouble engaging in projects with low probability of success. Perhaps it's a fact about how our motivational system works that we feel unmotivated to engage in projects that we think will most probably fail. Thus, the optimism bias causes us to believe that, contrary to the evidence, the project is likely to succeed. And this makes it more likely that we will find the mental strength and will to engage in the project.

In light of all this, what should I think about my optimism bias? Is embracing the bias the only psychologically available route to counter irrationality with regard to risky decisions? I tend to think that we can and should try to overcome both irrationalities. We should try to be more modest about ourselves and be more willing to take risks. Perhaps, though, my belief that we can do so is naïve and is itself a product of an optimism bias. If I can't overcome these biases, not much harm will be done in trying, since I will end up engaging in the risky projects nonetheless due to my optimism bias. So it seems to me that there's nothing to be lost by accepting norms that recommend being truthful about myself.

More generally, realistic examples of false beliefs that have practical or epistemic benefits, at least all of the ones that I can think of, are such that a cognitive bias is helpful in countering a practical irrationality. If we were not practically irrational, there would be nothing good about

such cognitive biases. Ideally, I would prefer to overcome both. Often, we can, at least to some extent, overcome both. If such an ideal is therefore helpful and if standard epistemic norms are needed to express it, that is a way in which those norms matter. Of course, that doesn't mean these norms are all that matter. I recognize the value of some of my cognitive biases and therefore have a partially positive attitude towards them even though they pull me away from following standard norms.

Summing up the main thread of this section, I can't think of any realistic example in which I think that the thing to do is to convince myself to believe an obvious falsehood. Therefore, that some otherwise good norm forbids tradeoffs is not enough to show that the norm doesn't matter.

3.2 How wrong norms can matter

There is another problem with the third premise of the argument from tradeoffs. Even if we actually and often enough face permissible tradeoffs, tradeoff-forbidding norms can still matter. Theories can be valuable even if they are, strictly speaking, false. Newtonian mechanics serves as an example of one reason why this can be the case. Strictly speaking, Newtonian mechanics is usually considered inaccurate. However, we still teach Newtonian mechanics, and the formulas are still used because they are simple and accurate enough for designing and predicting the behavior of midsized objects traveling at speeds that don't approach the speed of light.

When thinking about theories that posit norms, we must ask ourselves what roles these norms are supposed to have. If the norms are supposed to be guiding-norms—that is, norms that we are supposed to try to follow—then there can be a variety of reasons why inaccurate norms can be better to follow than accurate ones. One was already mentioned: that the accurate norms might be too complicated for us to follow. We might therefore be better off following heuristics, norms that are simpler and that in most ordinary circumstances guide us correctly. Another possibility is that, due to facts about us, such as psychological facts, it might be the case that when we try to do one thing, we predictably end up doing something else. A famous example is sometimes called the *pleasure paradox*—that is, the claim that one who seeks explicitly to increase their own pleasure will end up with less pleasure than one whose explicit goal is, for example, to

please others.² Put in terms of norms, even if the thing for us to do is to increase our pleasure, the lesson from the pleasure paradox is that we probably should not try to follow the norm that says we should do so. Similarly, it could be that tradeoffs are sometimes what we should do in order to acquire various benefits, but we are better off not explicitly deciding to make such tradeoffs. Alternatively, it may be psychologically impossible for us to decide to make tradeoffs, and that all of the actual tradeoffs we make are subconscious. Both possibilities imply that tradeoff-forbidding norms will give us all the guidance we can use, even if permissible tradeoffs exist.

Interestingly, Talbot himself makes such a comment in a different context. In section 3.4 of the book, he discusses the worry that we can't conform to epistemic norms anyhow, and that if ought implies can, then there can't be epistemic oughts. One of his lines of response (section 3.4.3) is that there is more than one way a norm can be a good guide. Most importantly, a norm that says that one ought to ϕ can guide in ways other than by getting one to actually ϕ . (An interesting example he gives is that a norm can guide one to compensate for their not ϕ -ing.) If so, a norm can be guiding even if we are incapable of fulfilling it, and this helps respond to the ought-implies-can worry. In addition, this line can potentially save standard theories from Talbot's own arguments: It may be that tradeoff-forbidding norms are good guides even if we should sometimes accept tradeoffs.

4 Conclusion

I am not convinced by Talbot's argument that in actual, familiar circumstances we should form beliefs that go against our evidence, are incoherent, or are clearly false. Nor am I convinced that the book puts an end to epistemology (the part of it concerned with norms) as we know it. But when was the last time philosophers agreed on anything anyhow? What's more important is that *The End of Epistemology As We Know It* is a gratifying stimulus for thinking about when and why epistemic norms matter.³

² Parfit (1986, p. 5) calls this phenomenon *indirect self-defeat*. A theory is indirectly self-defeating if it is the case that when one tries to achieve the aims of a theory, those aims are worse achieved.

³ Thanks to Brian Talbot and Joshua Schechter for very helpful discussions and comments on drafts.

Bibliography

Baras, D. (2022). *Calling for Explanation*. Oxford University Press.

<https://doi.org/10.1093/oso/9780197633649.001.0001>

Joyce, R. (2001). *The Myth of Morality*. Cambridge University Press.

<https://doi.org/10.1017/CBO9780511487101>

Kahneman, D. (2011). *Thinking, Fast and Slow*. Farrar, Straus and Giroux.

Parfit, D. (1986). *Reasons and Persons*. Oxford University Press.

<https://doi.org/10.1093/019824908X.001.0001>

Pettigrew, R. (2016). *Accuracy and the Laws of Credence*. Oxford University Press.

<https://doi.org/10.1093/acprof:oso/9780198732716.001.0001>

Scott, M., & Malcolm, F. (2018). Religious fictionalism. *Philosophy Compass*, 13(3).

<https://doi.org/10.1111/phc3.12474>

Sharot, T. (2012). *The Optimism Bias: A Tour of the Irrationally Positive Brain*. Vintage Books.