

Canadian Philosophical Association 2017 Essay Prize Winner

Is Epistemic Normativity Value-Based?

CHARLES CÔTÉ-BOUCHARD *Rutgers University*

ABSTRACT: What is the source of epistemic normativity? In virtue of what do epistemic norms have categorical normative authority? According to epistemic teleologism, epistemic normativity comes from value. Epistemic norms have categorical authority because conforming to them is necessarily good in some relevant sense. In this article, I argue that epistemic teleologism should be rejected. The problem, I argue, is that there is no relevant sense in which it is necessarily good to believe in accordance with epistemic norms, including in cases where the matter at hand is completely trivial. Therefore, if epistemology is normative, its normativity won't come from value.

RÉSUMÉ : Quelle est la source de la normativité épistémique? En vertu de quoi les normes épistémiques possèdent-elles une autorité normative catégorique? Selon la réponse téléologique, la normativité épistémique provient de la valeur: Les normes épistémiques ont une autorité catégorique parce qu'il est nécessairement bon de s'y conformer. Dans cet article, je soutiens que le téléologisme épistémique doit être rejeté. Le problème est qu'il n'y a pas de sens pertinent dans lequel croire en conformité avec les normes épistémiques est nécessairement bon, même lorsque la croyance en question est complètement triviale. Si l'épistémologie est normative, sa normativité ne provient pas de la valeur.

Keywords: epistemic normativity, metaepistemology, value, epistemic norms, epistemic value, normativity, metanormativity

1. Introduction: Value as the Source of Epistemic Normativity

Epistemology is widely seen as a *normative* discipline like ethics. On such a conception, epistemic claims like attributions of knowledge, epistemically justified belief, rational belief, and warranted belief are normative claims, much like moral claims. Taken seriously, this idea implies two theses. First, epistemology

Dialogue 56 (2017), 407–430.

© Canadian Philosophical Association/Association canadienne de philosophie 2017

doi:10.1017/S0012217317000506



and epistemic discourse have to do with norms or ‘oughts.’ More specifically, they have to do with *epistemic* norms or ‘oughts.’ Just like ethics is about figuring out what we morally ought to do, epistemology is about figuring out what we ought to believe, epistemically speaking.¹ Second, it means that these epistemic norms or ‘oughts’ have *necessary* or *categorical* normative authority for agents. Just like moral norms, epistemic norms are automatically reason-giving or authoritative. Why include this second condition?

For any norm, N, we can recognize the fact that N *says* that we should ϕ —i.e., that we should ϕ *relative to* or *according to* N—and still ask whether N has genuine normative authority for us, i.e., whether there is any *normative reason* for us to conform to N. Importantly, most norms are such that the answer to this question can be ‘no.’ Most norms, in other words, do not have necessary or automatic normative authority. For example, there is not necessarily a normative reason to conform to the norms set by etiquette, club rules, honour codes, fashion, the law, tradition, gender stereotypes, and so on. There can plausibly be situations where there is no normative reason to do as these norms require. So, although they imply norms, those claims and domains are not essentially normative since they don’t necessarily imply normative reasons. They imply what we might call ‘norm-relativity,’ but not genuine normativity.²

By contrast, some norms and domains do seem to have necessary normative force. The standard example is morality. For many metaethicists, if you morally ought to ϕ , then there must be a normative reason for you to ϕ . Normative authority, in other words, is a necessary or essential component of moral claims. Plausibly, then, to say that epistemic claims are normative like moral ones is not only to say that they imply norms, but also that these epistemic norms have necessary or categorical authority.³ Unlike the norms set by etiquette, fashion, tradition, and the law, epistemic norms are necessarily reason-giving. Just like moral claims, epistemic claims have normative authority as a necessary or essential component.

¹ As opposed to non-epistemically (i.e., ‘morally,’ ‘prudentially,’ ‘aesthetically,’ ‘legally,’ etc.) and all-things-considered. Although I only mention belief, epistemic claims and epistemic norms also govern disbelief and withholding or suspension of belief. Epistemic claims, then, are claims about what doxastic attitude someone ought to have, from an epistemic standpoint.

² The expression ‘norm-relativity’ as opposed to genuine normativity is from Hattiangadi (2007). Other labels for this distinction include mere ‘rule-implying’ versus ‘reason-implying’ normativity (Parfit, 2011), merely ‘formal’ versus ‘robust’ normativity (McPherson, 2011; Maguire and Woods, MS.), ‘weak’ versus ‘strong’ categoricity (Joyce, 2001), ‘mere’ requirements versus ‘normative’ requirements (Broome, 2013), and merely ‘reducible’ versus ‘irreducible normativity’ (Olson, 2014).

³ For explicit defences of this idea, see Kim (1988), Kelly (2003), Cuneo (2007), Grimm (2009), Rowlands (2013), and Gibbons (2013).

But, if epistemology is normative in this way, then one crucial question arises. What *explains* or *grounds* this epistemic normativity? *Why* or *in virtue of what* is there necessarily a normative reason to follow the evidence, avoid error, believe the truth, i.e., to conform to epistemic norms? One tempting answer is that doing so is necessarily *good* or *valuable*. Epistemic normativity, in other words, is grounded in facts about what is good. I will refer to this kind of account of the source of epistemic normativity as ‘epistemic teleologism’ or ‘teleologism,’ for short. To a first approximation, the core claim of this kind of account is the following:

Epistemic teleologism (ET): epistemology is normative because it is necessarily pro tanto good to conform to epistemic norms.⁴

That is, it is in virtue of the fact that conforming to epistemic norms is necessarily valuable, in some relevant sense, that there is necessarily a normative reason to follow the evidence, avoid error, believe the truth, and the like.⁵

In this paper, I argue that ET is implausible. More specifically, I argue that there is no relevant sense in which it is necessarily pro tanto good to conform to epistemic norms. ET should therefore be rejected.⁶ The starting point of my argument, which I introduce in Section 2, is what I label the ‘triviality challenge.’ The challenge, in short, is to explain why there seems to be cases where conforming to epistemic norms *does not matter at all*. It is to explain, in other words,

⁴ Here and in what follows, I take ‘good’ to mean either finally (or intrinsically) or instrumentally good. I discuss additional senses of ‘good’—i.e., epistemic, non-epistemic, attributive, prudential, and predicative value—below.

⁵ As I construe it, ET grounds epistemic normativity in value, where value is not grounded in our desires or valuing attitudes. I therefore take ET to be distinct from *desire-based* or *instrumentalist* accounts, according to which epistemic normativity comes from our desires or valuing attitudes. On desire-based accounts, epistemology is normative because conforming to epistemic norms necessarily promotes our desires or goals. Several philosophers have argued that such instrumentalist accounts are hopeless—see, e.g., Kelly (2003), Cuneo (2007), and Lockard (2013). However, many philosophers do not ground value in desires or valuing attitudes. Therefore, value-based accounts that do not ground value in desires—what I call ‘ET’—require a separate evaluation. The goal of this paper is to offer such an evaluation. I discuss desire-based accounts of epistemic normativity in Côté-Bouchard (2015).

⁶ Note that I do not take ET to be equivalent or necessarily committed to epistemic *consequentialism*, namely the view that the epistemic good is prior to the epistemic right. Although they are closely related, I do not think that the failure of the former entails the failure of the latter. I return to this issue below.

why it seems like there is sometimes *nothing* good in believing the truth, avoiding error, following the evidence, and the like. As I argue in Section 3, tel-eologists cannot plausibly meet the triviality challenge. To say that conformity to epistemic norms is necessarily good can either mean that it is good in a *prudential* sense (good *for us*), in a *predicative* sense (good *simpliciter*), or in an *attributive* sense (good *qua member of a kind*). The problem, I argue, is that (i) conforming to epistemic norms is not necessarily good *for us*, (ii) it is not necessarily just plain good either, and (iii) even if it is necessarily *attributively* good, that fact alone cannot vindicate epistemic normativity. I end in Section 4 by briefly exploring what the failure of ET means for metaepistemology and normativity theory more generally.

2. The Triviality Challenge

2.1 *The Triviality Intuition*

According to ET, it is necessarily good to conform to epistemic norms. However, an obvious worry with this suggestion is that we can easily imagine cases where there seems to be nothing good in believing the truth, avoiding error, following the evidence, and the like.

First, we can imagine cases where believing in conformity with epistemic norms would be overwhelmingly bad and would not lead to anything good (and where violating epistemic norms would be overwhelmingly good and would not lead to anything bad). Here is one such example:

Illness. After making several tests, Rita's doctor has bad news. She has a very serious illness and only has a 5% chance of survival. This particular illness is highly sensitive to patients' anxiety and stress levels. Since Rita is frightened of death and illness, following the evidence and believing the truth about her prognosis will lower her chances to almost 0. On the other hand, believing falsely and unjustifiably that she will almost certainly survive will dramatically increase her chances.

Second, there are cases where there is seemingly nothing bad and a lot of good in violating epistemic norms as well as nothing good in conforming to epistemic norms, given the triviality of the matter at hand. Here is one example:

Even stars. Myriam is fascinated by even numbers. She thinks they are the most harmonious and aesthetically pleasing thing in the universe. Realizing that something is in an even quantity always fills her with awe and joy. One night, Myriam looks at the sky and suddenly realizes something incredible: the number of stars could very well be even. This thought fills her with such excitement and wonder that she decides to start pretending that the number of stars really is even. Why not, she thinks? It is not as if being wrong about the number of stars could have any adverse consequences. Weeks pass by and her pretence unconsciously and gradually turns into a genuine, but epistemically unjustified, belief.

Third, there are cases where there is seemingly nothing good in conforming to epistemic norms as well as nothing bad in violating them because of the triviality of the matter at hand. For example:

Dream. Ten years ago, Vincent had a dream in which he saw a historian on TV saying that a pub in London named The Red Lion closed its doors on 1 February 1748. Today, something randomly reminds him of the content of that dream. However, it has been so long since he had that dream that he cannot remember if he really saw that on TV or if it was a dream. But since he does not care at all about such a trivial and insignificant issue he just assumes that it really happened. He thus forms the false and epistemically unjustified belief that a pub named The Red Lion closed its doors on 1 February 1748 as a result. He goes on to believe that for the rest of his life.

We could multiply the examples. There does not seem to be necessarily something good in believing the truth, avoiding error, or following the evidence concerning, e.g., the number of blades of grass on your neighbour's lawn at 11:59 am on 11 May 2016, how long it took the 14th customer served today at the third most popular coffee shop in New Jersey to drink her coffee, the average number of threads in the carpets of all the blue houses in Ireland, your grandfather's favourite sexual position when he was 39 years old, and so on. The lesson of such triviality cases is what I will refer to as the 'triviality intuition':

Triviality intuition: there seems to be cases where there is nothing good in conforming to epistemic norms.

Of course, the triviality intuition alone does not suffice to falsify ET. However, it does pose a pressing explanatory challenge or question for teleologists.⁷ If, as ET states, it is necessarily good to conform to epistemic norms, then why are there cases where there is seemingly nothing good in doing so?⁸ The triviality intuition directly challenges ET. So teleologism is only viable if it can give a satisfactory answer to that challenge.

⁷ See also Whiting (2013) for relevant discussion.

⁸ For those who might not share the triviality intuition, the challenge can be reformulated as the following two-part question. First, why are there cases where it seems *to many of us* that there is nothing good in conforming to epistemic norms? Second, why or in what sense is it supposed to be necessarily good to conform to epistemic norms, even in those triviality cases? The problem, as I will argue below, is that there is no sense of 'good' in which it is plausible to conclude that conforming to epistemic norms is good even in triviality cases.

2.2 *The Challenge: Explaining the Triviality Intuition*

Before examining teleologists' most promising response to the triviality challenge, it is necessary to first rule out one natural, but inadequate, answer. Call it the 'epistemic value' explanation:

The epistemic value explanation: the intuition is due to a failure to distinguish between epistemic and non-epistemic value. Once we make this distinction, we see that, although there is nothing non-epistemically good in conforming to epistemic norms in triviality cases, doing so is still epistemically good.

Although initially tempting, this answer won't do.

Merely pointing out the necessary *epistemic* value of conforming to epistemic norms is not sufficient for vindicating ET and epistemic normativity. To see this, note that when it comes to value and value claims, we need to make a distinction that is analogous to the one between genuine normativity and mere norm-relativity. Call this the distinction between 'genuine' or 'normativity-grounding' value on the one hand, and mere 'domain-relative' value on the other. The former is goodness, which entails normativity and normative reasons. It is the kind of value that is such that if ϕ -ing is good in that sense, then there is a normative reason to ϕ . The latter only implies the trivial claim that something is good *relative to* or *according to* the norms of a particular domain. Unlike genuine goodness, mere domain-relative goodness alone does not entail genuine normativity. The sole claim that ϕ -ing is good relative to some domain does not entail that there is a normative reason to ϕ .

For example, there is always something legally good in respecting the law and something legally bad in breaking it. Conformity to the law, in other words, is trivially good from a legal standpoint and violating the law is bad from a legal standpoint. However, it is a further question whether it is necessarily genuinely good to do what is legally good. Similarly, unfashionable outfits are, trivially, bad from a fashion standpoint and fashionable outfits are good from a fashion standpoint. They respectively have fashion disvalue and value. But it is a further question whether it is genuinely good to do what has fashion value and to avoid being unfashionable.

In the same way, it is trivial to claim that there is necessarily something *epistemically* good in conforming to epistemic norms. Obviously, violating epistemic norms is bad from an epistemic standpoint. But this only vindicates domain-relative value. That is, it only shows that conformity to epistemic norms is good relative to the epistemic domain. Even if we admit that, it is a further question whether it is always genuinely good to do what is epistemically good. It is a further question, in other words, whether epistemic value necessarily constitutes genuine, normativity-grounding value.

At this stage, I am not (yet) claiming that conforming to epistemic norms is not necessarily genuinely valuable. Instead, I am making two points. First, merely

associating epistemic norms with *epistemic* value leaves open the question whether it is necessarily genuinely good to conform to epistemic norms. Second, the kind of goodness that is relevant to ET and epistemic normativity is normativity-grounding goodness, not mere domain-relative goodness. Hence, the more precise definition of ET is the following:

Epistemic teleologism (ET): epistemology is normative because it is necessarily pro tanto *genuinely* good to conform to epistemic norms.

Epistemic value, in other words, necessarily constitutes normativity-grounding value according to ET. However, what the triviality intuition suggests is precisely that it is not always *genuinely* good to believe what is *epistemically* good to believe. This is why the epistemic value explanation won't do as a response to the triviality challenge.⁹

A more promising avenue for teleologists is to invoke the difficulty to tell whether something really has no value at all or if it just has *very little* value. Call this the 'minimal value' explanation:

The minimal value explanation: we have the triviality intuition because (i) the triviality cases are ones where conformity to epistemic norms is only very minimally (genuinely) good, and (ii) it is easy to mistake very little pro tanto (normativity-grounding) value with no value at all.

Simply invoking the minimal value explanation is not sufficient for getting teleologists off the hook, however. They must also show that such an explanation is appropriate in the case of epistemic norms. This is because one could invoke the same kind of explanation to defend the necessary authority *any* norm, including norms that clearly lack necessary normative authority.

Consider examples like the law or etiquette. In response to cases where there is seemingly nothing genuinely good in conforming to the law or etiquette—e.g., cases of immoral, absurd, or obscure laws or etiquette rules—one could give a

⁹ This is also why ET is distinct, at least in principle, from epistemic consequentialism, i.e., the view that the epistemic good is prior to the epistemic right or that what we epistemically ought to believe is solely a matter of conduciveness to epistemic value. What I call 'ET' in this paper is an answer to the metanormative and metaepistemological question 'what is the source of epistemic normativity?' But this question is distinct from the substantive *epistemological* question that consequentialists answer, namely 'what should we believe, epistemically speaking?' and 'what is the relation between the right and the good *within the epistemic domain*?' An epistemic consequentialist answer to this substantive question is compatible with the conclusion of this article, namely that it is not necessarily genuinely good to conform to epistemic norms.

minimal value explanation and say that these are simply cases where conformity to the law or etiquette is only very minimally good, which is easy to confuse with no value at all. But such an explanation would obviously not be sufficient for vindicating the necessary genuine value of following the law or etiquette. Hence, in addition to invoking the minimal value explanation, teleologists must tell us why, unlike in the case of laws or etiquette, that explanation is warranted in the case of epistemic norms.

In the rest of this paper, I will argue that the minimal value explanation is not warranted in the epistemic case.¹⁰ It is not plausible to think that necessarily, conforming to epistemic norms is minimally genuinely good.¹¹

3. Against the Minimal Value Explanation

In what sense exactly is conformity to epistemic norms supposed to be necessarily good according to the minimal value explanation? Philosophers standardly distinguish between three main senses of ‘good’ or ‘valuable.’ These are the *prudential* or ‘good for’ sense, the *predicative* or ‘good *simpliciter*’ sense, and the *attributive* or ‘good *qua member of a kind*’ sense. Accordingly, we can distinguish between three possible versions of the minimal value explanation. Its central claim is either that conforming to epistemic norms is necessarily (i) good *for* us, (ii) good *simpliciter*, or (iii) good in an attributive sense. The problem, I will argue, is that all three versions are implausible responses to the triviality challenge. The minimal value explanation is not convincing or successful under any of the three main senses of ‘good.’

3.1 Minimal Prudential Value

Something can be good in the sense of being *prudentially* good or good *for* someone. That is, something can be good in the sense that it makes one’s life go well and contributes to one’s *well-being* or *interests*. Therefore, a first possible interpretation of the minimal value explanation is a *prudential* version, which holds the following thesis:

Minimal prudential value: conforming to epistemic norms is necessarily good *for* us, at least minimally.

¹⁰ See also Wrenn (2015b) for relevant discussion.

¹¹ To be clear, my claim won’t be that there is *never* anything good in conforming to epistemic norms. My rejection of ET is compatible with the claim that doing so is *almost always* good or good *in general*. Similarly, it is compatible with the claim that it is good that we have epistemic norms and epistemic concepts. After all, we can plausibly say similar things about domains that lack categorical normativity. Following social conventions is generally or very often good. It is also good that we have social conventions. Yet there is not necessarily a normative reason to follow convention.

As the examples from Section 2 already indicate, however, it is clearly not always good for us to believe the truth, avoid error, follow the evidence, and the like.

In *Dream*, the matter at hand is so trivial and disconnected from any of Vincent's interests that conforming to epistemic norms and suspending his judgement about whether the event in question really happened would have had absolutely no impact on his well-being. Similarly, in *Even stars*, Myriam's violating epistemic norms and believing that the number of stars is even only has positive impacts on her well-being and interests. Having that epistemically unjustified belief seems to only make her life go better. Finally, in examples like *Illness*, not only would violating epistemic norms be overwhelmingly in Rita's interests, conforming to epistemic norms would be highly *detrimental* to her well-being. Believing what is epistemically justified for her to believe would only make her life go dramatically worse.

One possible reply is that, although it is not always *actually* good for us to conform to epistemic norms, it is always *potentially* so. It might be claimed, for example, that conforming to epistemic norms necessarily provides us with potential premises for successful good actions. It might also be that violating epistemic norms is always potentially *bad* for us because, for example, any false belief is a potential obstacle to successful good actions. Could teleologists ground epistemic normativity in such potential prudential value? There are three problems with this strategy.

First, this response shows, at best, that conforming to epistemic norms is always *potentially* good. But this does not mean that there is necessarily a normative reason to conform to epistemic norms. Rather, it means that there is always *potentially* a normative reason to conform to epistemic norms, which falls short of categorical epistemic normativity. The fact that there *might* be a reason for you to ϕ does not mean that there *is* a reason for you to ϕ .

Second, the argument proves too much. Even if we concede that conforming to epistemic norms is always potentially good for us, *violations* of epistemic norms are also potentially good for us. Even if my belief that my train is leaving at noon is false and epistemically unjustified, there *could* be a schedule change to noon at the last minute, in which case my unjustified belief would allow me to catch my train. Similarly, any true or epistemically justified belief is also potentially harmful. Therefore, if the potential usefulness of conforming to epistemic norms sufficed to vindicate epistemic normativity, it would also establish, implausibly, that there is necessarily a normative reason to *violate* epistemic norms.

Third, and most importantly, it is not even the case that conforming to epistemic norms is always potentially good for us. Imagine that in *Dream* there is an evil demon that will kill Vincent right away if he conforms to epistemic norms and suspends his judgement about the question at hand. In such a situation, it is not even potentially good for Vincent to conform to epistemic norms. And yet, suspending judgement about the question at hand is still the epistemically justified doxastic attitude for Vincent.

Another common line of defence of minimal prudential value is to construe ET not as a *case-based* thesis—i.e., not as the claim that conforming to epistemic norms is good for us *in all particular cases*—but rather as a *rule-based* thesis about what *policies* or *rules* of belief formation are good to adopt in the long run.¹² Epistemic normativity, on this view, is grounded in the prudential value of a policy of systematically conforming to those norms. So, even if we can be in particular cases where conforming to epistemic norms is not good for us, it remains good for us, over time, to systematically conform to epistemic norms, as a matter of general rule or policy.¹³

This rule-based strategy is not promising for two reasons. First, for agents like those in triviality cases, a policy of always conforming to epistemic norms *except in this one triviality case* would clearly be prudentially better than a policy of systematically conforming to such norms. Even in the long run, they would be better off following the former policy since it would give them the benefits of the latter without its costs. Adopting the latter would not add any prudential value. It is not clear, then, why it would be good for these agents to *systematically* conform to epistemic norms instead of going by a policy that allows them to slip this one time. One might object that a policy of systematically conforming to epistemic norms except in triviality cases is too unrealistic and difficult to follow. But a policy of systematically conforming to epistemic norms is no less difficult to follow. It is just as unrealistic to expect to stick by a policy that requires you to systematically follow your evidence, avoid error, and the like. So this reply is not available to teleologists who want to employ a rule-based strategy.

Second, triviality cases can be modified in ways that cause problems for both case-based and rule-based ET. Suppose that in *Dream*, Vincent is reminded of the content of his dream after drinking poison while jumping off a plane without a parachute alone in the middle of nowhere. As in the original case, the triviality of the question means that, in the short term, no prudential value would come out of conforming to epistemic norms. But since he is seconds away from a certain death, it would not be good for Vincent in the long run either. Yet, even in this extraordinary situation, Vincent's belief that his dream really happened would still be epistemically unjustified.

A final possible defence of a prudential version of ET is that having beliefs that conform to epistemic norms is good for humans, viewed as a group. It is plausible, for instance, that having such beliefs and norms allow us to pool and share reliable information. But since having reliable information is necessary for our survival and well-being, epistemic norms and epistemic goods are clearly good for us.¹⁴

However, this line of thought is not a defence of categorical epistemic normativity. It is not, in other words, a defence of the claim that there is

¹² I borrow the labels 'case-based' and 'rule-based' from Lockard (2013).

¹³ See, for instance, Kornblith (1993) and Leite (2007). See also Kelly (2003), (2007), and Lockard (2013) for discussion.

¹⁴ Thanks to Michael Hannon for bringing this possibility to my attention.

necessarily a reason for agents to conform to epistemic norms. At best, it shows that it is good, for the human race, to, for example, have the practice of making epistemic claims, use epistemic concepts, study epistemology, conform to epistemic norms *most of the time*, etc. But none of these conclusions support the necessary normative authority of epistemic norms. Even if we accept that it is good for humans to use epistemic concepts, to make epistemic claims, to generally conform to epistemic norms, and so on, the triviality challenge could still be successful. It could still be the case that conforming to epistemic norms is not necessarily good for agents. After all, the same kind of story could be told about many norms or domains that lack categorical authority. It is probably good for the human race to have developed norms like etiquette, laws, social conventions, and the like. Yet it is not necessarily good for us to conform to these norms.

At this point, one might respond by questioning the notion of epistemic normativity we have been working with. Why think of epistemic normativity as the notion that epistemic norms have necessary or categorical authority for agents? Do we really want or need more than, for example, the conclusion that conforming to epistemic norms is good most of the time? Perhaps we don't. Perhaps epistemology does have the same kind of authority as etiquette or the law,¹⁵ and perhaps this weaker kind of authority also deserves the label 'normativity.' But whether or not that is the case, it remains true that contemporary epistemologists and normativity theorists typically view epistemology as normative in the categorical or necessary sense. And my concern in this article is only epistemic normativity so understood.¹⁶

3.2 Minimal Predicative Value

A second possible sense of 'X is good' is a *predicative* sense, according to which X is good *simpliciter* or *just plain* good, even though it might not be good for anyone. Given this second sense of 'good,' the minimal value explanation of the triviality intuition can also be given a predicative interpretation that implies the following claim:

Minimal predicative value: conforming to epistemic norms is necessarily good simpliciter, at least minimally.

¹⁵ See, e.g., Maffie (1990), Laudan (1990a), (1990b), Papineau (2013), Hazlett (2013), and Maguire and Woods (MS.). See also Grimm's (2009) reading of Sosa (2007).

¹⁶ In addition to its popularity, another reason to focus on this sense of epistemic normativity is its metaethical or metanormative relevance. Although morality is widely thought to raise metaphysical problems, domains like etiquette, fashion, and the law are not. One natural explanation for that difference is that, unlike the latter domains, normative authority is essential to morality. So the question of whether the authority of epistemology is like that of morality has significant metaethical and metanormative implications. For relevant discussions, see, e.g., Cuneo (2007), Rowlands (2013), and Olson (2014).

This seems to be what many teleologists have in mind, especially those who hold that true belief and knowledge are *intrinsically* good or good *for their own sake*.¹⁷ This second option is also problematic, however. The claim that conforming to epistemic norms is necessarily good simpliciter leads to at least two deeply counterintuitive conclusions.

First, if conforming to epistemic norms necessarily led to plainly good states of affairs, there would necessarily be a normative reason to go change the world to make it fit our beliefs. In particular, there would necessarily be such a normative reason precisely and only because it would make our beliefs true or instances of knowledge. As Joseph Raz explains:

[I]magine that in all cases, if we have a belief about a certain matter then it is *pro tanto* better to have a true rather than a false belief, just because it is true. Consider an example: A month ahead of time I believe that Red Rod will win the Derby or that the Social Democrats will win the elections in Denmark. There may be ways to increase the likelihood that my belief is true. Perhaps I could give valuable advice to Red Rod's jockey, or lend my expertise to the Social Democrats. Is the fact that that will make it more likely that my beliefs are true a reason to do so? If there is value in one's beliefs being true as such then there should be no difference between making reality conform to the belief and making the belief conform to how things are.¹⁸

But this is a deeply counterintuitive result. Plausibly, there is not necessarily a normative reason to act so as to make our beliefs true just because it would give us true beliefs or knowledge. Most of the time, this is no reason to go change the world. Suppose I falsely and unjustifiably believe that the water in Toronto is poisonous. That belief would become knowledge if I poisoned the water. But that is obviously not a reason for me to poison the water.

Second, minimal predicative value leads to the conclusion that there is necessarily a normative reason for anyone to do *anything*. This second objection starts from the popular Anscombean thesis that you cannot count as ϕ -ing intentionally unless you know that you are ϕ -ing. As Kieran Setiya explains:

[...] if I have no idea that in humming Beethoven's Ninth I am driving my wife crazy, I simply cannot be driving her crazy *intentionally*—at least not so far as my humming goes. And if I am ignorant of the impatient tapping of my foot, as I pore over a draft of these pages, it too must be unintentional.¹⁹

¹⁷ See, e.g., Zagzebski (2003), Lynch (2004), Horwich (2006), and Kvanvig (2008). See, e.g., David (2005), McGrath (2005), Whiting (2013), and Wrenn (2015a), (2015b) for critical discussion.

¹⁸ Raz 2011, p. 45.

¹⁹ Setiya 2003, p. 343.

Part of what it is to ϕ intentionally, then, is knowing that you are ϕ -ing. But, if knowledge is necessarily plainly good, it follows that there is necessarily something good in ϕ -ing intentionally, namely the knowledge that you are ϕ -ing. Hence, ϕ -ing intentionally necessarily brings about something good, namely knowledge. Epistemic teleologists are therefore led to an absurd conclusion: since it is necessarily good to intentionally do *anything*, there is necessarily a normative reason for anyone to intentionally do anything, including the most horrible, absurd, or pointless acts you can imagine.

Of course, one could reject the claim that intentional action essentially involves knowledge. However, even opponents to the Anscombean thesis would have to admit that knowledge is still part of *most* cases of intentional action. But this weaker claim—that knowledge is involved in most cases of intentional action—suffices to generate an almost equally counterintuitive conclusion, namely that there is *almost always* a normative reason for anyone to do anything.

Could teleologists bite these two bullets in response? They could, but it would put ET in an unfavourable position in the current dialectic. This would only add to the intuitive cost of ET, which already includes rejecting the triviality intuition. The problem is that the teleological minimal value explanation is up against an alternative explanation, the denial of ET, which does not have these costs. It is hard to see, in that context, why the minimal value explanation would be preferable.

Are there considerations that *favour* minimal predicative value? Michael Lynch and Jonathan Kvanvig have each offered prominent cases for the claim that true beliefs are intrinsic or final goods. If sound, these arguments could perhaps tip the balance in favour of ET and the minimal predicative value.

According to Lynch, we all care about truth for its own sake and not for the sake of something else. We are committed, in other words, to the final value of believing the truth and avoiding error. To support this claim, Lynch offers the following two thought experiments:

Some super neuroscientists give you the choice between continuing to live normally, or having your brain hooked up to a supercomputer that will make it *seem* as if you are continuing to live normally (even though you're really just floating in a vat somewhere). When in the vat, you will continue to have all the same experiences you would have in the real world. Because of this, you would believe that you are reading a book, that you are hungry, and so on. In short, your beliefs and experiences will be the same, but most of your beliefs will be false. If we didn't really prefer true beliefs to false ones, we would be simply ambivalent about this choice. Vat, no vat; who cares? But we don't say this. We don't want to live in the vat, even though doing so would make no difference to what we experience or believe. This suggests that we have a basic preference for truth.²⁰

²⁰ Lynch 2004, p. 17.

Suppose that, unbeknownst to us, the world began yesterday—it *seems* older, but it isn't. If I really lived in a 'Russell world,' as I'll call it, almost all my beliefs about the past would be false. Yet my desires would be equally satisfied in both worlds. This is because the future of both worlds unfolds in exactly the same way. [...] In other words, whatever plans I accomplish now, I would also accomplish if the world had begun yesterday, despite the fact that in that case, my plans would be based on false beliefs about the past. Yet, given the choice between living in the actual world and living in a Russell world, I strongly prefer the actual world. Of course, once "inside" that world, I wouldn't see any difference between it and the real world; in both worlds, after all, events crank along in the same way. But that is beside the point. For the fact remains that thinking about the worlds only insofar as they are identical in instrumental value, *there is difference right now* between the two worlds that matters to me. Even when it has no effect on my other preferences, I—and presumably you as well—prefer true beliefs to false ones.²¹

If Lynch is right that we care about truth for its own sake, then we have the basis for a promising argument in favour of ET. In general, the fact that we value something for its own sake is strong evidence that it *really is* valuable for its own sake.²² That argument can be summarized as follows:

1. We all value the truth for its own sake.
2. If we all value the truth for its own sake, then, in all likelihood, the truth really is valuable for its own sake.
3. Therefore, in all likelihood, the truth is valuable for its own sake.
4. Therefore, in all likelihood, it is necessarily good to conform to epistemic norms.²³

There are two main problems with this argument, both of which concern premise 1.

First, even if we admit that Lynch's thought experiments provide evidence that we value the truth for its own sake, our data also includes the triviality intuition. While Lynch's thought experiments might be informative, we must also remember the triviality cases from Section 2, which suggest precisely that we do *not* value the truth for its own sake. Hence, even if Lynch's examples indicate that we value the truth for its own sake, their force is mitigated by the triviality intuition.

²¹ Lynch 2004, p. 18.

²² See Whiting (2013) for a helpful discussion of this point.

²³ For the sake of argument, I take for granted that conforming to epistemic norms is the best means to the goal of believing the truth and hence that if believing the truth is necessarily good, then so is conforming to epistemic norms.

Second, and more importantly, these thought experiments do not convincingly show that we care about the truth for its own sake. Lynch is surely right that, in his first example, most of us would decline the neuroscientists' offer. However, the reason is that it would rob us of many of the things we value most. If I entered the machine, I would no longer, for instance, be in contact with the people I love. I would stop being involved in their lives and they would stop being involved in mine. I would not help anyone ever again, I would stop improving the lives of others, I would not achieve anything anymore, I would not witness the great historical events of my lifetime, I would stop travelling, going to concerts, and so on. Granted, many of these things probably require knowing the truth. However, that does not mean that I value knowing the truth for its own sake. What I value is being in contact with loved ones, helping others, travelling, going to concerts, and so on. Not *knowing* that I am.

Regarding his second example, I should say that it is not clear to me that we would all prefer the actual world to a world that—unbeknownst to us—just came into existence. After all, a mind-boggling amount of horrible things happened up until yesterday in the actual world. But, if we are in the world imagined by Lynch, then none of these horrible things actually happened. No war, no illness, no genocide, no torture. Clearly, some of us might prefer that.

But, even if we set that aside, it is not clear that this preference would essentially have to do with knowing the truth. I would prefer the actual world because the other option would mean that most of the things that give my life value and meaning never actually happened. It would mean, for instance, that I never improved the life of anyone, that I never fulfilled promises, that I never finished my PhD, that I never built a relationship with my spouse, that the Montreal Canadiens never won the Stanley Cup, and so on.

The point is that, in both examples, our preference for the actual world can plausibly be explained without invoking our valuing the truth for its own sake. Of course, that alone does not entail that we do not, in fact, value the truth for its own sake. However, it does mean that Lynch's thought experiments fail to show that we do.²⁴

Kvanvig also argues that true belief is valuable for its own sake, but via reflection on what the *cognitively ideal* agent would be like:

[I]magine a world where no practical needs are left unmet and where no limitation of cognitive power creates any need for informational content to trump any value for truths with little or no content. [...] We should ask ourselves, regarding possible individuals in such a cost-free environment, what the cognitive ideal would involve. [...] Part of the cognitive ideal, whatever else it may involve, is knowledge of all

²⁴ For more on Lynch's arguments, see David (2005), McGrath (2005), Whiting (2013), and Wrenn (2015a), (2015b).

truths; omniscience, for short. But for omniscience to be part of the ideal, no truth can be pointless enough to play no role at all in the story of what it takes to be cognitively ideal.²⁵

Kvanvig seems to offer the following argument:

1. In cost-free environment where no practical needs are left unmet, a cognitively ideal agent would know all truths.
2. The cognitive ideal in such a world would not consist in knowing all truths if true belief was not valuable for its own sake.
3. Therefore, true belief is valuable for its own sake.

What does ‘cognitive’ in ‘cognitive ideal’ refer to exactly? One natural suggestion is that it just means ‘epistemic.’ On that reading, the cognitively ideal agent just is the agent who is ideal from an epistemic point view. This is a natural reading given that the cognitive is typically cashed out in terms of the ‘mind-to-world’ direction of fit. If ‘cognitive’ is defined in terms of the aim of representing reality accurately, then it is virtually equivalent to ‘epistemic.’

The problem with this reading of ‘cognitive ideal,’ however, is that it makes Kvanvig’s argument invalid. As we have already seen, merely showing that something is epistemically good is not sufficient to show that it is genuinely good and not merely domain-relatively good. It might be true that the epistemic ideal would involve omniscience. However, it is a further question whether it is necessarily genuinely good to do what the epistemic ideal being would do, i.e., to believe what is epistemically best to believe.

Compare this with the idea of the *etiquette* ideal. Trivially, the ideal agent from the standpoint of etiquette would always conform to every rule of etiquette. But that does not mean that it is always genuinely good to behave in the way that the etiquette ideal would behave. In the same way, no matter how we characterize the epistemic ideal, the crucial question is whether it is always genuinely good to do what the epistemically ideal agent would do. What the triviality intuition suggests is precisely that it is not.

An alternative reading of the first premise is that the ‘cognitive ideal’ does not refer merely to the ‘epistemically’ ideal agent, but rather to the cognition of the *genuinely ideal* agent. In other words, perhaps the first premise of Kvanvig’s argument should be interpreted as the claim that an agent who always does what is genuinely best always behaves as the epistemically ideal agent would behave. However, if this is how the first premise should be read, then the argument becomes question begging since it takes ET for granted. Instead of showing that epistemic value necessarily constitutes normativity-grounding

²⁵ Kvanvig 2008, pp. 209-210.

value, this second reading of the argument uses this claim as part of its first premise. So Kvanvig's argument is either invalid or question begging.

But perhaps I have misinterpreted Kvanvig's passage. So far, I have taken his 'cognitive ideal' point to be a deductive argument in favour of the value of truth. However, another interpretation is that Kvanvig only aims to provide an intuition pump. Under this interpretation, his point is simply that, when we try to imagine the cognitively ideal agent in a cost-free environment with no practical needs unmet, our intuitions tell us that such an agent is omniscient. Suppose Jon and Joan live in such an environment. Jon knows all truths, except for one completely inconsequential and obscure fact, about which he is wrong. Joan, in contrast, really knows all truths. Isn't there something better with Joan compared to Jon? Isn't Jon a little further from the ideal than Joan? But, if that is our intuition, then aren't we committed to the claim that there is at least something good in any true belief, however pointless?

A first thing to note is that if Kvanvig is only offering an intuitive case, then it is doubtful that it can suffice to vindicate minimal predicative value in the present context. Even if we admit Kvanvig's intuition, the predicative version of ET still has the counterintuitive implications outlined above and we still have the triviality cases. So, even if it provides *some* support for minimal predicative value, Kvanvig's intuition likely won't be enough to tip the scale in teleologists' favour.

What should we make of Kvanvig's intuition? Once again, we must distinguish between what it takes for an agent to be ideal from an epistemic standpoint and what the cognition of a genuinely ideal agent would be like. As we have already seen, there is plausibility to the idea that the agent who's ideal from an epistemic standpoint would be omniscient. So, in a sense, it is plausible that there is something worse or less than ideal with Jon compared to Joan: he is epistemically worse than her. He is less than epistemically ideal. But, once again, this claim only entails the domain-relative claim that knowing the truth is *epistemically* valuable.

What teleologists need is the intuition that Joan is *genuinely* better than Jon. But, once we set aside domain-relative value, it becomes less obvious that there must be something just plain worse with Jon. If the only difference between the two is really just a single, completely pointless piece of knowledge, then it is hard to see why Joan's situation would be genuinely better. One possible response is that, if we had to choose between being in Jon's situation and Joan's situation, most of us would choose Joan's. This might be true, but it is not clear that our answer to this question can be trusted as a guide to what is genuinely valuable.

To see this, consider an alternative version of the Jon and Joan case. Imagine that Jon conforms to all the rules of Victorian table etiquette, except for one completely inconsequential, unnoticeable, and practically unenforceable rule. Joan, in contrast, conforms to all the rules of Victorian table etiquette, including the obscure one that Jon violates. Once again, both agents are in a cost-free environment where no practical needs are left unmet. Hence, there are no costs

or benefits associated with following or violating the obscure rule in question. But what would we do if we had to choose between being in Jon's situation and Joan's situation? I suspect that most of us would choose Joan's. After all, there is no practical cost associated with conforming to that additional rule and no benefit to violating it. So why not? Yet conforming to Victorian table etiquette is obviously not necessarily good simpliciter.

So, for all these reasons, Kvanvig's 'cognitive ideal' point does not show that conforming to epistemic norms is necessarily good simpliciter.²⁶ More generally, given the failure of the two most prominent cases in its favour as well as the problems outlined above, I conclude that the predicative version of the minimal value explanation is implausible. Conformity to epistemic norms does not appear to be necessarily just plain good.

3.3 *Minimal Attributive Value*

This leaves the third, attributive sense of 'good.' This is the sense of 'good' that we use to refer to something being good *qua member of a kind K*. To be a good member of K in this sense is to do or be able to do what members of K are supposed to do *qua members of K*. It is, in other words, to achieve or be able to achieve the defining aim, function, or telos of members of that kind. We use this attributive sense of 'good' when we judge, for instance, that X is a good toaster, a good football player, a good teacher, and so on. Corresponding to this third kind of value is a final possible version of the minimal value explanation, which holds the following:

Minimal attributive value: there is necessarily something attributively good in conforming to epistemic norms, at least minimally.

There are several possible ways to flesh out this idea, but the most obvious one is in terms of good and bad *beliefs*. Perhaps beliefs that conform to epistemic norms are necessarily good, well-functioning, or successful *qua beliefs*. Such beliefs, the suggestion goes, are as they are supposed to be *qua beliefs*. Just like toasters that toast well, beliefs that conform to epistemic norms fulfil the constitutive aim or function of beliefs.²⁷

The problem, however, is that, even if it is true, this attributive value claim alone cannot vindicate ET and epistemic normativity. This is because goodness

²⁶ For more on Kvanvig's argument, see Whiting (2013) and Wrenn (2015a), (2015b).

²⁷ This is a common way to interpret the popular idea that 'belief constitutively aims at truth.' See, for instance, Velleman (2000), Burge (2003), Bergmann (2006), Bird (2007), Fassio (2011), McHugh (2011), Graham (2012), Littlejohn (2013), and Nolfi (2015). I discuss the connections between the aim of belief and epistemic normativity more extensively in Côté-Bouchard (2016).

qua member of a kind is not necessarily genuine or normativity-grounding goodness. The mere fact that X is well-functioning or successful *qua* member of kind K does not entail that it is genuinely good to be X, to have X, or to promote X. It leaves open the question of whether it is genuinely good to be a good member of K or to do what good members of that kind do. Hence, it only entails mere domain-relative value.

For example, part of being a good torturer is to make one's victims suffer intensely for a long time. This is what torturers are supposed to do *qua* torturers. Similarly, part of being a good hired killer is to murder one's victims quickly and without leaving a trace. This is what hired killers are supposed to do *qua* hired killer. But those attributive value claims do not make it genuinely good to make people suffer intensely for a long time or to murder swiftly and without leaving a trace. Even if we accept that this is what good torturers and hired killers do, such attributive value claims leave open the question of whether it is genuinely good to be a good torturer or a good hired killer. This means, in turn, that it leaves open the normative question of whether there is any normative reason to do what good torturers or good hired killers do.

Therefore, showing that conforming to epistemic norms means having well-functioning beliefs won't suffice for vindicating ET and epistemic normativity. Even if the underlying attributive value claim is true, it still leaves open the normative question of whether it is necessarily genuinely good to have beliefs that are good *qua* beliefs and to avoid beliefs that are bad *qua* beliefs. But this is what teleologists need to show. What the triviality intuition suggests is precisely that there is sometimes no normativity-grounding value in having well-functioning beliefs and in avoiding defective beliefs.

The same problem arises if we claim instead that agents who conform to epistemic norms are necessarily good *qua* *believers*, *qua* *inquirers*, *qua* *epistemic agents*, and the like. Even if they are true, these attributive value claims only entail domain-relative goodness since they leave open the normative question of whether it is necessarily genuinely good to be a good believer, a good inquirer, a good epistemic agent, and the like. Once again, this is precisely what the triviality intuition calls into question.

Could teleologists claim, instead, that conforming to epistemic norms necessarily makes us well-functioning *qua* *human beings*? This may be more promising since it is admittedly harder to deny that being good *qua* human agent is necessarily genuinely good. The problem, however, is that it is implausible that violating epistemic norms necessarily makes people defective *qua* human beings and that conforming to epistemic norms necessarily contributes to making one a well-functioning human agent. As we have already seen, to say that X is a good member of a kind K is to say that it does what members of that kind are supposed to do. It is to say, in other words, that it can achieve the constitutive, essential, or defining aim, function, or *telos* of members of that kind. But what is the essential *telos* of human agents? When are humans defective or well-functioning?

One natural answer is that we are well-functioning insofar as we *flourish*. As we have already seen, however, conforming to epistemic norms is not necessarily good *for* us and violating epistemic norms is not necessarily bad for us. But it is hard to see how ϕ -ing could be necessary for our flourishing without being also good for us. The reasons that it is not necessarily good for us to conform to epistemic norms are also plausibly reasons that it is not necessarily conducive to our flourishing.

Another natural answer would be that the essential aim of human agency is autonomy and rational action. Agents are well-functioning, in other words, insofar as they can intentionally do what they think they ought to do. But you can very well violate epistemic norms and still be autonomous in that way. Perhaps you cannot be autonomous if you are *generally* incompetent at conforming to epistemic norms. However, there can certainly be cases where failing to follow the evidence about particularly trivial and insignificant questions won't interfere with your autonomy or ability to act as you think you ought. Hence, it won't suffice to vindicate the necessary normative authority of epistemic norms.

Could conformity to epistemic norms itself be a constitutive telos of human beings? Could knowing the truth and avoiding error about any matter (including complete trivialities) be a defining aim of human agency? It is not clear why this would be the case. It cannot be because we all have the capacity to systematically conform to epistemic norms. We do not. Most people have deeply entrenched biases and limitations that routinely prevent them from, for example, following their evidence. And at least some of those people are unable to get rid of those biases because they do not care or because it is too difficult.

Neither can it be because we all actually *want* to know the truth and avoid error about any question with which we might be confronted. At least some human beings do not care to know the truth about, say, completely trivial and obscure questions. Once we add to this the fact that conforming to epistemic norms can often be *detrimental* to people's interests, well-being, and flourishing, it is simply implausible to suppose that human agents qua human agents are supposed to systematically conform to epistemic norms. If human agency must really have an epistemically relevant telos, other candidates seem more plausible. Why isn't it to, e.g., conform to epistemic norms often enough to flourish, be autonomous, or maintain optimal physical and psychological health? After all, this is the kind of thing that most of us would want and could at least come close to achieving.

So the claim that it is necessarily attributively good to conform to epistemic norms is either false or insufficient for grounding epistemic normativity. Hence, minimal attributive value cannot be invoked to meet the triviality challenge either.

4. Conclusion

According to ET, epistemic normativity comes from value. Epistemic norms have categorical authority because it is necessarily good to conform to them.

I have argued that ET should be rejected because it cannot meet the triviality challenge. The most promising strategy for teleologists is to argue that, even in triviality cases, there is still a minimal amount of value in conforming to epistemic norms. I have argued that this explanation is unconvincing because there is no sense of ‘good’ in which it is plausible to conclude that conforming to epistemic norms is necessarily genuinely good, even minimally. So, if epistemology is normative, its normativity won’t come from value.

What does this mean for metaepistemology and normativity theory more generally? One important upshot is that three popular ideas form an inconsistent triad. These are the idea that (i) epistemology is categorically normative, (ii) normativity is a unitary phenomenon which has a single source, and (iii) at least some kinds of normativity—e.g., practical normativity—are grounded in value. Since, as I have argued, teleologism fails in the epistemic case, we cannot keep all three claims. We must give up at least one. If epistemology is normative, then either normativity is unitary, in which case no normativity is value-based, or some normativity (e.g., practical normativity) is value-based, in which case, there must be different kinds of normativity with different sources. We could also keep theses (ii) and (iii), and view normativity as both unitary and value-based. But we would then have to give up the idea that epistemology is categorically normative and view it instead as analogous to domains like etiquette.

Acknowledgements: I am especially grateful to David Black, Simon-Pierre Chevarie-Cossette, Marc-Kevin Daoust, Michael Hannon, Anandi Hattiangadi, Hallvard Lillehammer, Clayton Littlejohn, and Jonathan Way for discussions and feedback on earlier versions of this paper. Thanks to audiences at the 2017 Congress of the Canadian Philosophical Association at Ryerson University, the 2016 *Understanding Value* conference at the University of Sheffield, the summer seminar in philosophy at King’s College London in 2016, the 2015 SOPHA Congress in Montréal, and the 2015 Stockholm graduate conference in Philosophy at Stockholm University. Thanks also to the FRQSC and the Rutgers Philosophy department for their support.

References

- Anscombe, G.E.M.
1957 *Intention*. Cambridge, MA: Harvard University Press.
- Bergmann, Michael
2006 *Justification without Awareness: A Defense of Epistemic Externalism*. Oxford: Oxford University Press.
- Bird, Alexander
2007 “Justified Judging.” *Philosophy and Phenomenological Research* 74 (1): 81–110.
- Broome, John
2013 *Rationality through Reasoning*. Chichester, UK: Wiley-Blackwell.

- Burge, Tyler
2003 "Perceptual Entitlement." *Philosophy and Phenomenological Research* 67 (3): 503–548.
- Côté-Bouchard, Charles
2015 "Epistemic Instrumentalism and the Too Few Reasons Objection." *International Journal of Philosophical Studies* 23 (3): 337–355.
- Côté-Bouchard, Charles
2016 "Can the Aim of Belief Ground Epistemic Normativity?" *Philosophical Studies* 173 (12): 3181–3198.
- Cuneo, Terence
2007 *The Normative Web: An Argument for Moral Realism*. Oxford: Oxford University Press.
- David, Marian
2005 "On 'Truth is Good.'" *Philosophical Books* 46 (4): 292–301.
- Fassio, Davide
2011 "Belief, Correctness and Normativity." *Logique et Analyse* 54 (216): 471–486.
- Gibbons, John
2013 *The Norm of Belief*. Oxford: Oxford University Press.
- Graham, Peter
2012 "Epistemic Entitlement." *Noûs* 46 (3): 449–482.
- Grimm, Stephen R.
2009 "Epistemic Normativity." In *Epistemic Value*, edited by Adrian Haddock, Alan Millar, and Duncan Pritchard, 243–264. Oxford: Oxford University Press.
- Hattiangadi, Anandi
2007 *Oughts and Thoughts: Rule-Following and the Normativity of Content*. Oxford: Oxford University Press.
- Hazlett, Allan
2013 *A Luxury of the Understanding: On the Value of True Belief*. Oxford: Oxford University Press.
- Horwich, Paul
2006 "The Value of Truth." *Noûs* 40 (2): 347–360.
- Joyce, Richard
2001 *The Myth of Morality*. Cambridge: Cambridge University Press.
- Kelly, Thomas
2003 "Epistemic Rationality as Instrumental Rationality: A Critique." *Philosophy and Phenomenological Research* 66 (3): 612–640.
- Kelly, Thomas
2007 "Evidence and Normativity: Reply to Leite." *Philosophy and Phenomenological Research* 75 (2): 465–474.
- Kim, Jaegwon
1988 "What is 'Naturalized Epistemology?'" *Philosophical Perspectives* 2: 381–405.

- Kornblith, Hilary
1993 "Epistemic Normativity." *Synthese* 94 (3): 357–376.
- Kvanvig, Jonathan
2008 "Pointless Truth." *Midwest Studies in Philosophy* 32 (1): 199–212.
- Laudan, Larry
1990a "Aim-Less Epistemology?" *Studies in History and Philosophy of Science Part A* 21 (2): 315–322.
- Laudan, Larry
1990b "Normative Naturalism." *Philosophy of Science* 57 (1): 44–59.
- Leite, Adam
2007 "Epistemic Instrumentalism and Reasons for Belief: A Reply to Tom Kelly's Epistemic Rationality as Instrumental Rationality: A Critique." *Philosophy and Phenomenological Research* 75 (2): 456–464.
- Littlejohn, Clayton
2013 "The Russellian Retreat." *Proceedings of the Aristotelian Society* 113 (3pt3): 293–320.
- Lockard, Matthew
2013 "Epistemic Instrumentalism." *Synthese* 190 (9): 1701–1718.
- Lynch, Michael
2004 *True to Life: Why Truth Matters*. Cambridge, MA: MIT Press.
- Maffie, James
1990 "Naturalism and the Normativity of Epistemology." *Philosophical Studies* 59 (3): 333–349.
- Maguire, Barry, and Jack Woods
Ms. "Explaining Epistemic Normativity."
- McGrath, Matthew
2005 "Lynch on the Value of Truth." *Philosophical Books* 46 (4): 302–310.
- McHugh, Conor
2011 "What Do We Aim At When We Believe?" *Dialectica* 65 (3): 369–392.
- McPherson, Tristram
2011 "Against Quietist Normative Realism." *Philosophical Studies* 154 (2): 223–240.
- Nolfi, Kate
2015 "How to be a Normativist about the Nature of Belief." *Pacific Philosophical Quarterly* 96 (2): 181–204.
- Olson, Jonas
2014 *Moral Error Theory: History, Critique, Defence*. Oxford: Oxford University Press.
- Papineau, David
2013 "There Are no Norms of Belief." In *The Aim of Belief*, edited by Timothy Chan, 64–79: Oxford: Oxford University Press.
- Parfit, Derek
2011 *On What Matters*. Oxford: Oxford University Press.

Raz, Joseph

2011 *From Normativity to Responsibility*. Oxford: Oxford University Press.

Rowland, Richard

2013 "Moral Error Theory and the Argument from Epistemic Reasons." *Journal of Ethics and Social Philosophy* 7 (1): 1–24.

Setiya, Kieran

2003 "Explaining Action." *Philosophical Review* 112 (3): 339–393.

Sosa, Ernest

2007 *A Virtue Epistemology*. Oxford: Oxford University Press.

Velleman, David

2000 *The Possibility of Practical Reason*. Oxford: Oxford University Press.

Whiting, Daniel

2013 "The Good and the True (Or the Bad and the False)." *Philosophy* 8 (2): 219–242.

Wrenn, Chase B.

2015a *Truth*. Cambridge: Polity.

Wrenn, Chase B.

2015b "Truth is Not (Very) Intrinsically Valuable." *Pacific Philosophical Quarterly* 97 (2): 108–128.

Zagzebski, Linda

2003 "The Search for the Source of Epistemic Good." *Metaphilosophy* 34 (1-2): 12–28.