

This article was downloaded by: [Kings College London]

On: 09 July 2015, At: 13:40

Publisher: Routledge

Informa Ltd Registered in England and Wales Registered Number: 1072954

Registered office: 5 Howick Place, London, SW1P 1WG



## International Journal of Philosophical Studies

Publication details, including instructions for authors  
and subscription information:

<http://www.tandfonline.com/loi/riph20>

### Epistemic Instrumentalism and the Too Few Reasons Objection

Charles Côté-Bouchard<sup>a</sup>

<sup>a</sup> King's College London, UK

Published online: 15 Jun 2015.



CrossMark

[Click for updates](#)

**To cite this article:** Charles Côté-Bouchard (2015): Epistemic Instrumentalism and the Too Few Reasons Objection, International Journal of Philosophical Studies, DOI: [10.1080/09672559.2015.1042007](https://doi.org/10.1080/09672559.2015.1042007)

**To link to this article:** <http://dx.doi.org/10.1080/09672559.2015.1042007>

PLEASE SCROLL DOWN FOR ARTICLE

Taylor & Francis makes every effort to ensure the accuracy of all the information (the "Content") contained in the publications on our platform. However, Taylor & Francis, our agents, and our licensors make no representations or warranties whatsoever as to the accuracy, completeness, or suitability for any purpose of the Content. Any opinions and views expressed in this publication are the opinions and views of the authors, and are not the views of or endorsed by Taylor & Francis. The accuracy of the Content should not be relied upon and should be independently verified with primary sources of information. Taylor and Francis shall not be liable for any losses, actions, claims, proceedings, demands, costs, expenses, damages, and other liabilities whatsoever or howsoever caused arising directly or indirectly in connection with, in relation to or arising out of the use of the Content.

This article may be used for research, teaching, and private study purposes. Any substantial or systematic reproduction, redistribution, reselling, loan, sub-licensing, systematic supply, or distribution in any form to anyone is expressly forbidden. Terms & Conditions of access and use can be found at <http://www.tandfonline.com/page/terms-and-conditions>

# Epistemic Instrumentalism and the Too Few Reasons Objection

Charles Côté-Bouchard

## Abstract

According to epistemic instrumentalism (EI), epistemic normativity arises from and depends on facts about our ends. On that view, a consideration  $C$  is an epistemic reason for a subject  $S$  to  $\Phi$  only if  $\Phi$ -ing would promote an end that  $S$  has. However, according to the Too Few Epistemic Reasons objection, this cannot be correct since there are cases in which, intuitively,  $C$  is an epistemic reason for  $S$  to  $\Phi$  even though  $\Phi$ -ing would not promote any of  $S$ 's ends. After clarifying both EI and the Too Few Epistemic Reasons objection, I examine three major instrumentalist replies and argue that none of them is satisfactory. I end by briefly sketching a fourth possible response, which is, I suggest, more promising than the other three.

**Keywords:** epistemic normativity; epistemic instrumentalism; epistemic reasons; normativity; meta-epistemology; epistemology

## 1. From Normative Instrumentalism to Epistemic Instrumentalism

One of the main issues that divide contemporary accounts of normativity is the question of its 'source'. *In virtue of what* is something a normative reason for someone to  $\Phi$ ? What is it that *grounds* the normative force that makes some consideration a reason for someone to  $\Phi$ ? Where does that normative or reason-giving force come from?

According to what I will call *normative instrumentalism* (NI), it comes from facts about our ends, aims, goals, or desires (e.g., Foot, 1972; Williams, 1981; Smith, 1994; Joyce, 2001; Finlay, 2006; and Schroeder, 2007).<sup>1</sup> More precisely, when a consideration  $C$  is a reason for an agent  $S$  to  $\Phi$ , that consideration has this normative force for  $S$  because or in virtue of the fact that  $\Phi$ -ing would realize or promote an end that  $S$  has.<sup>2</sup> Suppose I want to sing karaoke and tonight happens to be karaoke night at my local pub. In such a situation, it seems that the fact that tonight is karaoke night at the pub is a reason for me to go to the pub. According to normative instrumentalism, that consideration has this normative force in virtue of the fact that I want to sing karaoke and that going to the pub would promote that end of mine.

Essential to NI is the thesis that normativity *depends* on our ends. According to this dependence thesis, C is a reason for S to  $\Phi$  only if  $\Phi$ -ing would realize or promote an end that S has. If  $\Phi$ -ing does not at least promote the realization of an end of S, then there is no reason for S to  $\Phi$ . In the above example, if I were to lose my desire to sing karaoke and if this was the only end that would have been promoted by going to the pub, then the fact that tonight is karaoke night at the pub would cease to be a reason for me to go to the pub tonight. My losing the relevant end, in other words, would thereby rob that consideration of its normative force for me.

Normative reasons come in different varieties or kinds including, for instance, moral, prudential, and epistemic ones. This paper is about epistemic reasons. In particular, I will focus on a specific subset of epistemic reasons, namely *evidential* or *truth-related* reasons – *evidential reasons* for short. By evidential reasons, I mean considerations (i) that count in favor of a doxastic state D – i.e., belief, disbelief, or withholding of belief with respect to a proposition P – and (ii) that do so by bearing on whether P is true. Suppose I drive by my mother's house and see her car in the driveway. That my mother's car is in her driveway is plausibly an evidential reason for me to believe that she is home since that consideration counts in favour of that belief by indicating that she is home.

Note that I am not taking 'evidential reasons' to be synonymous with 'evidence'. In particular, not all evidential reasons for D-ing with respect to P are evidence for or against P.<sup>3</sup> A consideration can bear on whether P is true without being evidence for or against P. Consider, for example, the proposition that the number of stars is even. We have no evidence for or against that proposition. But that consideration – that there is no available evidence for or against that proposition – actually counts in favor of a particular doxastic state, namely that of withholding belief with respect to whether the number of stars is even. Moreover, it counts in favour of that doxastic state by bearing on the truth of that proposition in some sense. Yet, that consideration is not itself evidence for or against the proposition that the number of stars is even.

As I explained above, normative instrumentalists locate the source of normativity and reasons in facts about our ends. If this is right, then our ends are also what grounds more specific varieties or kinds of normativity and reasons, including epistemic reasons. NI thus entails what I will call *epistemic instrumentalism* (EI), i.e., the view that epistemic reasons come from facts about our ends. More precisely, when C is an epistemic reason for S to  $\Phi$ , that fact has this normative force for S because or in virtue of the fact that  $\Phi$ -ing would realize or promote an end that S has. Focusing on the subset of epistemic reasons that will be the concern of this paper, EI entails that when a consideration C is an evidential reason for S to have some doxastic state D with respect to some proposition P, that consideration has this normative force for S because or in virtue of the fact that D-ing would realize or promote an end that S has.

Thus, what I called the dependence thesis – the idea that normativity depends on our ends – is also an essential upshot of EI. If EI is true, then C is an epistemic reason for S to  $\Phi$  only if  $\Phi$ -ing would realize or promote an end that S has. To use the example above, that my mother's car is in her driveway is an evidential reason for me to believe that she is home only if believing that proposition would realize or promote an end that I have. If I have no such end, then that consideration cannot have the normative force that would make it a reason for me to believe that proposition. Call this application of the dependence thesis to the case of epistemic reasons the *epistemic dependence thesis*.<sup>4</sup>

Although the issue of the source of normativity has been less discussed in epistemology, epistemic instrumentalism can be attributed to at least some epistemologists and philosophers of science (e.g., Foley, 1987, 1993; Maffie, 1990; Laudan, 1990, 1991; Kornblith, 1993; and Leite, 2007).<sup>5</sup> EI can also be attributed to those who defend instrumentalism about normativity in general. However, many (if not most) instrumentalists only explicitly discuss *practical* normativity and reasons for action. Nevertheless, there is reason to think that these philosophers should endorse epistemic instrumentalism as well. First, one plausible desideratum in the philosophy of normativity is to provide a unified account of normativity and normative reasons. Hence, the inapplicability of instrumentalism beyond the practical domain would itself be a problem for the theory. Second, it is not clear why, if true, instrumentalism would only apply to one species of normativity. What instrumentalism provides is an account of the source of the normative force that makes some considerations normative reasons. If other reasons besides practical ones are also genuine normative reasons, why would the source of their normative force be any different?

The main motivation for adopting NI and EI is that it promises to help make normativity and reasons unproblematic from a naturalistic point of view. Normativity is arguably less mysterious if it ultimately has to do with our ends. After all, there clearly are knowable facts about what ends we have and about what would realize or promote them. These notions and relations do not seem involve anything incompatible with a scientific worldview. Relatedly, EI might be seen as the key to reconcile Quine's (1969) project of naturalizing epistemology with the apparently normative nature of that discipline (e.g., Laudan, 1990, 1991; Maffie, 1990; Quine, 1998). It could, for example, serve to elaborate Quine's (1998) so-called 'engineering' reply to critics like Kim (1988).<sup>6</sup>

This paper is about one particular kind of objection against epistemic instrumentalism. It is an instance of what Schroeder (2007, Ch. 6) calls the Too Few Reasons objection – TFR for short – against NI. According to TFR, there are cases in which, intuitively, a consideration C is a reason for S to  $\Phi$  even though  $\Phi$ -ing would not promote any of S's ends. But given its dependence thesis, NI has to deny, implausibly, that these are reasons for S to  $\Phi$ . NI should thus be rejected because it entails the existence of too few reasons.

Similarly, according to the Too Few *Epistemic* Reasons objection – TFER for short – against epistemic instrumentalism, there are cases in which, intuitively, a consideration *C* is an epistemic reason for *S* to  $\Phi$  even though  $\Phi$ -ing would not promote any of *S*'s ends. But given its epistemic dependence thesis, EI has to deny, implausibly, that these are reasons for *S* to  $\Phi$ . EI should thus be rejected because it entails the existence of too few epistemic reasons. After making a case for TFER in section 2, I consider – and ultimately reject – various instrumentalist replies in sections 3 to 5. I end by briefly suggesting an alternative reply to TFER in section 6.

## 2. The Too Few Epistemic Reasons Objection Against EI

If epistemic instrumentalism (EI) is true, then what I called the epistemic dependency thesis is true: *C* is an epistemic reason for *S* to  $\Phi$  only if  $\Phi$ -ing would promote an end that *S* has. However, evidential reasons seem to pose a problem for that thesis. As Thomas Kelly (2003) pointed out, there are cases in which a fact is an evidential reason for *S* to have a doxastic state *D* even though *D*-ing would not promote any of *S*'s ends.<sup>7</sup> Kelly identifies two kinds of such counter-examples. First, there are cases where following your evidential reasons would not promote your ends because you are utterly *indifferent* about the matter at hand. Call these cases of *epistemic indifference*. Second, there are also cases in which following your evidential reasons would not promote your ends because you want to *avoid* learning the truth about *P*. Call these cases of *truth-avoidance*. Here is an example of epistemic indifference:

***Delaware's Beverage*** Like most people, Nancy has absolutely no desire to know what the official beverage of the state of Delaware is. Unbeknownst to her however, her friend Brett, whom she knows to be very reliable, has recently developed a deep obsession for Delaware. One day he comes up to her, grabs her by the shoulders and says: 'Listen to me, Nancy. I've got to tell you something. I just found out that the state beverage of Delaware is milk! Isn't that amazing?'

Here is an example of truth-avoidance:

***Spoiler Alert*** Max missed the finale of his favorite television series, which aired last night and revealed whether character *X* ends up dying or surviving. Given his passion for the show, Max really wants to avoid learning the answer to that question before he watches the rerun tonight after work. But to his dismay, one of his colleagues – who is convinced that Max did watch the show last night – comes up to him in the morning shouting: 'I knew it! I told you character *X* would die!'

Intuitively there is evidential reason for both Nancy and Max to believe the proposition at hand. However, it seems that EI must deny this. After all, given Nancy's indifference and Max's end of avoiding the truth, neither of them has the end of believing the truth about the matter at hand. Following their evidential reasons and believing the truth would not promote any of their ends.

One thing to note, however, is that cases of mere epistemic indifference and truth-avoidance are not sufficient to falsify EI. To think that they are is to attribute to epistemic instrumentalists a view that they do not have to have, namely that C is an epistemic reason for S to  $\Phi$  only if  $\Phi$ -ing would realize or promote a distinctly *epistemic* or *cognitive* end that S has. Such a commitment is not essential to EI. There is no reason think that EI necessarily makes epistemic reasons depend on corresponding *epistemic* ends. All that epistemic instrumentalists claim is that just like any kind of normative reasons, epistemic reasons depend on our ends.<sup>8</sup> More generally, NI is not committed to the view that having a reason of a particular kind depends on having an end of the corresponding kind.

To be fair, it does seem true that whenever a fact is a reason of kind K for S to  $\Phi$ ,  $\Phi$ -ing would, in effect, promote an end of the corresponding kind K whether or not S *has* that particular end of kind K. Take the following prudential example. For most people, the fact that exercising regularly decreases the risk of becoming seriously ill is a prudential reason to exercise regularly. In those cases, following that prudential reason – i.e., exercising regularly – would, in effect, promote a distinctly prudential end, namely that of, e.g., preserving one's long-term health and well-being. However, suppose that unlike most people, I do not have any such prudential ends: I do not care at all about my long-term health and well-being. Does that mean that, according to normative instrumentalism, that prudential reason has no normative force for me? Not necessarily. I might still have some other end that would be promoted by my preserving my long-term health and well-being. Perhaps I really want to read the not yet published novels from George R.R. Martin's series *A Song of Ice and Fire*. Normative instrumentalism certainly allows that, in such a case, the fact that exercising regularly decreases the risk of serious illness *does* constitute a reason for me to exercise regularly. According to NI, that fact has normative force me in virtue of the fact that my end of reading Martin's upcoming novels would be promoted by the realization of the prudential end of preserving my long-term health and well-being, which would itself be promoted by exercising regularly.

In the same way, it does seem true that whenever a fact is an evidential reason for S to D with respect to P, D-ing with respect to P would, in effect, promote an end of the corresponding kind. What end? A plausible suggestion is the end of *believing that P if and only if P*.<sup>9</sup> Suppose the consideration that my mother's car is in her driveway – call it M – is an evidential reason for me to believe that she is home. Believing in accordance with that reason – i.e., believing that my mother is home – would, in effect, promote the end of

believing something if and only if it is true. However, this does not mean that I need to have that specific kind of end in order for M to be an evidential reason for me to believe that she is home. I could simply have another end that would itself be promoted by my believing the truth. Perhaps I want to see my mother in person today before I leave for a long vacation and believing the truth about whether she is home is a necessary means to that end.

Hence, EI is not falsified by mere cases of epistemic reasons for S to  $\Phi$  without S having any *epistemic* end that would be promoted by  $\Phi$ -ing. It can only be falsified by cases where there is an epistemic reason for S to  $\Phi$  even though  $\Phi$ -ing would not promote any of S's ends. But the problem for epistemic instrumentalists is that there does seem to be such cases. Consider Delaware's Beverage again. It seems perfectly possible that believing the truth regarding Delaware's beverage would further *none* of Nancy's ends. This is especially plausible given the trivial and inconsequential character of that question. Call this modified version of the first example *Delaware's Beverage\**. Unfortunately for EI, this makes no difference for our intuition regarding Nancy's evidential reason. It still seems that the fact that Brett tells her that Delaware's beverage is milk is an evidential reason for her to believe that that beverage is milk.

We can also imagine cases in which believing the truth and avoiding error about a question would not only fail to promote any of S's ends, but also severely go *against* S's ends. Here is a version of an often-used example:

***Illness*** After making several tests, Rita's doctor has bad news: she has a very serious illness and she is very unlikely to survive it. However, as it happens, believing (falsely) that she *is* likely to survive – or, even better, that she *will* survive – would considerably increase her chances of surviving. On other hand, believing (truly) that she is very unlikely to make it will actually decrease her chances of survival even more.

Intuitively, the fact that Rita's doctor told her that she is very unlikely to survive her illness is strong evidential reason for Rita to believe that she is very unlikely to survive. The problem for EI is that it does not look like having that true belief would promote any of Rita's ends. In fact, it would actually *frustrate* many of her ends, including that of surviving her illness. Hence, it seems that given the epistemic dependence thesis, EI must say, implausibly, that the fact that Rita's doctor told her that she is unlikely to survive does *not* count as a reason for her to believe it. EI should thus be rejected. This is the Too Few Epistemic Reasons objection (TFER) against EI.

Instrumentalists have three broad options in response to TFER. One is to respond in the same way that some normative instrumentalists have reacted to the apparent inability of NI to accommodate *moral* reasons, which has been to adopt an *error theory* about those reasons. On such a view, it is true that NI cannot accommodate the existence of moral reasons. But since instrumentalism

is true, this just means that there really are no moral reasons (see e.g. Mackie, 1977; Garner, 1990; Joyce, 2001; Olson, 2011a, 2011b). Similarly, epistemic instrumentalists can reply to TFER that although their view cannot accommodate evidential reasons, this just means that there are no such reasons. Call this first option *error theory*.

A second possible option is to deny that in the problematic cases, epistemic reasons are genuine normative reasons at all. When we say that C is an ‘epistemic reason’ for S to  $\Phi$  (at least in cases where  $\Phi$ -ing would not promote any of S’s ends), all we mean is that  $\Phi$ -ing is ‘correct relative to’ or ‘required by’ epistemic norms. What we call ‘epistemic reasons’ in those cases are thus mere ‘pseudo’ reasons, i.e., reasons that follow trivially from some norm or requirement but which lack the normative force of genuine normative reasons.<sup>10</sup> On that view, epistemic reasons are analogous to the reasons that follow from e.g., etiquette, fashion, or convention.<sup>11</sup> Call this second option *pseudo reasons*.

The third option is to respond that contrary to what it first seems, EI can indeed deliver the intuitively right result in the apparently problematic cases. Epistemic instrumentalists, in other words, do have the resources to yield the result that C is an evidential reason (with genuine normative force) for S to D in the alleged counter-examples. The only way they can do this is by arguing that contrary to what it first seems, D-ing does, in some way, promote an end that S has. Call this the *optimistic response*.

I only consider versions of the optimistic response in this paper. This is motivated both by limits of space and by the relative popularity and desirability of that option. The availability of a satisfactory optimistic response would arguably outweigh any motivation to adopt the other two options given the considerably higher intuitive cost of *error theory* and *pseudo reasons*. In the next three sections, I consider and ultimately reject three optimistic responses to TFER that have recently been given. I briefly suggest an alternative reply in section 6.

### 3. Mark Schroeder’s Reply

Mark Schroeder suggests an optimistic response to TFER in *Slaves of the Passions*.<sup>12</sup> He would say that EI is able to deliver the right result in the examples above because following evidential reasons and avoiding false beliefs necessarily promotes our ends.

His key idea is that because of the interconnectedness of our beliefs, believing falsely *any* proposition can potentially start a chain of errors that could end up frustrating *any* of one’s ends. As a result, believing any proposition only if it is true promotes *all* of S’s ends to some extent. Any of S’s ends, in other words, is necessarily promoted by avoiding error with respect to any proposition. Schroeder’s striking conclusion is thus that *each and every* end that S has grounds each and every evidential reason that there is for S to D. To use one

of his examples, take, on the one hand, Mary's end of buying a new pair of shoes, and on the other, the number of moons that Jupiter has. As Schroeder (2007, p. 114) puts it:

Being in error about it might lead to being in error about other things, such that being in error about them might lead to being in error about other things, and so on until something might lead to Mary having trouble getting new shoes. If this is right, then for any proposition, Mary's desire to get a new pair of shoes will serve to explain why there is a reason for Mary to believe it only if it is true.

Schroeder's claim thus entails that there is a reason for Mary to avoid error with respect to how many moons Jupiter has because it would actually promote Mary's end of buying a new pair of shoes. Moreover, given the close connection between our beliefs, that same desire grounds any possible evidential reason that there is for Mary to have a doxastic state D. Whenever C is an evidential reason for S to D, it is true that C has that normative force in virtue of the fact that Mary wants to buy new shoes. Moreover, 'to have new shoes' can be replaced by any of Mary's ends. More generally, each of S's ends grounds each of S's evidential reasons. In Delaware's Beverage\* and Illness, avoiding error with respect to the matter at hand would promote each and every ends that Nancy and Rita have.

This is, on the face of it, highly counterintuitive. Even if we admit that, in some sense, being mistaken about the number of Jupiter's moons could somehow lead to the frustration of Mary's end of buying shoes, it still sounds very strange to say that it actually *promotes* that end, especially if we claim that this is what grounds the normative force of Mary's evidential reason. As Schroeder himself emphasizes, his argument relies on an extremely weak understanding of what it is for something like an act or a doxastic state to *promote* someone's end. Schroeder's (2007, p. 113) account of the promotion relation is this:

*X's doing A promotes p* just in case it increases the likelihood of *p* relative to some baseline. And the baseline, I suggest, is fixed by the likelihood of *p* conditional on *X's* doing nothing – conditional on the status quo.

But adopting such a weak understanding of the promotion relation has problematic consequences when coupled with normative instrumentalism.

The main one is that it seems to entail what McPherson (2012) calls an *explosion* of reasons, i.e., the consequence that, at least in the case of reasons for action, each of my ends makes it the case that I have a reason to do every act. As McPherson (2012, p. 447) puts it:

Consider again the desire to get home on time. For just about anything that I might do – scratching my ear, running the other way, trying to drive my car off a cliff – there is some possible scenario in which this act leads me to get home on time.

For any act A and end E, in other words, doing A would increase the likelihood of E being realized since we can always come up with a possible scenario (however remote) in which A leads to E being realized. In McPherson's example, his desire to get home on time thus makes it the case that he has a reason to scratch his ear, run the other way, try to drive his car off cliff, and so on. Why would such an explosion be problematic? First, there is its prima facie implausibility. That there are things that there is no reason for us to do seems like a truism or, at least, an intuitive datum that any account of normativity should aim at accommodating. Is there really a reason for me to, e.g., jump out of my window while eating two bananas and singing happy birthday to John Travolta in a batman costume? It seems like such an explosion of reasons would trivialize normativity. After all, if it is so easy for consideration to have normative force and be reasons, why does it matter?

Relatedly, there is an intuitive distinction between requirements that necessarily entail or give normative reasons and those that do not. It is often claimed that there is necessarily a reason to do as morality and prudence requires. But apart from these and a few other controversial cases, it seems like most norms or requirements do not necessarily imply or give reasons. Take for example the requirements of etiquette, fashion, convention, the law, language rules, game rules, and so on. Intuitively, there is not necessarily a reason for agents to follow these requirements. As Foot (1972) famously pointed out, the fact that the rules of etiquette apply to you and require you to  $\Phi$  does not mean that you have a reason to  $\Phi$ . The problem is that if, as Schroeder's account seems to entail, there is always a reason to do anything, we lose this intuitive datum. It would mean that contrary to what Foot and others have claimed, there is necessarily a reason to follow the requirements of etiquette, fashion, convention, the law, rules of languages, rules of games, and so on.

A second problem is that it seems possible to modify the counter-examples to EI in a way that actually blocks Schroeder's reply. Take Delaware's Beverage\* for instance. Suppose that an evil demon will kill Nancy instantly if she forms the true belief that milk is Delaware's state beverage. In that case – call it Delaware's Beverage\*\* – it is hard to see how avoiding error with respect to the identity of Delaware's state beverage would promote any of Nancy's ends. Yet, it still seems true that there is an evidential reason for her to believe that that beverage is milk. For all these reasons, Schroeder's reply to TFER does not seem convincing.

#### 4. The Aim-of-Belief Reply

Another kind of optimistic reply to TFER is to show that we necessarily have a particular end that is promoted by following evidential reasons. A prominent version of this strategy starts from the popular idea that belief – or more precisely, the formation and regulation of doxastic states – has a *constitutive aim*. To say that some activity A has a constitutive aim is roughly to say that it is part of what makes something an instance of A-ing that it is performed with that aim. An activity that would not be performed with that aim would thereby not count as an instance of A-ing. Games like Sudoku, for example, seem to have such an aim.<sup>13</sup> It seems that one does not count as doing a Sudoku puzzle unless one has the aim of filling out every square with a number from 1 to 9 without repeating the same number in the same row, column, or square. An activity performed without that aim would simply not count as doing a Sudoku puzzle. One thus has that aim every time one does a Sudoku puzzle.

Analogously, to say that doxastic states have a constitutive aim is to say it is part of what makes something an instance of forming and regulating a doxastic state that it is performed with that aim.<sup>14</sup> An activity that would not be performed with that aim would thereby not count as an instance of forming or regulating a doxastic state. If that is correct, then one has that aim every time one forms or regulates a doxastic state. What would be that aim? One popular suggestion is that it is the aim of believing that P if and only if P. Thus, following that suggestion, it is part of what makes something an instance of forming or regulating a doxastic state with respect to P that it is performed with the aim of believing that P if and only if P. S thus necessarily has that aim every time she forms or regulates a doxastic state.

If all this is correct, then we have a basis for defending the optimistic response to TFER. In both Delaware's Beverage\* and Illness the agent is involved in forming or regulating a doxastic state with respect to P. So if what we have said above is correct – i.e., if doxastic states have the constitutive aim of believing that P if and only if P – then the agents in those examples actually aim at believing that P if and only if P. But as we have seen, following one's evidential reasons tends to promote the end of believing that P if and only if P. Therefore, even in cases like Delaware's Beverage\* and Illness, following the evidential reason does turn out to promote a subject that the subject has, i.e., the constitutive aim of belief.<sup>15</sup>

A first worry with this response is that even if we admit that belief has such a constitutive aim, this aim or end is one that a subject can have only while she is actually forming or regulating a doxastic state. The problem is that intuitively, there can be evidential reasons for S to D with respect to P even when she is not in the process of forming or regulating a doxastic state.

Second, it seems highly unlikely that doxastic states have a constitutive aim in the sense specified above. At least two conditions would have to be fulfilled in order for the aim-of-belief thesis to be of any help against TFER. It would

have to be true that (i) doxastic states have an aim in a literal sense (as opposed to a metaphorical and normative sense) and (ii) that aim is one that agents have. But both of these conditions are highly controversial.

The first one is important because some authors help themselves with the claim that belief has a constitutive truth-aim, but interpret it merely as a metaphor for the claim that belief has a constitutive *norm* such as ‘a belief is correct if and only if it is true’.<sup>16</sup> Some also argue that beliefs *cannot* have an aim in the literal non-normative sense (Owens, 2003; Shah, 2003; Shah and Velleman, 2005).<sup>17</sup> The success of the aim-of-belief reply depends on the falsity of those conceptions and arguments since according to EI, normativity depends on our literal ends.

We would also have to show that the constitutive aim of belief is not only a literal one, but also one that *agents* have and not merely one that our doxastic states themselves have in some sense. Call this distinction the agential versus non-agential construal of the aim-of-belief thesis. Most of those who accept that belief has a constitutive aim in the literal sense seem to understand this as something like a function that defines what the state of belief is, but not as an aim that agents have. On this non-agential construal of the aim of belief, doxastic states aim truth and avoidance of error in roughly the sense that, e.g., the pancreas aims at producing insulin (see e.g., Velleman, 2000; Bird, 2007; and McHugh, 2011, 2012). However, such a construal of the aim-of-belief thesis won’t do for epistemic instrumentalists since according to EI, normativity comes from the ends that *agents* have.<sup>18</sup>

The problem, however, is that the agential construal of the aim-of-belief thesis seems implausible. More precisely, it seems unlikely that *we* have, say, the aim of believing the truth every time we form or regulate a doxastic state with respect to P. First, it remains true that the agents in the examples above seem to lack the aim of believing the truth and avoiding error regarding the matter at hand. In other words, it remains counter-intuitive to say that Nancy and Max are in a situation where they literally aim at believing the truth respectively about Delaware’s state beverage and about the fate of character X.

Second, what is it to  $\Phi$  with an aim that you have? On one natural understanding, it is to  $\Phi$  *intentionally*. The problem is that many (if not all) of our beliefs are not things we do intentionally. Perceptual beliefs and wishful thinking, for example, are clearly not states that we form intentionally. Intuitively, coming to have these beliefs is an unintentional process that just happens to us. Suppose I form the belief that a rock is coming at me as a result of seeing a rock rapidly coming towards my face. Or suppose I stumble upon conclusive evidence that my father just died and form the belief that my father is dead as a result. Saying that I am aiming at believing the truth in these cases seems implausible. It is not as if I was *trying* to get at the truth in those cases. The formation of these beliefs is something that automatically happens to me.

One possible reply is that  $\Phi$ -ing with an aim that we have is not the same as  $\Phi$ -ing intentionally. Perhaps it is possible to form e.g., perceptual beliefs

with an aim without forming them intentionally. This would be the case if we understood the aim of belief not as a *particular* aim that we necessarily have every time we form or regulate a doxastic state, but instead as a *global* or *general* epistemic aim that agents necessarily have. Perhaps agents necessarily have the general or global goal of e.g., avoiding false beliefs and believing true things (generally speaking). If this is were the case, then following evidential reasons would always promote this general end that we all have.

The problem is that this suggestion seems even less plausible than the previous one. Even if we admit that there are some people who have the general end of, e.g., believing the truth and not the false, to claim that all agents necessarily have it seems absurd. There are clearly some agents who do not care about believing the truth regarding at least some matters. Moreover, even if it turned out that everyone actually has such a general end, it would still seem that there *could* be an agent without that goal. It just seems implausible that there could not be an agent that would not have the end of e.g., avoiding false beliefs and believing true things. Hence, it appears that the aim-of-belief reply to TFER is not convincing either.

### 5. Hilary Kornblith's Reply

Hilary Kornblith (1993) suggests a different kind of reply to TFER. It differs from the previous two in that it is best interpreted as a *rule-based* rather than *case-based* reply to TFER (see also Leite, 2007).<sup>19</sup> Case-based replies attempt to show that in every particular case where C is intuitively an evidential reason for S to D, *that particular D-ing* would promote an end that S has.

By contrast, according to the rule-based approach, it would be enough if instrumentalists could show that for any agent, *systematically* following evidential reasons – and thus believing that P if and only if P – as a general *policy* or *rule* would itself promote one's ends. It concedes that there are particular cases in which believing that P if and only if P would not promote any of S's ends. However, it adds that even in those particular cases, it remains true that systematically believing that P if and only if P, as a rule, promotes S's ends. In the alleged counter-examples then, evidential reasons are genuine normative reasons and they inherit their normative force from that of the rule or, more precisely, from the fact that the policy of systematically following evidential reasons necessarily promotes our ends.

Kornblith's rule-based proposal is that the systematic following of evidential reasons best promotes *all* of our ends over time. We all want to achieve our ends. But achieving our ends requires successful practical reasoning. But since successful practical reasoning requires true beliefs, systematically believing a proposition if and only if it is true will promote, over time, the achievement of our ends. As Kornblith (1993, pp. 371–2) puts it:

It seems that someone who cares about acting in a way which furthers the things he cares about, and that includes all of us, has pragmatic reasons to favor a cognitive system which is effective in generating truths, whether he otherwise cares about truth or not. We should thus adopt a method of cognitive evaluation which endorses truth-conducive processes.

Since we all want to achieve our ends and since systematically following evidential reasons promotes that end, evidential reasons have normative force – and are thus genuine reasons – for all of us.

The main problem with Kornblith's reply is that it seems overly and unjustifiably optimistic about the usefulness of following evidential reasons. Successful practical reasoning does not necessarily require true or even rational belief. In fact, for most of our ends, it is possible to imagine a way in which having an irrational or false belief in our practical reasoning would improve our chances of realizing that end. Suppose I want to get from London to Edinburgh today and my only way to do so is to take a train that exceptionally departs at three o'clock today. Of course, one good way to reach my end would be to check the train schedule and form the true belief that the train departs at three today. But an even more straightforward and cost-effective way of realizing the same end would be to believe that the train departs at 3pm as the result of wishful thinking. If we can imagine an example of that kind for most of our ends, then it is not clear that systematically following evidential reasons is the optimal way of generally realizing our ends. Perhaps the optimal belief-forming system for promoting the ends of beings like us in the long run would include *some* unreliable belief-forming and belief-regulating mechanisms, i.e., mechanisms that are not truth-conducive and not responsive to evidential reasons like wishful thinking, cognitive biases, delusions, self-deception, and so on.<sup>20</sup> It is an obvious fact, after all, that such processes constantly affect the formation and regulation of our doxastic states. But why did we end up that way? Could it not be because things like wishful thinking and cognitive biases tend, to a certain extent, to help us reach our ends better than if we always followed our evidential reasons perfectly?

Of course, what belief-forming system would be the optimal one for promoting the ends of beings like us is ultimately an empirical question. In fact, one worry with Kornblith's proposal is that he gives an answer to this empirical question without backing it with empirical evidence. Furthermore, there are reasons to doubt that the instrumentally optimal belief-forming system for us would be completely devoid of unreliable or reasons-unresponsive belief-forming mechanisms. Here are some examples.

First, it does not seem far-fetched to claim that in some cases of intense traumatic experience, not following our evidential reasons (at least for a little while and to a certain extent) is instrumentally better, on the whole, than

following them correctly. There might very well be situations in which responding correctly to your evidential reasons would generate psychological suffering and damage that would frustrate your ends more, in the long run, than by being just a little less responsive to evidential reasons.

Second, it does not seem far-fetched either to claim that at least in certain situations, slightly overestimating your own qualities and capacities is instrumentally better, on the whole, than judging yourself in perfect accordance with your evidential reasons. It seems plausible that you could overestimate yourself by just the right amount, i.e., an amount that would allow you to better realize your ends (e.g., by giving you a significant confidence boost) without any significant instrumental cost.

Finally, there are plausibly cases in which being slightly more optimistic than your evidential reasons warrant is instrumentally better, on the whole, than being perfectly responsive to your evidential reasons. There are clearly cases in which your evidential reasons strongly support a very pessimistic view of your own situation. But since adopting such a pessimistic outlook is likely to make you depressed and discourage you from taking the necessary steps to improving your situation, not being perfectly responsive to your evidential reasons might very well be better, overall, for achieving your ends.

As I said above, what belief-forming system would be the optimal one for promoting the ends of beings like us is ultimately an empirical question. Nevertheless, as these examples suggest, it is highly doubtful that such a system would only include belief-forming and belief-regulating mechanisms that are perfectly responsive to evidential reasons (see also Bortolotti, 2015). It seems that for beings like us, the optimal belief-forming system would include at least *some* unreliable mechanisms. Hence, it appears that Komblith's reply to TFER is not convincing either.

## 6. An Alternative Reply: The Aim of $\Phi$ -ing for Reasons

None of the optimistic replies considered so far have been satisfactory. However, before giving up the optimistic strategy and, possibly, EI, I would like to briefly sketch what I think is a more promising strategy.

We all do, believe, and feel certain things for reasons that are at least *apparent* reasons. More precisely, it seems like an essential characteristic of agents that they often  $\Phi$  for an at least apparent reason that P. I say 'at least' apparent because these apparent reasons sometimes – hopefully often – turn out to be real or genuine ones. As I understand them, genuine reasons are facts or true propositions.<sup>21</sup> Consequently, a reason that P is genuine – and not merely or only apparent – only if P is true. If P is false, then S's reason that P is *merely* an apparent one.<sup>22</sup>

Moreover, it seems that whenever we  $\Phi$  for an at least apparent reason that P, we *believe* that P. After all, it would be absurd to assert, e.g., 'my reason for

bringing an umbrella is that it is raining outside, but it is not raining outside'. It would be equally absurd to assert, e.g., 'Myriam's reason for being angry is that someone stole her cupcakes, but she does not believe that someone stole her cupcakes'. We are sometimes wrong, of course, since the relevant proposition sometimes turns out to be false. When that happens, we  $\Phi$  for a merely apparent reason that P rather than for a genuine one. As a result, whenever we  $\Phi$  for a *genuine* reason that P, we believe that P, and P is true. Therefore, having a true belief that P is a necessary condition for  $\Phi$ -ing for a genuine reason that P.<sup>23</sup>

Given all this, it seems that  $\Phi$ -ing for an at least apparent reason that P necessarily involves believing that P is a *genuine* reason for you to  $\Phi$ .<sup>24</sup> But why is this? Why do we always believe that P is a genuine reason to  $\Phi$  when we  $\Phi$  for an at least apparent reason? My suggestion is that this is because whenever we  $\Phi$  for an at least apparent reason, we *aim* at  $\Phi$ -ing for a genuine reason and not for a merely apparent one. All agents often  $\Phi$  for an at least apparent reason that P, and when they do, they aim at  $\Phi$ -ing for the genuine reason that P.<sup>25</sup> To use John Hyman's (2010, pp. 411–13) metaphor, we want to be 'guided by the facts' in those situations.

How does that help EI answering the Too Few Epistemic Reasons objection? If what I have said so far is correct, then we all have the end of  $\Phi$ -ing for a genuine reason that P whenever we  $\Phi$  for an at least apparent reason that P. But recall that one can  $\Phi$  for a genuine reason that P only if one has a true belief that P.<sup>26</sup> This, however, is not sufficient to help EI since there can be particular cases in which following evidential reasons and believing a particular proposition if and only if it is true won't promote the end of  $\Phi$ -ing for a genuine reason that P whenever we  $\Phi$  for an at least apparent reason that P.

But here, the obvious move for proponents of EI is to invoke the rule-based strategy. Recall that according to that strategy, epistemic instrumentalists only need to show that for any agent, *systematically* following evidential reasons and thus believing that P if and only if P as a general *policy* or *rule*, would itself promote one's ends. But crucially, it seems very plausible to say that systematically following evidential reasons and thus believing that P if and only if P as a general rule would promote the end of  $\Phi$ -ing for a genuine reason that P whenever we  $\Phi$  for an at least apparent reason that P. On the one hand, doing so would promote the avoidance of false beliefs which could both be mere apparent reasons or the cause of mere apparent reasons. On the other, it would also promote the formation of true beliefs about facts that could be both one's genuine reasons and the cause of other of one's genuine reasons.

My suggested reply to TFER is therefore this: Epistemic instrumentalists can deliver the right intuitive result in cases like Delaware's Beverage\* and Illness because they can say that in those examples, there is evidential reason for S to have a doxastic state D in virtue of the fact that D-ing is prescribed by a rule, the systematic following of which promotes S's end of  $\Phi$ -ing for a genuine reason that P whenever she  $\Phi$ -s for an at least apparent reason that P. It

follows that these are not counter-examples to EI and its thesis that epistemic reasons depend and come from on our ends.

## 7. Conclusion

I have shown that evidential reasons pose a problem for epistemic instrumentalism because there seems to be cases in which C is an evidential reason for S to D even though D-ing would not promote any of S's ends. This is what I called the Too Few Epistemic Reasons objection (TFER) against EI. I considered three major instrumentalist replies and found them unconvincing. I ended by briefly sketching a fourth solution which appears more promising. Its three key claims are the following. First, we can only  $\Phi$  for a *genuine* reason that P if we have a true belief that P. Second, we all have the end of  $\Phi$ -ing for a *genuine* reason that P whenever we  $\Phi$  for an *at least apparent* reason that P. Third, the optimal way to realize that end is to systematically follow our evidential reasons. If this is true, then EI does have the resources to deliver the right intuitive result in the alleged problematic cases provided by TFER. It is thus an avenue worth exploring for those who think that reasons and normativity come from our ends.

*King's College London, UK*

## Acknowledgements

I would like to thank the organizers and participants of the 2014 Cracow Workshop in Analytic Philosophy. Special thanks also to the participants of the 4th Edinburgh Graduate Epistemology Conference and of the 2014 European Epistemology Meeting. I am especially grateful to Clayton Littlejohn, Matthew McGrath, Teemu Toppinen, and Terence Cuneo for their very helpful comments. Thanks also to David Papineau, Mike Coxhead, Benjamin Davies, Giulia Felappi, and Sophie Stammers.

## Disclosure statement

No potential conflict of interest was reported by the author.

## Notes

- 1 I use these terms interchangeably in what follows. Common alternative labels for what I call normative instrumentalism include 'desire-based' or 'ends-based' accounts of normativity, 'subjectivism', 'internalism', 'neo-Humeanism', and 'hypotheticalism'.
- 2 In addition to Joyce, most moral error theorists rely on NI more or less explicitly. See for instance Mackie, 1977; Garner, 1990, and Olson, 2011a, 2011b.
- 3 Nor am I committed to the claim that all pieces of evidence are also evidential reasons.

- 4 The literature on epistemic normativity sometimes conflates EI with *value-based* accounts that locate the source of epistemic normativity in value, i.e., in facts about what is valuable, good, or desirable rather than facts about what S desires or values. On that alternative view, when C is an epistemic reason for S to  $\Phi$ , that fact has this normative force for S because or in virtue of the fact that  $\Phi$ -ing would realize or promote something valuable or good. EI and value-based accounts must be distinguished since the latter can understand ‘valuable’ or ‘good’ in an objective, external, or mind-independent way. Since such views would locate the source of epistemic normativity outside of us rather than within us, they would reject the dependence thesis and would thus be opposed to EI.
- 5 See also Kelly, 2003; Cuneo, 2007, Ch. 7, and Lockard, 2013, for helpful critical presentations of EI.
- 6 See Wrenn (2006) for a careful examination both Kim’s objection and Quine’s ‘engineering’ reply.
- 7 See also Cuneo, 2007, and Lockard, 2013.
- 8 This is recognized by Lockard (2013).
- 9 Following evidential reasons probably also promotes other ends such as ‘believing that P *only if* P’, ‘believing that P only if you *know* that P’, etc. The bottom line is that it promotes epistemic success and the avoidance of error.
- 10 This is roughly what Hattiangadi (2007) calls the distinction between normativity and mere ‘norm-relativity’. See also Foot, 1972; Parfit, 2011 (esp. vol. 2 s. 88) and Broome, 2013 (pp. 26–7) for a similar distinction.
- 11 Although he rejects the label ‘pseudo reasons’ and the distinction behind it, Finlay (2006) can be read as defending such a response to TFER.
- 12 See esp. Schroeder, 2007, pp. 113–15.
- 13 I borrow this example from McHugh, 2012.
- 14 Velleman, 2000; Owens, 2000; Steglich-Petersen 2006, 2009; Bird, 2007; and McHugh, 2011, 2012 defend the idea that belief has a constitutive aim in roughly that sense.
- 15 Such a strategy can be attributed to Velleman, 2000; O’Hagan, 2005; Steglich-Petersen, 2009; Tubert, 2010, 659–660, and Wiland, 2012, Ch. 6.
- 16 E.g., Wedgwood, 2002; Shah, 2003; Gibbard, 2003; Shah and Velleman, 2005; Boghossian, 2003; Engel, 2007.
- 17 For replies, see e.g. Steglich-Petersen, 2006, 2009; McHugh, 2012, 2013.
- 18 This point is emphasized by Owens, 2003, and Whiting, 2013.
- 19 I borrow the label ‘rule-based’ from Lockard, 2013.
- 20 See Stich, 1990, and Wrenn, 2010, for related criticisms of the idea behind Kornblith’s strategy.
- 21 Since I have no space to defend that claim, I will simply note that this is the dominant view in the philosophy of normativity. To give just a few examples: Finlay, 2006; Schroeder, 2007; Kearns and Star, 2009; Alvarez, 2010; Raz, 2011; Parfit, 2011; Littlejohn, 2012; Broome, 2013.
- 22 I take this distinction to be equivalent to that between subjective and objective reasons. On that alternative terminology, we all  $\Phi$  for what are, at least, subjective reasons. But sometimes – hopefully often – these subjective reasons turn out to be objective reasons.
- 23 See, e.g., Littlejohn, 2012, Ch. 4, for a defence of that claim.
- 24 In the subjective-objective reason terminology, this means that to  $\Phi$  for an at least subjective reason that P is, in part, to take P to be an objective reason to  $\Phi$ . This is congenial to Schroeder’s definition of subjective reasons. He claims that ‘[f]or R to be a subjective reason for X to do A is for X to believe R, and for it to be the case

that *R* is the kind of thing, if true, to be an objective reason for *X* to do *A*': Schroeder, 2007, p. 14.

25 See Whiting, 2014, for a similar proposal.

26 Here is a possible worry with this. When *P* is *someone's* reason for  $\Phi$ -ing or, equivalently, the reason *for which* someone  $\Phi$ -s, *P* is a *motivating* reason. The problem is that according to many philosophers, motivating reasons are not (or not necessarily) facts; *P* can be both *S's* motivating reason and be false. If that's true, then contrary to what I claim, believing truly that *P* is not necessary for  $\Phi$ -ing for the genuine reason that *P*. Although I do not share this view of motivating reasons, my account can be made compatible with such a view by replacing ' $\Phi$ -ing for a genuine reason that *P*' by ' $\Phi$ -ing for a *true* reason that *P*'. The core idea, which I think is equally plausible, would then be that whenever we  $\Phi$  for a reason that *P*, we aim at  $\Phi$ -ing for a true reason that *P*. After all, whether or not motivating reasons are facts, it remains true that whenever we  $\Phi$  for the reason that *P*, we believe that *P* is true.

### References

- Alvarez, M. (2010) *Kinds of Reasons*, Oxford: Oxford University Press.
- Bird, A. (2007) 'Justified Judging', *Philosophy and Phenomenological Research* 74(1): 81–110.
- Boghossian, P.A. (2003) 'The Normativity of Content', *Philosophical Issues* 13: 31–45.
- Bortolotti, L. (2015) 'The Epistemic Innocence of Motivated Delusions', *Consciousness and Cognition* 33: 490–499.
- Broome, J. (2013) *Rationality Through Reasoning*, Malden, MA: Wiley-Blackwell.
- Cuneo, T. (2007) *The Normative Web*, Oxford: Oxford University Press.
- Engel, P. (2007) 'Belief and Normativity', *Disputatio* 2(23): 179–203.
- Finlay, S. (2006) 'The Reasons that Matter', *Australasian Journal of Philosophy* 84: 1–20.
- Foley, R. (1987) *The Theory of Epistemic Rationality*, Cambridge, MA: Harvard University Press.
- . (1993) *Working Without A Net*, Oxford: Oxford University Press.
- Foot, P. (1972) 'Morality as a System of Hypothetical Imperatives', *Philosophical Review* 81: 305–16.
- Garner, R. (1990) 'On the Genuine Queerness of Moral Properties and Facts', *Australasian Journal of Philosophy* 68(2): 137–46.
- Gibbard, A. (2003) 'Thoughts and Norms', *Philosophical Issues* 13: 83–98.
- Hattiangadi, A. (2007) *Oughts and Thoughts: Rule-Following and the Normativity of Content*, Oxford: Oxford University Press.
- Hyman, J. (2010) 'The Road to Larissa', *Ratio* 23(4): 393–414.
- Joyce, R. (2001) *The Myth of Morality*, Cambridge: Cambridge University Press.
- Kearns, S., and D. Star (2009) 'Reasons as Evidence', in R. Shafer-Landau (ed.) *Oxford Studies in Metaethics: Volume 4*, Oxford: Oxford University Press.
- Kelly, T. (2003) 'Epistemic Rationality as Instrumental Rationality: A Critique', *Philosophy and Phenomenological Research* 66(3): 612–40.
- Kim, J. (1988) 'What is "Naturalized Epistemology"?', in J.E. Tomberlin (ed.) *Philosophical Perspectives: Volume 2*, Atascadero, CA: Ridgeview, pp. 381–405.
- Kornblith, H. (1993) 'Epistemic Normativity', *Synthese* 94: 357–76.
- Laudan, L. (1990) 'Normative Naturalism', *Philosophy of Science* 57: 44–59.
- . (1991) 'Aimless Epistemology?', *Studies in History and Philosophy of Science* 21: 315–22.
- Leite, A. (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): 456–64.

- Littlejohn, C. (2012) *Justification and the Truth-Connection*, Cambridge: Cambridge University Press.
- Lockard, M. (2013) 'Epistemic Instrumentalism', *Synthese* 190: 1701–18.
- Mackie, J.L. (1977) *Ethics: Inventing Right and Wrong*, Harmondsworth: Penguin.
- Maffie, J. (1990) 'Naturalism and the Normativity of Epistemology', *Philosophical Studies* 59: 333–49.
- McHugh, C. (2011) 'What Do We Aim at When We Believe?', *Dialectica* 65(3): 369–92.
- . (2012) 'Belief and Aims', *Philosophical Studies* 160(3): 425–39.
- . (2013) 'Normativism and Doxastic Deliberation', *Analytic Philosophy* 45(4): 447–65.
- McPherson, T. (2012) 'Mark Schroeder's Hypotheticalism: Agent-Neutrality, Moral Epistemology, and Methodology', *Philosophical Studies* 157(3): 445–53.
- O'Hagan, E. (2005) 'Belief, Normativity and the Constitution of Agency', *Philosophical Explorations* 8(1): 39–52.
- Olson, J. (2011a) 'Error Theory and Reasons for Belief', in A. Reisner and A. Steglich-Petersen (eds) *Reasons for Belief*, Cambridge: Cambridge University Press.
- . (2011b) 'In Defense of Moral Error Theory', in M. Brady (ed.) *New Waves in Metaethics*, Basingstoke: Palgrave Macmillan.
- Owens, D. (2000) *Reason Without Freedom: The Problem of Epistemic Normativity*, London: Routledge.
- . (2003) 'Does Belief Have an Aim?', *Philosophical Studies* 115: 283–305.
- Parfit, D. (2011) *On What Matters: Volume One*, Oxford: Oxford University Press.
- Quine, W.V. (1969) 'Epistemology Naturalized', in *Ontological Relativity and Other Essays*, New York: Columbia UP.
- . (1998) 'Reply to Morton White', in L.E. Hahn and P.A. Schilpp (eds) *The Philosophy of W.V. Quine*, 2nd edn, Chicago, IL: Open Court.
- Raz, J. (2011) *From Normativity to Responsibility*, Oxford: Oxford University Press.
- Schroeder, M. (2007) *Slaves of the Passions*, New York: Oxford University Press.
- Shah, N. (2003) 'How Truth Governs Belief', *Philosophical Review* 112: 447–82.
- Shah, N., and J.D. Velleman (2005) 'Doxastic Deliberation', *The Philosophical Review* 114: 497–534.
- Smith, M. (1994) *The Moral Problem*, Oxford: Blackwell.
- Steglich-Petersen, A. (2006) 'No Norm Needed: On the Aim of Belief', *The Philosophical Quarterly* 56: 499–516.
- . (2009) 'Weighing the Aim of Belief', *Philosophical Studies* 145: 395–405.
- Stich, S. (1990) *The Fragmentation of Reason*, Cambridge, MA: MIT Press.
- Tubert, A. (2010) 'Constitutive Arguments', *Philosophy Compass* 5(8): 656–66.
- Velleman, J.D. (2000) 'On the Aim of Belief', in *The Possibility of Practical Reason*, Oxford: Oxford University Press.
- Wedgwood, R. (2002) 'The Aim of Belief', *Philosophical Perspectives* 16: 267–97.
- Whiting, D. (2013) 'Truth: The Aim and Norm of Belief', *Teorema* 32(3): 121–36.
- . (2014) 'Reasons for Belief, Reasons for Action, the Aim of Belief, and the Aim of Action', in C. Littlejohn and J. Turri (eds) *Epistemic Norms*, Oxford: Oxford University Press.
- Willard, E. (2012) *Reasons*, New York, NY: Continuum.
- Williams, B. (1981) "Internal and External Reasons" [1979], *Moral Luck*, Cambridge: Cambridge University Press, pp. 101–13.
- Wrenn, C. (2006) 'Epistemology as Engineering?', *Theoria* 72(1): 60–79.
- . (2010) 'True Belief is Not Instrumentally Valuable', in Cory D. Wright and Nikolaj J.L.L. Pedersen (eds) *New Waves in Truth*, New York, NY: Palgrave Macmillan.