

ARTICLE

Against Universal Epistemic Instrumentalism

James Bernard Willoughby 

School of Philosophy, Australian National University, Canberra, Australia
Email: james.willoughby@anu.edu.au

Abstract

Beliefs should conform to some norms. Epistemic instrumentalism holds that your beliefs should conform to these epistemic norms just because conforming is useful. But there seems to be cases where conforming to the epistemic norms isn't useful at all, as in so-called "too-few-reasons" cases. In response to these cases, universal epistemic instrumentalists argue that despite first appearances, it is always useful to conform to the epistemic norms. I argue that all current versions of this universalist response are objectionable. I conclude with a conjecture about why no version of universalism will succeed.

Keywords: Epistemic normativity; epistemic instrumentalism; too-few-reasons objection; optimistic instrumentalism; epistemic norms; normativity; nonauthoritative reasons

1. Introduction

Conforming to the epistemic norms is remarkably useful. I hold that this fact can explain all the relevant phenomena associated with epistemic normativity. This makes me an epistemic instrumentalist. However, there is a certain kind of epistemic instrumentalism, 'universal epistemic instrumentalism,' that I find unacceptable. That I find it unacceptable is not so unusual. There is a long list of criticisms against this position. However, I am unsatisfied with existent criticisms. They fail to correctly identify the deep flaws in the universal position. I aim to rectify that.

The plan: in [section 2](#), I describe epistemic instrumentalism, universal epistemic instrumentalism, and what motivates some philosophers to adopt the universal position. In [sections 3, 4, 5, and 6](#), I describe various ways the universal position has been developed. In each of these sections, I identify the deep problems with each version. Finally in [section 7](#), I conjecture that no universal position will succeed.

2. Universal epistemic instrumentalism

Why conform to the epistemic norms? Unsurprisingly, philosophers have a host of answers to this question. 'Epistemic instrumentalism' is the thesis that you should conform to the epistemic norms because beliefs that do so promote the achievement of your goals or other valuable ends. Beliefs that conform are useful tools. There are some theoretical motivations to be an epistemic instrumentalist. Firstly, conforming to the epistemic norms is in fact a useful thing to do—at least in a lot of cases. So instrumentalism has a simple explanation about why we should conform to the epistemic norms and why we might learn to pay special attention to them. Secondly, this isn't just a theory about mere semantics, rather it is a metaphysical claim about the world (Kornblith 1993, 358–63). It explains why epistemic normativity is important, about why it is philosophically interesting, regardless of the content of our terms or concepts. Thirdly, instrumentalism is an avenue to

complete Quine's project of naturalising epistemology (Quine 1969). Instrumentalism, as it proposes a natural ontology, requires a much less suspicious metaphysics and epistemology than nonnatural rivals (for some metaphysical and epistemic concerns about normativity, see Joyce 2001; Mackie 1977; Street 2006). This is a serious advantage over those nonnatural rival theories (for rival theories, see Kelly 2003; Parfit 2011; Scanlon 1998; Wedgwood 2007 cited in Cowie 2014, 4004; also see Enoch 2011; Moore [1903] 1968; Scanlon 2014). And many philosophers have expressed sympathy for naturalising epistemology with instrumentalism (Dyke 2021; Foley 1987; Laudan 1990; Nozick 1993). The promise of a natural metaphysics is one of, if not the, central motivation for instrumentalism.

While instrumentalism has some things going for it, there is a traditional objection to it, the 'too-few-reasons objection'. The canonical statement of the too-few-reasons objection is by Thomas Kelly (2003; also see 2007). A similar argument is made by Harvey Siegel (1990; 2019). Charles Côté-Bouchard describes a case close to one of Kelly's (Côté-Bouchard 2015, 340):

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!'

On first pass, none of Max's goals are promoted by believing his colleague. But the epistemic norms demand that his colleague be believed. We can stipulate the colleague is an excellent testifier, having never led Max astray in the past. The problem for instrumentalists is that philosophers judge that Max has epistemic reasons to conform to the epistemic norms in this case.¹ But if none of Max's goals are promoted, there are no instrumental reasons to conform. So the instrumentalist has too few reasons to account for the epistemic reasons in this case. Let's call cases like this 'too-few-reasons cases.' Of course, there are many possible too-few-reasons cases. So there will be many cases where the instrumentalist can't find relevant instrumental reasons. This is the too-few-reasons objection.

The first thing to note: successful too-few-reasons cases are much more unusual than is commonly supposed. Finding relevant instrumental reasons in **Spoiler Alert** is not particularly hard. Supposing Max is anything like a normal human, he will not want to risk upsetting his colleague by not trusting his testimony. So Max has instrumental reasons to conform to the epistemic norms in **Spoiler Alert**. I suspect this isn't generally pointed out because philosophers can easily anticipate a response. We can modify the case and stipulate that Max doesn't care about other peoples' feelings. But even in this modified case, Max still has relevant instrumental reasons to conform. If Max is anything like a normal human, he will want to be seen as reasonable, as not weird. Flouting the epistemic norms risks causing Max's colleagues to see him as unreasonable, as weird. But suppose Max doesn't care about other peoples' opinions either. This modified too-few-reasons case is still not successful. If Max cares about getting a promotion or being invited to parties say, there are reasons not to risk others' seeing him as weird because it will lower the chances of him getting the promotion or the invitation. But suppose Max doesn't care about promotions or parties.... I am sure there is a successful too-few-reasons case. My point is that it will be a lot more unusual, a lot more unrealistic than **Spoiler Alert**. It will have to involve an agent that has a very impoverished set of goals or some other type of unrealistic element. Too-few-reasons cases aren't an everyday problem.

¹At least, let's say so for the sake of argument. Without this judgment, the motivation to be universal or noninstrumentalist is weakened. So universalists and noninstrumentalists should welcome it. In my experience, a lot of philosophers will make this judgment. But it would be worth checking experimentally just how uniform this judgment is among philosophers and nonphilosophers.

The strangeness of successful too-few-reasons cases is an issue for the objector (and the universalist; see section 7). Instrumentalism is a metaphysical theory. When the instrumentalist claims that “epistemic reasons are instrumental reasons,” they are claiming something about the world, not merely about the relationships between concepts. To use Derek Parfit’s taxonomy, instrumentalism can be interpreted as a form of nonanalytical naturalism (2011, 643). Think of the instrumentalist’s claim like you think about the claim “water is H₂O.” So the evidential relationship between the too-few-reasons objection—that is the judgments—and instrumentalism isn’t simply that there is a counter example to an analysis, like a Gettier case against a justified, true belief analysis of knowledge (Gettier 1963). It is something more complicated. And as these metaphysical judgments are about more and more unrealistic scenarios, it becomes harder to justify treating those judgments as true or reliable.²

Rather than question the evidential value of the judgments, ‘universal epistemic instrumentalists’ instead offer an extremely bold reply in response to too-few-reasons cases: in all possible cases, there are instrumental reasons to conform to the epistemic norms whenever we judge that there are epistemic reasons. This is the sense in which they are *universal* epistemic instrumentalists. The relevant instrumental reasons are everywhere. The fact they offer such a bold reply to such an esoteric set of cases is the first clue that something about this position is problematic. But the issues do not end there. In the next four sections, I will describe four different versions of the universal position and the serious problems with each version.

3. Cheap universal epistemic instrumentalism

3.1 Schroeder’s cheap reasons

The first theory I will consider was described by Mark Schroeder (2007, chap. 6). Consider the situation in **Spoiler Alert**. Even if Max doesn’t know of any goals that will be directly promoted by conforming, there is some chance that the correct information will be important later. Suppose later that day, Max’s crush wants to speak about the spoiler. Knowing about the spoiler will increase the chances that the conversation with his crush will go well. Suppose later that day, Max is listening to the radio and a competition comes on requiring knowledge of that spoiler. Knowledge of that spoiler may help Max win the competition. Who knows what the future will hold for Max?

Cheap universalists posit a glut of instrumental reasons. The central idea is that any possibility of promoting your goals gives rise to instrumental reasons and there is always *some* possibility that conforming to the epistemic norms will promote your goals. So instrumental reasons are very cheap. Cheap enough that, supposedly, there will always be relevant instrumental reasons in any too-few-reasons case.

I, however, doubt that there are cheap reasons in every possible case. We can modify **Spoiler Alert**, making it extremely unrealistic. Suppose Max only cares about working through the day and watching his show; he doesn’t care about anything else. Also, Max believes with certainty that an all-powerful god will intervene to ensure that conforming to or flouting the epistemic norms will not impact any of Max’s goals and we can stipulate that this belief is true. In this case, Max’s goals aren’t affected by not conforming to the epistemic norms. So there won’t even be cheap reasons to conform. But even if there is a response to this modified case, there is still a deeper flaw with the position.

²For arguments that odd cases make for poor evidence, see Machery (2017), Van Inwagen (1998), and Kilov and Hendy (2022). Whether you agree with the arguments or not, all I am pointing out here is that the epistemic basis for justifying judgments from esoteric cases is controversial at best, and the instrumentalist is not compelled to adopt that controversial epistemology, so it is questionable whether the cases create a strong objection in the first place. I will argue in section 7 that even more than the objector, the universalist needs the objection to carry significant epistemic weight.

3.2 *The otiose objection*

My objection has two parts: (1) positing cheap reasons creates a lacuna and (2) filling the lacuna makes cheap reasons otiose. Let's call this the 'otiose objection.'

3.2.1 *Cheap reasons create a lacuna*

Philosophers make too-few-reasons judgments—that is, they judge there are reasons to conform to the epistemic norms even when there are no apparent instrumental reasons to conform. Cheap universalism vindicates these judgments by claiming that there are cheap epistemic reasons in those cases, so philosophers are judging correctly.

But cheap reasons aren't rare. Almost every norm in almost every situation has *some* possibility of promoting your goals. For example, as you are reading this, there is some possibility that flouting the epistemic norms will promote your goals; let's call the associated reasons 'cheap anti-epistemic reasons.' We don't make too-few-reasons judgments about cheap anti-epistemic reasons. Indeed, we don't make too-few-reasons judgments in the vast majority of cases where we have some cheap reasons (Schroeder 2007, chap. 5).³ So by the lights of cheap universalism, some of our judgments are mistaken. Perhaps we judge correctly that there are reasons in too-few-reasons cases. But we judge incorrectly that we don't have reasons in cases where we have all sorts of cheap reasons, such as cheap anti-epistemic reasons. So cheap universalism only vindicates our too-few-reasons judgments by positing a massive error in our judgments about when reasons do not exist. If you were motivated to adopt universalism to vindicate judgments, to "save the appearances" as it is sometimes called, this is a problem.

Moreover, and more importantly for my objection, there is also a gap in the cheap universalist's story that needs explanation. If we don't normally judge cheap reasons to exist but we do normally judge cheap epistemic reasons to exist, what is special about cheap epistemic reasons? This lacuna is a deep flaw in the position.

3.2.2 *Filling the lacuna makes cheap reasons otiose*

To fill this lacuna, cheap universalists need to explain why we make the relevant judgments in too-few-reasons cases. But as those judgments were the motivation to posit cheap reasons in the first place, explaining why we make those judgments with some other feature of reality leaves nothing left for cheap reasons to explain. Filling the lacuna makes cheap reasons otiose.

For example, Schroeder—responding to a slightly different problem—tried to fill the lacuna (2007, chaps. 5, 6, 7). His explanation is extremely complicated. The gist is that reasons we have in virtue of a task itself are more important than any old cheap reasons. We have cheap epistemic reasons in virtue of believing alone. The reasons to conform to the epistemic norms are special in this way, so we pay extra attention to them. But the same is not true for our cheap anti-epistemic reasons. Those reasons we do not have in virtue of believing alone. So we make too-few-reasons judgments about cheap epistemic reasons but not cheap anti-epistemic reasons because the former are reasons we have in virtue of believing alone.

I reject this explanation. But suppose it is right. If we have reasons to conform to the epistemic norms in virtue of believing alone, then we don't need to posit cheap reasons. In too-few-reasons cases, as people are believing, they should already have relevant instrumental reasons to conform to

³Again, let's grant this to the universalist and the objector. If philosophers do make too-few-reasons judgments in response to other cases of cheap reasons, such as cheap anti-epistemic reasons, this would be evidence that too-few-reasons judgments aren't particularly important for epistemic normativity. If we made too-few-reasons judgments about all cheap reasons, it would make as much sense to be a universalist about the cheap reasons to negligently swing a running chainsaw around your head as it would to be a universalist about epistemic reasons. In this case, universalism would not be an important part of the theory of epistemic normativity.

the epistemic norms. Whether there are additional cheap reasons is beside the point. Schroeder has undercut the original motivation for positing cheap reasons by explaining too-few-reasons judgments with these in-virtue-of-the-task-itself special reasons.

This isn't a coincidence. If cheap universalists explain why we treat cheap epistemic reasons differently to other cheap reasons, then that explanation also responds to too-few-reasons cases. This explanation can't depend on epistemic reasons being cheap because we don't track all kinds of other cheap reasons, such as cheap anti-epistemic reasons. So a successful explanation of why we treat cheap epistemic reasons differently makes positing cheap reasons redundant. Cheap reasons don't actually move the debate forward at all because they require an independent response to too-few-reasons cases. So cheap reasons are otiose.

I have found this point difficult to communicate so it is worth summarising the argument. This is the otiose objection:

P1: Cheap reasons are supposed to reply to the too-few-reasons objection.

P2: Positing cheap reasons incurs an explanatory debt: Why do we track cheap epistemic reasons but not other cheap reasons, such as cheap anti-epistemic reasons?

P3: A successful explanation can't depend on the cheapness of epistemic reasons because both types of reasons are equally cheap.

P4: A successful explanation of why we track cheap epistemic reasons is also a response to the too-few-reasons objection.

Therefore,

C1: Cheap reasons are otiose. They are supposed to reply to the too-few-reasons objection, but they require an explanation that constitutes an independent reply to the too-few-reasons objection.

Therefore,

C2: Cheap universal epistemic instrumentalism isn't the correct theory of epistemic reasons.

This is a serious problem for cheap universalism, one that can't be resolved by rejecting a case, denying a judgment or finagling the theory. The too-few-reasons objection is a challenge to explain why we treat epistemic reasons differently than other instrumental reasons. The answer is not that epistemic reasons can be cheap.

4. Rule universal epistemic instrumentalism

4.1 Rule-based reasons

Note that up until now we have been putting forward a version of universalism analogous to act-utilitarianism. We have been looking for a way that conforming to the epistemic norms promotes at least one goal in every case. But we could offer a version of universalism analogous to rule-utilitarianism. Adam Leite offers this type of theory (2007).⁴

Why conform to the epistemic norms? Leite offers this answer (for him, conforming to the epistemic norms means believing in proportion with your evidence):

... it serves the overall satisfaction of one's goals (cognitive and perhaps otherwise) to accept a system of norms according to which believing on good evidence is epistemically laudable even if one lacks good instrumental reasons for so believing in the particular case. (463)

⁴Leite isn't himself an instrumentalist. However, Hilary Kornblith endorses a similar view, with epistemic normativity being based on what makes for a good epistemic policy, rather than on what achieves your specific goals in any particular case (Kornblith 1993). So even if Leite doesn't endorse the view, other philosophers do.

Rule universalists hold that we should adopt a policy of conforming to the epistemic norms, even if in particular cases conforming to the epistemic norms doesn't promote any specific goals. Considering **Spoiler Alert**, suppose Max doesn't have any traditional instrumental reasons to conform. The rule universalist maintains that Max would still have other instrumental reasons to conform to the epistemic norms as a rule. Supposedly, maintaining that rule will promote Max's future goals. So, supposedly, the rule provides instrumental reasons for Max, and instrumental reasons in all possible too-few-reasons cases.

I also doubt rule instrumentalism can provide relevant instrumental reasons in all possible too-few-reasons cases. We can imagine contrarian humans, or aliens, if you prefer, whose only goal is to not conform to the epistemic norms at any time or as a rule. It is hard to see how these beings have relevant instrumental reasons to conform. But even if there is a response to such a case, there is a deeper flaw in this position.

4.2 *The collapse dilemma*

Rule universalism is like rule consequentialism; it is open to similar objections.⁵ Brad Hooker outlines an objection to rule consequentialism (2016, sec. 8). Rule consequentialism either collapses into act consequentialism or it is incoherent. A rule either maximises the good or it doesn't. If the rule maximises the good, then act consequentialism also holds that you should follow the rule, in which case act consequentialism and rule consequentialism are equivalent in an important sense. If following a rule doesn't maximise the good, then the rule consequentialist seems committed to not maximising the good, which, at least, appears incoherent.

There is a similar, although not exactly the same, issue with rule universalism. Suppose you are faced with a choice about whether to conform to the epistemic norms. If there are instrumental reasons to conform to the epistemic norms, then rule universalism and 'act universalism' (let's call the alternative) are equivalent. They both tell you to conform. If we do well by following a rule, then that is what act universalism tells us to do. If there are no instrumental reasons to conform to the epistemic norms, then rule universalism is advocating that we conform without any instrumental reasons. While this may not be strictly incoherent, there is something odd about it.

Indeed, we can put a finer point on the 'collapse dilemma' for rule universalism. The first horn: rule universalism collapses into act universalism. Suppose there are instrumental reasons for following a rule. The act universalist can respond to those reasons as much as the rule universalist can. So we haven't gained any ground by shifting to rule universalism. The second horn: rule universalism isn't an instrumentalist theory. Suppose there are no instrumental reasons to conform as a rule. If there are reasons to conform in this case, they are not instrumental reasons. This implies epistemic instrumentalism is false. So, at least on first pass, adopting either horn isn't a viable gambit for universal instrumentalists.

4.2.1 *Avoiding the collapse: Realistic rule universalism*

Hooker offers a response on behalf of rule consequentialism (2016, sec. 8). His strategy is to include some pragmatic considerations when deciding on the correct rules. So, for example, when assessing the best rules, we should consider how hard they are to follow, to learn, and to enforce. These considerations limit the complexity of rules that we can have. We can't have a rule that says: "maximise the good" because it would be very hard to know what maximises the good in every situation and it would be very hard to enforce. Rules that are too specific will also be a problem. Rules like "whenever it costs you almost nothing, help an eighty-three-year-old woman named Mabel across the street" might be easier to know when to apply, but rarely come up. So, according to Hooker, we should have rules that look more familiar. Rules like, "Don't lie," "Be nice to people,"

⁵For an introduction to rule consequentialism, see Hooker (2016).

“Be charitable,” etc. This list of rules would be longer than “maximise the good” but would be more enforceable and easier to apply. So the best rules will look more like the ten commandments rather than a general imperative to maximise the good.

Most importantly, for Hooker, once we accept the rules, the right action to take in any particular case is to conform with the rules. So, if the rules say “Don’t lie,” then the right action to take in any situation is not to lie.⁶ Note that following the rule will not be the same as maximising the good. Perhaps there is a situation where lying could maximise the good. But, according to Hooker, rule consequentialists hold there is some reason to conform to a rule even when it doesn’t maximise the good.

You might be tempted to make a similar response on behalf of rule universalists. Rule universalism doesn’t collapse into act universalism. There will be some occasions where there are no instrumental reasons to conform to the epistemic norms. Rule universalism will still direct you to conform to the epistemic norms because conforming to the norms is what makes a belief right.

No instrumentalist should be satisfied with Hooker’s response. Most importantly, Hooker’s response to the collapse dilemma distinguishes rule consequentialism from act consequentialism by giving up on consequentialism. It holds that conforming to a rule gives rise to right action. When applied to rule universalism, the analogous claim is that conforming to the epistemic norms gives rise to right belief regardless of whether there are instrumental reasons to conform or not. This response amounts to giving up instrumentalism. Of course, if there are instrumental reasons, act universalists can respond to those. But if there aren’t any instrumental reasons, then rule universalism is no longer instrumentalist. While Hooker’s response avoids the charge of inconsistency (which is how he originally described the second horn of the collapse dilemma), a Hooker-style response runs head-first into the second horn of the collapse dilemma as I describe it for rule universalists.

4.2.2 *Accepting the collapse: The otiose objection, part deux*

Rather than trying to avoid the collapse dilemma, the universalist might just accept that rule universalism collapses into act universalism. If you think of rule universalism as a concession to too-few-reasons objectors, then accepting the collapse is a problem. If the rule universalist concedes that there are cases where there are no instrumental reasons to conform to the epistemic norms, then the fact that rule universalism is a type of act universalism makes the concession incoherent. Act universalism just is the claim that there are always some instrumental reasons to conform to the epistemic norms.

But rather than a concession, you might be tempted to view rule universalism as a change of emphasis. Traditionally, act universalism focuses on how a belief that conforms to the epistemic norms can promote your goals. You can think of rule universalism as making a similar point but focusing on how conforming in particular cases can help cultivate good habits. Rule universalists are worried that each deviation from the epistemic norms is a risk that you will deviate at some other point. And they are hopeful that each time you conform to the epistemic norms you increase the chances you will conform in the future. Indeed, in this way rule universalism complements act universalism. Sometimes conforming matters because the belief is important. Sometimes conforming matters because the conforming is important. Sometimes both are important. Moreover, this move forces too-few-reasons cases to be rather extreme to rule out all relevant instrumental reasons. A successful too-few-reasons case must rule out any possibility that conforming in one case will

⁶There is a complication here. You might wonder what happens if not lying would lead to the deaths of one thousand people or even one person? We can put in a caveat at the end of all rules: ... unless doing so leads to catastrophe. Of course, this caveat creates more problems for rule consequentialism, problems I am not sure I can solve let alone solve in this essay. So let’s put aside catastrophes for this discussion.

affect the chances of conforming in the next case. These too-few-reasons cases will not look very realistic. This is a promising way to construe rule universalism.

But universalists aren't satisfied with pointing out that successful too-few-reasons cases are esoteric. They could have done that well before invoking rule-based reasons. Rather they want to claim that *in all possible* cases there are instrumental reasons to conform. This change-of-emphasis rule universalism inherits all the problems of act universalism, in particular the otiose objection. Considering a modified **Spoiler Alert**, imagine that Max is about to die and that he has good evidence that he is about to die. All he wants is to finish his day at work (say his day is ending in a minute) and go home to watch his show, spoiler free (say his home is right next door to his work). Does he have instrumental reasons to believe his spoiler-uttering colleague? Rule universalism says yes. Conforming in this case increases the chances of conforming in the next case. But Max is about to die. It is extremely unlikely that there are going to be any important cases after this. So the probability that conforming will actually matter is extremely low. At the same time, there are many equally cheap reasons in this case, even other rule-based reasons. Perhaps blowing raspberries in this case will help in the next case where blowing raspberries is the right action to take. But we don't make the too-few-reasons judgments in response to reasons to blow raspberries. So the rule universalist, who considers their theory to be a version of act universalism, has the same explanatory lacuna as cheap universalism. We need some explanation about why we treat rule-based epistemic reasons so differently than other rule-based reasons. The fact the reasons are rule-based can't be the basis of the explanation because there are plenty of equally cheap rule-based reasons we don't offer the too-few reasons judgment about. So we can offer the full otiose objection as presented in section 3.2 against this version of rule universalism, too.

Therefore, rule universalism either isn't a form of instrumentalism or is open to the otiose objection. It is not a step forward for universal instrumentalists despite looking so promising.

5. Nonauthoritative universal instrumentalism

5.1 Nonauthoritative reasons

Recently there has been a lot of interest in a position about epistemic normativity that distinguishes between two types of reasons, 'authoritative' and 'nonauthoritative reasons'⁷ (Maguire and Woods 2020; Cowie 2019, cited in Murphy 2021; Côté-Bouchard and Littlejohn 2018). None of the philosophers who defend this view defend it as a version of instrumentalism, so nothing I say here is necessarily a criticism of the actual view. Nonetheless, we can entertain this view as a version of universal instrumentalism. Let's call this hypothetical position 'Nonauthoritative universal instrumentalism.' In this section, I articulate two criticisms. First, the central distinction at the heart of this position isn't compatible with universal instrumentalism, so this hypothetical position is in tension with itself. Second, this distinction doesn't provide a good explanation of too-few-reasons judgments.

But before I make my criticism, let me introduce the view. As I said above, nonauthoritative universal instrumentalism relies on a distinction between different types of reasons: authoritative and nonauthoritative reasons. To identify the difference between these types of reasons Barry Maguire and Jack Woods point to the different functions of these reasons (Maguire and Woods 2020, 225–30). Authoritative reasons contribute to what you ought to do, all things considered. For example, "that it will cause your imminent death" is an authoritative reason against jumping off a tall building; everything else being equal, it tells against jumping off a tall building. Nonauthoritative

⁷Maguire and Woods use the terms "authoritatively normative" and "nonauthoritatively normative" to refer to different types of reasons (2020, 217). Côté-Bouchard and Littlejohn use the terms "thinly normative" and "robustly normative" to refer to the same distinction (Côté-Bouchard and Littlejohn 2018, 154). Other terms that are used to point to one side of the distinction are "genuine," "substantive," and "generic" (Maguire and Woods 2020, n13).

reasons contribute to when something is correct according to a standard. For example, it might be that castling early in a game of chess is correct from the standard of winning games of chess, but it might not be what you ought to do from an all-things-considered point of view; if you were teaching the game to someone, for instance, perhaps you shouldn't introduce castling until they understand how the pieces typically move.

This distinction is supposed to offer a convincing reply to too-few-reasons cases (Maguire and Woods 2020, 239–41). In these cases, there are nonauthoritative reasons that make it epistemically correct to believe in certain ways even if there are no authoritative reasons that make it all-things-considered rational to believe in that way. So, in **Spoiler Alert**, it is epistemically correct for Max to believe the spoiler, but he lacks authoritative reasons to believe it, so all-things-considered, he ought not believe the spoiler. According to nonauthoritative universal instrumentalism, there are still reasons in all too-few-reasons cases; they just lack authority. So we can, supposedly, vindicate the too-few-reasons judgment while preserving the core of an instrumentalist position, that instrumental reasons are doing the normative heavy lifting.

What gives rise to relevant standards? There are many possibilities. Here are two. First, it might be that humans can collectively or individually declare standards. This is implied by Barry and Woods (2020, 217–18). They hold that we have nonauthoritative reasons to play games in certain ways, like moving the king one space at a time during chess. This implies that we can generate nonauthoritative reasons as easily as we can create games. So in a similar way that children can declare the standard of a novel game on the playground, humans may be able to declare standards and thereby create nonauthoritative reasons. Second, it might be that the relevant standards were created by evolution (Nolfi 2018). Perhaps beliefs have a proper function. A 'proper function' is the function that beliefs were selected for when they evolved. The relevant standards for belief, a proponent of the view might hold, is whatever fulfils that proper function. In this way, evolution may have bestowed upon us the relevant epistemic standards, and these standards generate nonauthoritative reasons—the reasons we are tracking when we make too-few-reasons judgments.

Whatever story you tell about how the standards are generated, there is one constraint. It can't be that the standards, and therefore the nonauthoritative reasons, are generated by usefulness alone. Why? Because the distinction is supposed to provide extra explanatory resources that the traditional instrumentalist doesn't have; the one thing you can't claim is that it is useful to conform to the epistemic standards in genuine too-few-reasons cases. So the nonauthoritative reasons have to come from somewhere else. This constraint creates a problem for nonauthoritative universal instrumentalism.

5.2 *Nonauthoritative universal instrumentalism is in tension with itself*

Nonauthoritative universal instrumentalism is either not universal or not instrumentalist. Either way, the position is in tension with itself. Are "nonauthoritative reasons" reasons? Answering either yes or no is incompatible with universal epistemic instrumentalism.

Suppose the nonauthoritative instrumentalist answers: "No, nonauthoritative reasons are not reasons." This implies that even though there are nonauthoritative reasons in too-few-reasons cases, there are no reasons, instrumental or otherwise. The nonauthoritative reasons are supposed to vindicate too-few-reasons judgments; they are supposed to be the things we are tracking with our judgments. So if the things we are tracking are not reasons, then the universal instrumentalist is wrong; it is not true that there are always instrumental reasons in these cases. Sometimes there are *only* nonauthoritative reasons, which, on this horn of the dilemma, aren't reasons, instrumental or otherwise. So answering no is incompatible with the universal-part of universal instrumentalism because it suggests the relevant reasons are not universal.

Suppose the nonauthoritative instrumentalist answers: "Yes, nonauthoritative reasons are reasons." This implies that the reasons in too-few-reasons cases aren't instrumental reasons. Recall that the only constraint on the origin of nonauthoritative reasons is that they don't come from being

instrumentally useful. That constraint creates a problem on this horn. Sometimes in too-few reasons cases there are *only* nonauthoritative reasons, which, on this horn of the dilemma, are not instrumental reasons. So answering yes is incompatible with the instrumental-part of universal instrumentalism because it suggests the relevant reasons aren't instrumental reasons. This dilemma shows that the hypothetical position 'nonauthoritative universal instrumentalism' is in tension with itself. Either the position isn't universal or it isn't instrumentalist, so in neither case is it universal instrumentalism.

Indeed, this dilemma is a version of the collapse dilemma I offered against rule universalism. The nonauthoritative distinction doesn't introduce any new resources to explain the short fall of instrumental reasons that too-few-reasons cases point to. So either it will collapse into the other instrumental views, where there are no relevant reasons in too-few-reasons cases, or it draws upon noninstrumental reasons, making it not an instrumentalist position. Making a distinction does not help you find the extra instrumental reasons universalists need.

5.3 *The distinction doesn't help with too-few-reasons cases*

In section 3.2, I argued that positing cheap reasons incurs an explanatory debt that makes cheap reasons otiose. Something similar happens in the case of nonauthoritative reasons. However, there are important differences between the philosophers who advocate for cheap reasons (Schroeder 2007) and those who advocate for nonauthoritative reasons (Maguire and Woods 2020; Cowie 2019, cited in Murphy 2021; Côté-Bouchard and Littlejohn 2018). Cheap reasons were posited to explicitly handle too-few-reasons cases (Schroeder 2007, chap. 6), whereas nonauthoritative reasons have other motivations and only incidentally are supposed to respond to too-few-reasons cases.⁸ So, unlike cheap reasons, even if nonauthoritative reasons fail to produce a response to the too-few-reasons objection, they may not be completely otiose; they may play other important theoretical roles.

Nonetheless, nonauthoritative reasons fail to produce a convincing response to the too-few-reasons objection. The response is that in the relevant cases there are still reasons, just not authoritative ones. So our judgments are correct. But as with cheap reasons, nonauthoritative reasons aren't rare and this leaves the position open to the otiose objection.

Consider games. We may have nonauthoritative reasons to play games in certain ways, like moving the king one space at a time during chess (Maguire and Woods 2020, 217–18). But I take it as common ground that there are enough instrumental reasons to account for the reasons to move the king one space—that is, we don't think there are too-few-reasons for the normativity of games, like chess and Twister.⁹ Indeed, if we can arbitrarily produce standards, say, according to the Made-Up Standard, it is correct that you should now be hopping on one leg; then nonauthoritative reasons can be extremely common.

But as I said above, perhaps the epistemic standards, and therefore the epistemic nonauthoritative reasons, are produced in some other way—by evolution or some other mechanism. Nonetheless, I can offer a schema for generating the same problem.

1. The first step is to locate some standard produced by that same mechanism that we don't make too-few-reasons judgments about. In the case of a proper function, we might point to the proper function of my hand, for example. Suppose the proper function of my hand is, at least in part, to dextrously manipulate objects with my fingers.

⁸For example, Maguire and Woods are primarily concerned with providing a response to objections based on the idea that you can't weigh epistemic reasons with other instrumental reasons (2020, 214). So it is possible that while nonauthoritative reasons fail to produce a convincing response to the too-few-reasons objection, they might provide a convincing response to these other objections. Although it also is possible that the otiose objection generalises.

⁹Again, let's grant the universalist and the objector (see n1, n3).

2. The second step is to generate a too-few-reasons-like case—a case where we have no instrumental reasons to achieve the relevant standard. For our proper function example, imagine I need to fall asleep to live and dextrously manipulating objects with my fingers would interfere with me falling asleep. Even though according to the standard generated by the proper function I have nonauthoritative reasons to dextrously manipulate objects, I have no instrumental reasons to do that.
3. The third step is to note that we don't make too-few-reasons judgments about the too-few-reasons-like case. For our example, I don't need to posit extra reasons in this case for me to dextrously manipulate objects with my fingers. Indeed, I bet that for most standards, once it is no longer useful for me to achieve that standard, in general, people judge that there are no relevant reasons to achieve the standard. There are exceptions, such as ethical and epistemic standards, but these are very rare, or at least I conjecture.
4. Then we note that we are making a mistake somewhere. If we are tracking the nonauthoritative reasons in canonical too-few-reasons cases by making too-few-reasons judgments, then we are making a mistake in not making similar judgments about nonepistemic standards in too-few-reasons-like cases. One way or another, we are treating the standards generated by whatever mechanism generates the relevant standards differently. Again, part of the challenge that too-few-reasons judgments offer is to explain what is special about epistemic norms. And that there are nonauthoritative reasons to conform to them isn't what makes them special; if there are any nonauthoritative reasons then there are heaps of them, many that we regularly judge not to exist.

Strictly speaking, this schema leaves open a possible response. Nonauthoritative instrumentalists could identify a mechanism that generates epistemic standards but no other standards we fail to make too-few-reasons judgments about. I doubt it is possible to make such a response plausible because such a response is heavily constrained.

Firstly, we have a pretty good idea that these standards evolved, either in a biological process as Sharon Street argues (2006; 2009) or perhaps in some cultural process. As the example above suggests, evolution doesn't only produce standards that we make too-few-reasons judgments about. So we need an additional mechanism to make this response. But whatever additional mechanism is proposed, it has to fit in with the evolutionary story about how we actually learned the standards.

Secondly, we also need a story about why humans are tracking standards produced by that mechanism. There are heaps of standards. And there are heaps of mechanisms by which standards are produced. There needs to be some explanation about why the particular mechanism is special, such that we developed too-few-reasons judgments about the standards it produced.

Thirdly, supposing there are beliefs, the mechanism needs to be extremely modally robust. Why? As I noted in section 2, successful too-few-reasons cases are already extremely unrealistic. So to the extent that you are using those unrealistic cases to motivate universal instrumentalism, you are already aiming to develop a theory that produces epistemic reasons in all (or perhaps almost all) cases where agents believe. Otherwise, you are not a universal instrumentalist. If the mechanisms or standards are not modally robust (given the existence of beliefs) then there will be too-few-reasons cases with no nonauthoritative reasons. If standards don't robustly exist given beliefs, we can imagine a human on a far planet or some protohuman in the deep past or some engineered human in our future, who has beliefs but does not have the relevant environment to produce epistemic standards. Nonauthoritative universalism will have to reject the idea that there are epistemic reasons for those agents in these cases, even though they are structurally similar in all other respects to too-few-reasons cases. Therefore the mechanism needs to exist robustly, so that when we imagine a protohuman in a too-few-reasons case (etc.), the relevant standards already exist.

So the possible response available to nonauthoritative instrumentalists requires some mechanism that plausibly generates our epistemic standards that is plausibly integrated with an evolutionary story about how we learnt about those standards, that is special in some respect so humans

pay attention to the standards produced by that mechanism (but that nevertheless doesn't give rise to authoritative reasons), that exists robustly across modal space and robustly produces those standards (given the existence of beliefs), and that only produces standards we make too-few-reasons judgments about. This possible response is very constrained, so constrained that I doubt there is a plausible story to tell that checks all the boxes.

Setting aside such a response, this argumentative schema is a problem for anyone that holds that the nonauthoritative distinction somehow vindicates our too-few-reasons judgments. Perhaps we are right to judge that there are reasons in too-few-reasons cases because we accurately track the nonauthoritative reasons, but then we are wrong to judge that there are no reasons to conform to the other standards generated by the relevant mechanism. One way or another, we can't save all the appearances.

Moreover, again setting aside the possible response to my argumentative schema, once we make the distinction between authoritative and nonauthoritative reasons, we still need an explanation about why philosophers treat epistemic standards differently to other standards. Whatever that explanation is, it can't depend on nonauthoritative reasons because those reasons are too common. For example, suppose the relevant difference is that conforming to the epistemic standards is almost always useful.¹⁰ Most of the time, we have authoritative reasons to conform so it is natural to expect that we pay special attention to the epistemic standards. But this just is the nonuniversalist instrumentalist response to too-few-reasons cases (Willoughby 2022; Sharadin 2018). I agree the overwhelming usefulness explains why we would make too-few-reasons judgments, an efficient if mistaken overgeneralisation. But then there is nothing left for the distinction to explain in too-few-reasons cases. I grant that the distinction makes too-few-reasons judgments true (at the cost of implying anytime we don't make too-few-reasons judgments in response to too-few-reasons-like cases we are mistaken). But the explanation of why we make the judgment in response to epistemic norms but not in response to other standards has nothing to do with tracking the nonauthoritative reasons. Indeed, this story implies that a fortunate coincidence took place. The reason we treat epistemic norms as special is because of the overwhelming (but not universal) usefulness of conforming. But the too-few-reasons judgment is made true by the presence of unrelated non-authoritative reasons. This is like Bertrand Russell's case of the broken clock fortunately telling you the right time (Russell [1948] 2009). Sure, your belief is true, but you just got lucky. It is not really vindicating the judgment. Indeed, positing that this fortunate coincidence took place is a cost for a theory.

This problem is a version of the otiose objection. Too-few-reasons judgments tell you that something is special about epistemic normativity. The purported explanans of the judgments, nonauthoritative reasons, don't tell you what is special about epistemic normativity. And an additional explanation about what is special, such as the overwhelming usefulness of conforming to epistemic norms, makes the purported explanans otiose, at least when it comes to responding to the too-few-reasons objection.¹¹

6. Common goal universal instrumentalism

6.1 Common goal reasons

So far, the versions of universalism we have been considering hold that there are instrumental reasons to conform to the epistemic norms, regardless of the goals you have. Let's change tack. Now

¹⁰This was suggested by an anonymous reviewer.

¹¹I remain uncommitted about whether you can combine nonuniversal instrumentalism to explain too-few-reasons cases with the nonauthoritative distinction to explain other cases. Although, as I argue above, this combination implies that a fortunate coincidence took place where our judgments were caused by something (usefulness) unrelated to what makes the judgments true (the presence of nonauthoritative reasons). So at a minimum, this combination of theories carries this cost.

we will consider if there is some universal goal that is promoted by conforming to the epistemic norms. If we all have such a goal, then we will always have instrumental reasons to conform to the epistemic norms.

Richard Foley describes the simplest version of common goal instrumentalism—that we all have the goal of conforming to the epistemic norms (for him, the epistemic norms are now believing the truth and now not believing a falsehood):

Might not [some person] lack the goal of now believing truths and now not believing falsehoods? Perhaps not. Perhaps this is *impossible*, at least if the notion of a goal is understood in a suitably broad sense. After all, *the vast majority* of us are intrinsically curious about the world; we intrinsically want to have true beliefs. (1987, 11; my emphasis)

If we all had that goal, then we would all have instrumental reasons to conform to the epistemic norms.

Indeed, we can make the required sense of having a goal quite minimal. Ernest Sosa holds it is at least plausible that “for any arbitrary belief of ours, we would prefer that it be true rather than not true, other things equal” (2003, 157). This minimal claim would be enough to establish common goal universalism. So long as you also held that satisfying preferences gave rise to instrumental reasons, then there would always be instrumental reasons to conform to the epistemic norms because conforming always promotes the achievement of one of your preferences.

This mere preference view has an analogous lacuna to cheap universalism, so it is open to a version of the otiose objection too. However plausible it is to claim we all have a mere preference to conform, it is as plausible to claim that all other things equal, we prefer to win a game of chess rather than lose. But chess reasons and epistemic reasons are treated very differently. So, the mere preference view still needs that additional explanation that cheap universalism needed. What is the difference between epistemic reasons and other instrumental reasons? And the explanation can’t depend on there being a mere preference for conforming because it is as plausible to say we have a mere preference for many goals, but we don’t judge the resultant reasons to exist the way we judge epistemic reasons to exist. A mere preference doesn’t explain the difference between epistemic reasons and other instrumental reasons.

That being said, it is not likely that we *all* have a mere preference for conforming. Humans exhibit a massive amount of variation (Hull 1986). As Nathaniel Sharadin argues (2018, 3794–96), common goal universalism posits an unlikely uniformity in nature.¹² For universalists, it won’t do if the “vast majority” of people intrinsically value truth or prefer truth. Universalism requires that every single person prefers truth. And it is hard to not believe that there are some avant-garde libertines, students of Nietzsche or just psychologically damaged people, who are utterly indifferent to truth other things equal. In general, humans are varied. Not all of us have two eyes. Not all of us want to keep on living. Not all of us like our children. There are even some humans who don’t enjoy philosophy! So we need incredibly strong evidence to believe a specific universal uniformity in humans, especially when the supposed uniformity is merely a contingent fact about our psychology.

6.2 *Côté-Bouchard’s common goal of acting for genuine reasons*

Côté-Bouchard employs a different strategy for establishing a common goal (2015). The strategy is to argue that agents necessarily have a goal that is promoted by conforming to the epistemic norms. This is a reasonable strategy to try because the element of necessity is exactly what the above versions of universalism lacked.

Rather than analyse Côté-Bouchard’s argument, I will just grant the central claim. Let’s grant that “agents” constitutively have the goal to “act for genuine reasons” and satisfying that goal often

¹²Sharadin calls common-goal instrumentalism “the special interests strategy.”

requires the agent to conform to the epistemic norms (Côté-Bouchard 2015, 350–51). Of course, conforming does not always promote that goal. We can imagine that by trying to conform, we end up acting for only apparent reasons because the evidence was misleading. In response, Côté-Bouchard claims that *as a rule* conforming to the epistemic norms promotes the goal of acting for genuine reasons (351). As I have argued above in section 4.2, rule universalism faces the collapse dilemma. That is reason enough to reject this position.

But there is a new problem with this position too. No constitutive claim could make common goal universalism plausible. The constitutive claim is supposed to be evidence that almost all humans in fact have the requisite goal. Côté-Bouchard's argument is supposed to proceed: almost all humans are agents; (constitutively) agents desire to act for genuine reasons; so almost all humans have a desire to act for genuine reasons. However, this is a roundabout way to argue that almost all humans have a goal. If a lot of humans did have this goal, there would be a lot more evidence of this fact, such as a lot of people declaring that they have such a desire or that they are working towards it, or that something thwarted their attempts to get it, or that attaining it is just around the corner, just like they do with widespread goals to be healthy and wealthy. I suspect that most humans have the goal to continue breathing. When the goal to continue breathing is thwarted, even temporarily, it is newsworthy; many people have said something to the effect, "I almost choked on that chip!" Indeed, I take it the goal to act for genuine reasons is supposed to be even more widespread than the goals to be healthy and wealthy. But almost no one mentions that they have this goal or mentions anything about this goal. So even if Côté-Bouchard is right about the constitutive claim, given the lack of evidence that people by and large have a common goal, we should hold that people by and large are just not "agents" in Côté-Bouchard's sense. So this version of common goal universalism has the same serious problem as all types of common goal universalism: humans don't share a common goal!

7. Universalism's Future

Universalism hasn't been successful. Will it succeed in the future? I conjecture that it will not. The space for a successful universal theory is very narrow, perhaps nonexistent. Moreover, it starts with a very unlikely claim and it does not have much evidence to support that claim. So I doubt there will ever be a successful universal instrumental theory.

We analysed four different versions of the universal position. But all these theories are trying to address the too-few-reasons objection with the same broad strategy. The strategy is to find a source of previously unaccounted for instrumental reasons. Let's call this strategy "extra-reasons universalism." This strategy has an incredibly narrow window of success.

Firstly, in general, instrumentalists don't lack reasons. The fact that successful too-few-reasons cases must be so unusual implies that, in everyday cases, there are instrumental reasons to conform to the epistemic norms. The problem is that whether something is useful is a contingent fact. So there will be some possibilities where it is not useful to conform. You don't change this by describing an additional source of contingent reasons. Perhaps the possibilities will become more fanciful, involving gods, agents with few goals, or aliens but they will still be possible. You need *extra necessary* reasons, not *extra contingent* reasons.

Secondly, to be an instrumentalist theory, the extra reasons need to be instrumental. But as we identify extra reasons originating from sources other than usefulness, such as reasons generated by rules or by standards, the reasons tend to lose their instrumentalist character. And if the extra reasons aren't instrumental reasons, then the position is no longer instrumentalist and it loses the theoretical advantages of that position. A problem for extra-reasons universalism is that there isn't some vast untapped reservoir of instrumental reasons. This is why the collapse dilemma often manifests. By pushing to identify extra reasons, universalists end up positing extra noninstrumental reasons or fail to posit any extra reasons at all.

Thirdly, the extra reasons can't be ubiquitous. If the extra reasons are everywhere, then you can't explain why we make the too-few-reasons judgment about epistemic reasons specifically. This is what gives rise to the otiose objection. And indeed, positing ubiquitous extra-reasons implies we are making a massive error by not having the too-few-reasons judgment more generally. So if your universalism posits extra reasons that are ubiquitous, then your position requires an independent response to the too-few reasons objection, making the extra reasons otiose.

So a successful extra-reasons universalism posits some type of extra necessary reasons, that are instrumental in character, that aren't ubiquitous. That is a tall order. So the question that now faces us: Is it possible to thread the needle by identifying some relevant extra necessary nonubiquitous instrumental reasons, or is there some other strategy that might be successful? I doubt it.

Universalism makes this claim: conforming to the epistemic norms is *always* useful. There is a lot of inductive evidence against this (Willoughby 2022, 19). For almost every set of norms,¹³ conforming to them is not always useful. Conforming to the rules of Twister doesn't help plan a typical funeral. For every tool, using it is not always useful. Wielding a knife doesn't help lead a relaxing meditation. And because the epistemic norms are not obviously different to other norms and tools, we should set our credence very low that conforming to the epistemic norms is *always* useful.

The main, perhaps only, source of evidence for universalism is our too-few-reasons judgments. But you must be extraordinarily optimistic about the reliability of our judgments to consider them especially evidentially weighty. The content of these judgments might be semantic,¹⁴ in which case you must believe that the content of our concepts is an incredibly accurate guide to metaphysical reality, even in very unusual cases, to overcome the initial implausibility of the universalists claim. Alternatively, the content of these judgments might be metaphysical,¹⁵ in which case you need to believe our metaphysical insight is incredibly accurate, even in very unusual cases, to overcome the initial implausibility of the universalist's claim. Finally, the content of the judgments might be some third thing, in which case you must believe this hitherto unspecified faculty is incredibly accurate, even in very unusual cases, to overcome the initial implausibility of the universalist's claim. Without believing these judgments are especially evidentially weighty, a belief that I bet isn't particularly widespread, the universalist is stuck with an implausible claim and little evidence to move the debate forwards.

In this way, the universalist depends on too-few-reasons judgments even more than the objector. The objector is arguing that epistemic normativity is something other than instrumental normativity because the epistemic norms always give rise to epistemic reasons, but conforming isn't always useful. It is not so unlikely that epistemic normativity is something other than instrumental normativity and it is quite likely that conforming to the norms isn't always useful. So the objector doesn't need the judgments to be that evidentially weighty. On the other hand, the universalist needs the judgments to tell us that conforming to the epistemic norms is *always* useful—a wildly unlikely claim. So I doubt universal epistemic instrumentalism will ever be successful.

Acknowledgments. Many thanks to Noga Gratvol, Alan Hájek, Thomas Kelly, Mathew Kopec, Hilary Kornblith, Toby Charles Penhallurick Solomon, Katie Steele, Jeremy Strasser, Shang Long Yeo, Preston J. Werner, and Timothy L. Williamson. I also thank Joan Keyes for editorial services.

James Bernard Willoughby is a PhD candidate at the Australian National University. His thesis offers an instrumental account of epistemic normativity, a topic he has published on previously. While he rarely encounters a philosophical topic that doesn't interest him, his current research interests include epistemology, metaphilosophy, experimental philosophy, and philosophy of law.

¹³Why not all? ... I anticipate some gerrymandered set of norms like: "do only what is useful."

¹⁴Broadly speaking, this is consistent with the method of cases presented by Frank Jackson (1998).

¹⁵Broadly speaking, this is consistent with the method of cases presented by Timothy Williamson (2007).

References

- Côté-Bouchard, Charles. 2015. "Epistemic Instrumentalism and the Too-few-reasons Objection." *Philosophical Studies* 23 (3): 337–55. <https://doi.org/10.1080/09672559.2015.1042007>.
- Côté-Bouchard, Charles, and Clayton Littlejohn. 2018. "Knowledge, Reasons, and Errors about Error Theory." In *Metaepistemology*, edited by Christos Kyriacou and Robin McKenna: Cham: Springer International.
- Cowie, Christopher. 2014. "In Defence of Instrumentalism about Epistemic Normativity." *Synthese* 191 (16): 4003–17. <https://doi.org/10.1007/s11229-014-0510-6>.
- Cowie, Christopher. 2019. *Morality and Epistemic Judgement: the argument from analogy*. New York: Oxford University Press.
- Dyke, Michelle M. 2021. "Could Our Epistemic Reasons Be Collective Practical Reasons?" *Nous* 55: 842–62. <https://doi.org/10.1111/nous.12335>.
- Enoch, D. (2011). *Taking Morality Seriously: A Defense of Robust Realism*. Oxford: Oxford University Press.
- Foley, Richard. 1987. *The Theory of Epistemic Rationality*. Cambridge, MA: Harvard University Press.
- Gettier, Edmund. 1963. "Is Justified True Belief Knowledge?" *Analysis* 23 (6): 121–23.
- Hooker, Brad. 2016. "Rule Consequentialism." In *The Stanford Encyclopedia of Philosophy* (Winter), edited by Edward N. Zalta. <https://plato.stanford.edu/archives/win2016/entries/consequentialism-rule/>.
- Hull, David L. 1986. "On Human Nature." *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association* 1986: 3–13.
- Jackson, Frank. 1998. *From Metaphysics to Ethics: A Defence of Conceptual Analysis*. Oxford: Oxford University Press.
- Joyce, Richard. 2001. *The Myth of Morality*. Vol. 452: Cambridge: Cambridge University Press.
- Kelly, Thomas. 2003. "Epistemic Rationality as Instrumental Rationality: A Critique." *Philosophy and Phenomenological Research* 66 (3): 612–40. <https://doi.org/10.2307/20140564>.
- Kelly, Thomas. 2007. "Evidence and Normativity: Reply to Leite." *Philosophy and Phenomenological Research* 75 (2): 465–74. <https://doi.org/10.2307/40041118>.
- Kilov, Daniel, and Caroline Hendy. 2022. "Pundits and Possibilities: Philosophers Are Not Modal Experts." *Australasian Journal of Philosophy*: 1–20.
- Kornblith, Hilary. 1993. "Epistemic Normativity." *Synthese* 94 (3): 357–76. <https://doi.org/10.1007/bf01064485>.
- Laudan, Larry. 1990. "Aim-less Epistemology?" *Studies in History and Philosophy of Science Part A* 21 (2): 315–22. [https://doi.org/10.1016/0039-3681\(90\)90029-8](https://doi.org/10.1016/0039-3681(90)90029-8).
- 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): 456–64. <https://doi.org/10.2307/40041117>.
- Machery, Edouard. 2017. *Philosophy within Its Proper Bounds*. Oxford: Oxford University Press.
- Mackie, John L. 1977. *Ethics: Inventing Right and Wrong*. Hammondsworth, UK: Penguin.
- Maguire, Barry, and Jack Woods. 2020. "The Game of Belief." *The Philosophical Review* 129 (2): 211–49. <https://doi.org/10.1215/00318108-8012843>.
- Moore, G. E. (1903) 1968. *Principia Ethica*. Cambridge: Cambridge University Press.
- Murphy, Alex. 2021. "Chris Cowie, Morality and Epistemic Judgement, (OUP), 2019." *Ethical Theory and Moral Practice* 24 (2): 629–32. <https://doi.org/10.1007/s10677-021-10172-w>.
- Nolfi, Kate. 2018. "Why Only Evidential Considerations Can Justify Belief." In *Normativity: Epistemic and Practical*, edited by Conor McHugh, Jonathan Way, and Daniel Whiting, 179–99. Oxford: Oxford University Press.
- Nozick, Robert. 1993. *The Nature of Rationality*. Princeton, NJ: Princeton University Press.
- Parfit, Derek. 2011. *On What Matters*. 3 vols. Oxford: Oxford University Press.
- Quine, W. V. 1969. "Epistemology Naturalized." In *Ontological Relativity & Other Essays*, 69–90. New York: Columbia University Press.
- Russell, Bertrand. (1948) 2009. *Human Knowledge: Its Scope and Limits*. London: Routledge.
- Scanlon, Thomas M. 1998. *What We Owe to Each Other*. Cambridge, MA: Belknap Press.
- Scanlon, Thomas M. 2014. *Being Realistic about Reasons*. Oxford: Oxford University Press.
- Schroeder, Mark Andrew. 2007. *Slaves of the Passions*. Vol. 3. Oxford: Oxford University Press.
- Sharadin, Nathaniel. 2018. "Epistemic Instrumentalism and the Reason to believe in Accord with the Evidence." *Synthese* 195 (9): 3791–809. <https://doi.org/10.1007/s11229-016-1245-3>.
- Siegel, Harvey. 1990. "Laudan's Normative Naturalism." *Studies in History and Philosophy of Science Part A* 21 (2): 295–313. [http://doi.org/10.1016/0039-3681\(90\)90028-7](http://doi.org/10.1016/0039-3681(90)90028-7).
- Siegel, Harvey. 2019. "Epistemic Rationality." *Metaphilosophy* 50 (5): 608–30. <https://doi.org/10.1111/meta.12392>.
- Sosa, Ernest. 2003. "The Place of Truth in Epistemology." In *Intellectual Virtue: Perspectives from Ethics and Epistemology*, edited by Michael DePaul and Linda Zagzebski. New York: Oxford University Press.
- Street, Sharon. 2006. "A Darwinian Dilemma for Realist Theories of Value." *Philosophical Studies* 127 (1): 109–66.
- Street, Sharon. 2009. "Evolution and the Normativity of Epistemic Reasons." *Canadian Journal of Philosophy* 39: 213–48. <https://doi.org/10.1080/00455091.2009.10717649>.

- Van Inwagen, Peter. 1998. "Modal Epistemology." *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition* 92 (1/2): 67–84. <http://doi.org/10.1023/a:1017159501073>.
- Wedgwood, Ralph. 2007. *The Nature of Normativity*. Oxford: Oxford University Press.
- Williamson, Timothy. 2007. *The Philosophy of Philosophy*. Vol. 3. Malden, MA: Blackwell.
- Willoughby, James Bernard. 2022. "Embedded Epistemic Instrumentalism: An Account of Epistemic Normativity." *Philosophers' Imprint* 22 (7). <https://journals.publishing.umich.edu/phimp/article/id/745/>.