

Embedded Epistemic Instrumentalism: An Account of Epistemic Normativity

James Bernard Willoughby

Australian National University

© 2022 James Bernard Willoughby
*This work is licensed under a Creative Commons
Attribution-NonCommercial-NoDerivatives 3.0 License.*
<www.philosophersimprint.org/007022/>

1. Introduction

Why conform to the epistemic norms?

Conforming to the epistemic norms is a useful undertaking. It is useful in many ways. The fact that it is useful explains the normativity of epistemic norms and epistemic reasons.¹ This is embedded epistemic instrumentalism in a nutshell. The rest of the paper explains, expands, and defends this position.

The theory that usefulness explains epistemic normativity is called **epistemic instrumentalism**. Epistemic instrumentalism is familiar (Côté-Bouchard, 2015; Dyke, 2020; Foley, 1987, 1993; Giere, 1989; Kitcher, 1992; Kornblith, 1993; Laudan, 1986, 1990a, 1990b; Leite, 2007; Nozick, 1993; Quine, 1969; Schroeder, 2007)² but not popular. We can lay at least part of the blame for this lack of popularity on the standard way epistemic instrumentalism is developed. Standard developments have two flaws: first, they assert that the epistemic norms exert normative force always and everywhere, that epistemic norms have universal application. This universal epistemic instrumentalism is not very plausible. Second, they focus on the way the epistemic norms can promote a particular end, the end chosen varying with the account. One proposal is that people have a universal end that is always promoted by conforming to the epistemic norms (Côté-Bouchard, 2015; Foley, 1987). Another proposal is that any arbitrarily chosen end will always be promoted by conforming to the epistemic norms (Kornblith, 1993; Schroeder, 2007). However, the strongest form of instrumentalism does not focus on a supposed universal end, or an arbitrarily chosen end, or even any particular way such an end is promoted by conforming to the epistemic norms. Rather, the strongest form of instrumentalism focuses on the plurality of ends and the plurality of ways conforming to the epistemic norms promotes those ends.

1. In this paper, I talk about 'epistemic reasons'. I do not mean to imply these are the fundamental normative terms. My view could be articulated equally well with other normative terminology. I use reason-talk because it is standard in the relevant literature.
2. The details of my account differ, sometimes dramatically, from previous instrumental accounts.

The development of epistemic instrumentalism was heavily influenced by Thomas Kelly (2003), offering what came to be known as the too-few-reasons objection.³ The gist of the objection is that philosophers judge that the epistemic norms have normative force in response to all proposed cases, even in response to cases where conforming to the epistemic norms doesn't appear to promote any goals. Philosophers judge as if the normative force of epistemic norms is universal.⁴ However, according to epistemic instrumentalism, epistemic norms only have force when a person has a relevant end. At least on first pass, ends are not universal. Therefore, instrumentalists can't say the epistemic norms have universal force, implying that the theory cannot be correct. In response, instrumentalists try to shoehorn instrumental reasons into playing a universal role, claiming that, one way or another, epistemic instrumentalism *does* deliver universal normativity. This is the wrong response and yet it is this response that has shaped the standard development of epistemic instrumentalism. My account is much simpler and offers a much more compelling response to the too-few-reasons objection — or so I will argue.

There are many theoretical motivations for holding an instrumental account. As Michelle M. Dyke asserts, instrumental accounts avoid the host of difficulties associated with *sui generis* normative accounts (Dyke, 2020, p. 2). Epistemic instrumentalism does not require any dubious metaphysics or epistemology, nor does the theory require fundamental normative properties or any epistemic miracles (c.f. Enoch, 2011, pp. 172–173). As Kelly argues, these accounts conform with the loose collection of motivations dubbed “naturalism” (Kelly, 2003, pp. 614–616). Instrumentalism coheres with scientific practice in that

3. Also see: Kelly (2007). For a similar argument see Siegel (1989, 2019). The ‘too-few-reasons’ objection was coined by Schroeder (2007, chapter 5) and followed by Côté-Bouchard (2015). Sharadin calls it the ‘Universality Challenge’ (Sharadin, 2018, p. 3793)
4. In my experience, philosophers do judge in this way. However, it would be worth checking experimentally just how uniform this judgment is. For this paper, I am happy to concede that philosophers judge this way. If they don't have too few reasons judgments, all the better for my theory.

everything it postulates can be studied scientifically. As Hilary Kornblith explains, these accounts are not just theories engaging in mere semantics, rather they make substantial claims about the world (Kornblith, 1993, pp. 359–363). Regardless of what you call it, the usefulness of conforming to the epistemic norms is an important fact about the world that should be studied. This isn't something that only matters because we happen to refer to it with the term ‘epistemic normativity’. If you are a philosopher who is motivated by such things, then this style of theory will appeal to you. Even ignoring the theoretical motivations, my version of instrumentalism is based on a few uncontroversial facts and is able to explain what it needs to explain. As such, I hold it is a compelling answer to the question: why conform to the epistemic norms?

An obstacle to answering this question is that there are many potential candidates for **epistemic norms**: believe in proportion with your evidence;⁵ believe what is supported by your other beliefs;⁶ have beliefs that are formed by a reliable belief-forming process;⁷ believe what can be derived from a secure foundation;⁸ believe only what you know.⁹ You might hold that only one of the above is the correct norm, or that some combination of them is correct, or another formulation entirely. In this paper, I remain agnostic. My position will work with any plausible account of the epistemic norms.

The plan: Section 2 will describe how conforming to the epistemic norms is useful; Section 3 will describe how this usefulness explains the normative force of epistemic norms; Section 4 will describe and

5. As an evidentialist holds, e.g., Conee & Feldman (2004). Indeed, Kate Nolfi (2018, p. 180) claims that almost all epistemologists accept something resembling this thesis.
6. As a coherentist holds, e.g., Quine & Ullian (1970).
7. As a reliabilist holds, e.g., Goldman (1967).
8. As a foundationalist holds, e.g., Rene Descartes' ‘Meditations on First Philosophy’ in Cottingham, Stoothoff, & Murdoch (1985, pp. 1–62).
9. As knowledge-first advocates hold, e.g., Williamson (2000). Note: the answers are not mutually exclusive. There may be interesting combinations of them. For a mix of foundationalism and coherentism, see Elgin (2005).

respond to the too-few-reasons objection; and Section 5 will articulate some advantages of this account over its competitors.

2. Conforming is Useful

In this section, I describe certain ways that conforming to the epistemic norms is useful. These descriptions are influenced by previous work on epistemic instrumentalism. While previous instrumental theories tend to focus on one specific way that conforming is useful (attempting to respond to the too-few-reasons objection), my theory holds that conforming is useful in many related, though somewhat distinct, ways. In a sense, embedded instrumentalism is a disjunctive theory.

First, conforming promotes having accurate representations of the world. Accurate representations help with almost everything you do in life. Want to write? It helps to know the locations of the relevant computer keys. Want to eat? It helps to have true beliefs about what is nutritious and, more importantly, what tastes great. Want to get a promotion? It helps to have accurate credences about the chances that currying favor with your boss will pay off. Conforming helps you accurately represent the world. Accurate representations help you do almost everything you want to do. This is the basic way conforming is useful. Broadly speaking, you could subsume all the different ways conforming is useful described below in this basic way. As such, what follows is not a series of distinct ways that conforming is useful, but rather a series of varying points of emphasis, demonstrating the broad usefulness of conforming. It is a mistake to only focus on one, missing the forest for the trees.

Second, conforming is useful because there is a chance that accuracy will matter in the future. Mark Schroeder bases his theory on this fact (2007, chapter 6).¹⁰ Odds are you don't currently care what the state beverage of Delaware is, but at some point in the future it might be a question on a quiz show, with the correct answer winning you a million dollars. There is always a chance that accuracy will matter.

10. For criticism, see Côté-Bouchard (2015, pp. 343–345), Sharadin (2018, pp. 3796–3801), and Willoughby (2022, section 3).

So rather than directly helping now, the chance that accurate beliefs will help in the future provides instrumental reasons to conform. To build on Schroeder's point: typically, the cost of conforming is small. In many instances, having beliefs that conform takes less effort than having non-conforming beliefs (try to believe you are not reading now, for example). Furthermore, we are indifferent to the content of many beliefs, so we may as well conform. This results in a wide variety of cases where you ought to conform because conforming might help in the future and the costs of conforming are negligible.

Third, conforming now can help you conform in the future. Hilary Kornblith and Adam Leite develop instrumental theories based on the idea that conforming to the epistemic norms as a rule is important (Kornblith, 1993; Leite, 2007).¹¹ A virtue of these accounts is that they remind us that it can be important to conform not because conforming is useful right now, nor because accuracy might matter in the future, but because being *able* to conform is useful. If you are in the habit of flouting the norms, you may not be able to conform when the stakes are high. As such, it is better to practice when nothing else is at stake so when it matters, you are able to conform.

Fourth, you have goals that require conforming. Some authors have suggested that all people share some common goal. Charles Côté-Bouchard argues that everyone necessarily has the goal of acting on the basis of genuine reasons and this requires that you conform as a rule (2015, pp. 350–352).¹² Richard Foley asserts that people may simply have a preference for conforming (1987, p. 11).¹³ While I doubt the universality of any single goal, there are some reasonably common goals that demand conformance. For example, many people want others to trust their judgment, and they aspire to deserve that trust. Such a goal requires people to conform to the epistemic norms.

11. For criticism, see Côté-Bouchard (2015, pp. 348–450), Lockard (2013, pp. 1712–1716), and Willoughby (2022, section 4).

12. For criticism, see Sharadin (2018, pp. 3794–3796) and Willoughby (2022, section 6).

13. Also see Sosa (2003, p. 157).

If people regularly flout the epistemic norms, then they are not worthy of being trusted. Indeed, the general usefulness of conforming and the fact that many common goals require it may be developmentally connected. Conforming is so useful that we can imagine natural evolutionary processes, be they biological or cultural, furnishing us with an appetite for conforming the way it furnished us with an appetite for calories. So while I doubt the universality of Côté-Bouchard's and Foley's claims, there are some common goals that require conforming to the epistemic norms.

Fifth, many social elements of our world require conformance or systematic deception. As the level of deception required is beyond most people (or they simply do not want to be deceptive), they are required to conform to achieve their social goals. For example, many people want to act in a way that makes those around them comfortable. The more you flout epistemic norms, the more people will feel uncomfortable. Imagine someone refusing to believe that there was a cup in front of them when it is clearly visible, or placing their hand in a fire because they don't believe it will hurt this time, etc. Perhaps there is some genius liar who could flout the norms while convincingly pretending otherwise, but for most of us, the only means we have to make those around us comfortable is to conform. I concede that there are some subjects such as politics, religion, or public scandals, where one can ignore epistemic norms with seemingly little repercussion, but start denying there are cars on the road and your life will change very quickly. Similarly, people don't want to be patronized; they don't want to be ignored; they want to be taken seriously. These social goals often require conforming to the epistemic norms.

Sixth, some people are part of groups with epistemic goals that require conformance. Michelle M. Dyke holds that the fact people are part of communities that have epistemic goals is the basis of epistemic normativity (Dyke, 2020). While I reject the theory overall, I agree that communities have goals and that some people want to further those goals. For example, I am part of the philosophical community. This community is trying to understand the world. I want to help. This

requires me to conform to the epistemic norms so that I can communicate and convince other members of the community, to help make progress. There are many communities like this: businesses, not-for-profit organizations, governments, etc. Many people want to help their communities, and conforming is required to achieve some community goals.

In summary, there are at least six related ways conforming is useful. I suspect there are more, but I take it these six ways are relatively uncontroversial. Moreover, they provide instrumental reasons in a lot of cases, so even if these are the only ways conforming is useful, they are enough to support my theory. Importantly, we begin to see one way that the standard developments of epistemic instrumentalism are deficient. By focusing only on one way conforming is useful, they miss the plurality of other ways. We can utilize this plurality to deliver a more compelling version of instrumentalism.

3. The Embedded Goal Model

I have described some ways conforming is useful. But this is not much of a theory, as it is not clear how the plurality of usefulness explains epistemic normativity. In particular, on first pass, instrumental reasons seem unlike epistemic reasons. Typically, we think of instrumental reasons as fickle. Sometimes, you have instrumental reasons to go for a walk. It might, right now, help achieve some explicit goal or valuable end. But in our complex environment, whether you have reasons to go for a walk will vary over time. Right now, you might be stressed and have reasons to walk, while later tonight, when you are trying to fall asleep, you won't have those reasons. This fickleness is normal for instrumental reasons, but it is a problem for an instrumental account of epistemic normativity. Typically, we think of epistemic reasons as robust. If you see a cat on a mat, then you have epistemic reasons to believe there is a cat on the mat, regardless of how tired you are, whether you are stressed or, at least on first pass, what your goals are. Typical instrumental reasons are much more sensitive to context than typical epistemic reasons. A challenge for the instrumentalist is to use

a theory that typically produces fickle reasons (instrumentalism) to account for a phenomenon that produces robust reasons (epistemic norms).

In this section, I will present a version of instrumentalism that does just that. I will present my theory as a version of pragmatic instrumentalism, concerned with achieving your explicit goals. However, I will have something to say about how to modify this account for other versions of instrumentalism as well. The central move is to posit a type of intermediary goal — a goal that will help you achieve a number of other goals (or valuable ends), whatever they might be. This intermediary goal is relatively stable, so even as your other goals (or valuable ends) change, the intermediary goal stays the same. This intermediary goal gives rise to robust epistemic reasons.

My account offers a different explanation for the stability of epistemic reasons to the standard version. Standard accounts aim to describe some universal relation between conforming to the epistemic norms and promoting at least one of your goals. The universality of the relation explains the robustness of epistemic reasons. In contrast, my account focuses on the aggregate. We are shifting from a particle physics-like focus to a statistical mechanics-like one. The robustness of epistemic reasons does not come from specific interactions between instrumental reasons and goals, but from the robustness of the aggregate of many such interactions. While the difference is subtle, I am offering an alternative explanation of epistemic normativity to the standard account. I am claiming that the robustness of epistemic normativity comes from the forest, not the trees.

To introduce the embedded goal model of epistemic normativity, I employ an extended analogy. The instrumentalist can easily account for less mysterious normative phenomena, like the reasons you have to be healthy. So initially, I will build an instrumental model of health normativity. Then, I will apply this model to the case of epistemic normativity.

3.1 *The Embedded Health Goal Model*

You should eat vegetables. You should exercise. You should relax. You should do these things because they are healthy. You have health reasons to do them. If you do not do these things, you increase your chances of having health problems — which is bad. It is not bad in some all-things-considered way, or some mysterious abstract way. Rather, being healthy is crucial for achieving most of the goals you have in life.

You might have explicit health goals: think of my-body-is-a-temple people. Those people place great importance on being healthy for its own sake. Being unhealthy hinders the achievement of their explicit health goals. This is one reason it is bad to be unhealthy.

However, being unhealthy is bad regardless of whether you have explicit health goals or not. You have all sorts of other goals that require, or are greatly promoted by, being healthy. You may want to achieve big things in a career. Being healthy helps. You may want a loving relationship. Being healthy helps there, too. Maybe you just want to bowl ripping leg-spinners or dunk a basketball or drive a golf ball or run long distances or dance the lambada. Yet again, being healthy helps. You may or may not have the explicit goal, but it is still in your interest to be healthy. Let's call being healthy **an embedded goal**, because achieving it promotes a variety of other goals.

Embedded goals give rise to more robust reasons than their explicit counterparts. Explicit goals depend on facts about human psychology, which is rather fickle. For example, when I get home from work, I sometimes want to watch TV, while at other times I want to go for a walk. The goal of going for a walk creates reasons to do other things, such as reasons to put on sunscreen (so I don't get sunburnt). These derivative reasons are fickle. They only exist while I have the goal of going for a walk, and whether I have that goal or not is incredibly variable. Of course, it is a matter of degree, but reasons based on explicit goals typically inherit the underlying fickleness of human psychology. Contrast this with embedded goals, which depend on more stable facts: facts about what it takes to achieve many different goals.

For example, even when I have no desire to go for a walk, I still have reasons to go for a walk based on the embedded health goal. Walking is good for my health and my health is good for achieving many different things. Reasons based on embedded goals are less fickle than their explicit counterparts.

Embeddedness comes in degrees. Consider the following two types of embedded goals. Most modern people have the embedded goal of earning money. It is useful for achieving any number of other goals. However, we can imagine some prehistoric woman, striking stones together to make a knife. She has no reasons to earn money because the institutions are not in place that would make money useful. However, she has a closely associated embedded goal, the goal of gathering resources. Indeed, almost any creature that has goals needs to gather resources to achieve those goals. So “gathering resources” is a more embedded goal than “making money” because it exists in all contexts where the goal “making money” exists, and additional contexts as well.

A goal can have a high degree of embeddedness, but still not give rise to reasons always and everywhere. Consider the well-embedded health goal. Suppose someone whose every moment of life brings them extraordinary pain and their only goal is to stop the pain as soon as possible. Achieving this goal will necessarily require their sudden death. They choose to decrease their pain accordingly. This is an unhealthy choice, but it is the only reasonable choice they have. Even the well-embedded health goal does not give us reasons to be healthy always and everywhere.

Nevertheless, notice how we typically treat “being healthy” as if it is good always and everywhere. For example, “being healthy” is usually taken as sufficient reason to explain an action. We might ask: “Why are you running or eating that or relaxing?” The reply “it is healthy” is usually taken as a sufficient reason. Similarly, we usually consider everyone to be bound by the normative power of being healthy. You should eat your vegetables, and you should exercise, and you should relax. We do not usually add the caveat “...if you want to be healthy”. Lastly, we usually consider it a great harm to deprive children of the

ability to be healthy, even if we thwart their preferences. Like all sensible humans, some children want to eat chocolate all day. However, letting a child fulfil that preference is inappropriate, in the absence of an extenuating circumstance. And if a child does not want to eat any vegetables, under usual circumstances we would coerce the child into eating them. In general, we can treat promoting the goal of “being healthy” as good always and everywhere, and we will rarely encounter any issues because it is such a well-embedded goal.

An important point of contrast between explicit goals and embedded goals is possible complexity.¹⁴ To understand this contrast, imagine a child playing his first game of chess against an adult. Suppose that both the child and the adult have the goal of winning the game. As a matter of fact, in the early game, you should develop your pieces, castle, and control the center spaces. The adult understands this, and has the explicit goal “develop my pieces, castle, and control the center spaces”. The child isn’t even aware that you can castle in chess, so cannot have this explicit goal. Nevertheless, the child still has the embedded goal “develop my pieces, castle, and control the center spaces” because embedded goals are based on facts about what will help the child win in a wide variety of contexts, not on the cognitive limitations of the ignorant child’s mind. So embedded goals can be extraordinarily complex. They can be based on information a person doesn’t possess. They can even be unimaginable. Embedded goals can be much more complex than their explicit counterparts.

Returning to our case study, consider the complexity of the goal “being healthy”. My best guess is that “being healthy” amounts to some complex combination of having a long- and well-functioning body and mind. This complexity might pose a problem for formulating an explicit health goal because explicit health goals must be cognitively tractable. But there is no complexity constraint on embedded health

14. This possible complexity avoids a problem that plagues some other accounts of epistemic instrumentalism: see Section 5. It also allows for responses to other objections. For the objections, see Berker (2013) and Friedman (2018). However, responding to these other objections is left for another time.

goals. Depending on how the facts fall, the health goal might be a very complex goal indeed. It might be disjunctive. It might be inconceivable. It might even be pluralistic. This last possibility deserves some attention. The embedded health goal might be pluralistic like so: “achieve longevity to the extent that it doesn’t overly reduce functionality and achieve functionality to the extent that it does not overly reduce longevity.” Such a goal could be pluralistic in the sense that, from a health perspective, there is no single correct way to trade off longevity and functionality; there is no single notion of “overly reduce”. Indeed, extreme cases aside, when faced with a choice between longevity and functionality, it is hard to say what is the healthier choice for a third party to make. Therefore, the embedded health goal might have this messy complexity.

To summarize the Embedded Health Goal Model: there are explicit health goals and there is an embedded health goal. The embedded health goal is more robust because it depends on facts concerning what contributes to us achieving our other goals, as opposed to facts concerning our psychology. We have a degreed notion of embeddedness: the more embedded a goal, the more there are contexts where it provides reasons. However, just because a goal has a high degree of embeddedness does not imply that it provides reasons always and everywhere; there will be cases where a highly embedded goal provides no reasons. Finally, embedded goals can be very complex because their complexity is not constrained by our cognitive limitations. This means that people can have goals that are highly disjunctive, inconceivable, or even pluralistic.

3.2 *The Embedded Epistemic Goal Model*

Returning to epistemic instrumentalism, I am going to argue that people have an embedded **epistemic goal** — a goal to conform to the epistemic norms. Some people might have an explicit epistemic goal, the equivalent of the my-body-is-a-temple people. Philosophers may be a prime candidate, I suppose. These epistemic champions explicitly aim to have beliefs that conform to epistemic norms. However, as

suggested by our Embedded Health Goal Model, explicit epistemic goals aren’t required to account for the robust and widespread normative force of the epistemic norms.

Almost all of us have a variety of goals. As I described above, achieving the epistemic goal helps. For example, many people want to be well respected members of their community. It is hard to be respected if your testimony is false most of the time. We have short-term mundane goals, like getting coffee and food. These require us to locate coffee and food. We have long-term goals concerning things that will happen many years in the future, often requiring us to accurately anticipate how people will react to various changes to their circumstances. And we will have any number of goals in the future — ones we have not even considered yet. These will likely require us to have beliefs that conform to the epistemic norms, too. So the epistemic goal is an embedded goal, intimately connected to the vast array of other goals we find ourselves possessing.

The epistemic goal has a very high degree of embeddedness. Even if we imagine aliens living on the planet Seti-546 — if they have a variety of goals, and they use beliefs to guide their behavior as we do — then they are going to have the embedded epistemic goal. Suppose there are animals without goals *per se*, like coral. Even then, they will have something like goals to eat and to reproduce, so even they will have something like the embedded epistemic goal, too. It is in something like their interest to have accurate representations of the world. The epistemic goal is among the most embedded goals we have.

The well-embedded epistemic goal provides a robust basis for epistemic normativity. Unlike relying on explicit goals, the facts that give rise to the embedded goal are consistent across time. Unlike most explicit goals, the embedded goal can be very complex, very disjunctive, inconceivable, and pluralistic. We haven’t settled what the epistemic norms are and how they trade off against one another. Nonetheless, we still have the embedded epistemic goal to conform to the epistemic norms.

Importantly, the embedded model isn't committed to any particular source of normativity. I have presented the embedded model paired with pragmatic normativity, ultimately concerned with achieving explicit goals. However, the embedded model pairs with most sources of normativity, be it pragmatic, moral, aesthetic, or some bespoke source. For example, the embedded model of epistemic normativity is compatible with holding that only morality generates genuine normativity. Even if morality alone generates genuine normativity — to achieve what morality demands of you, to treat people with respect, or to maximize the good — you are required to achieve a wide variety of goals, more than enough variety to embed the epistemic goal (c.f. Clifford, 1877 [1999]). If a source of normativity gives rise to a rich set of ends, then it is compatible with the embedded model of epistemic normativity.¹⁵

Recall our central question: why believe in accordance with the epistemic norms? My answer: you have the embedded goal of conforming to the epistemic norms; achieving that embedded goal promotes a wide variety of other goals or valuable ends you have. Today, you may want to be a respectable epistemic peer. Tomorrow you may want to discover a vaccine for a pandemic. Conforming helps achieve those things. But the specific individual goals or valuable ends you have do not really matter. More importantly, you have a plurality of goals or ends and conforming promotes them in a plurality of ways. This creates a robust normative pressure that permeates our lives. This is the normative grip that the epistemic norms exert on us.

4. The Too-Few-Reasons Objection

So what's the problem? For epistemic instrumentalism, epistemic normativity requires conforming to the epistemic norms to be useful. According to the embedded epistemic goal model, while conforming is widely useful, it is not useful always and everywhere. In cases where

conforming isn't useful, you have no epistemic reasons to conform. These are known as **too-few-reasons cases**. For many, this is a deal breaker. They hold that the epistemic norms exert normative force always and everywhere. So we need, they suppose, some further fact about the epistemic norms — something beyond usefulness — to explain epistemic normativity in these cases.

Before examining a specific too-few-reasons case, I will point out a weakness with the objectors' position. That conforming is widely useful is uncontroversial (even if there are questions about how wide and how useful). As such, in a great many cases, non-instrumentalists are positing a glut of reasons. They say that in addition to the instrumental reasons, there are some non-instrumental reasons, too. The extra source of normativity might be a form of some *sui generis* normativity,¹⁶ some essential property of beliefs, the result of a proper function (Nolfi, 2018), or one of many other possibilities. But this extra source of normativity is posited in addition to the usefulness of conforming because the general usefulness is hard to deny. The problem is that usefulness alone is enough to explain why we would develop epistemic norms. It is enough to explain why we would pay attention to them, why we would track who conforms and who doesn't, the social esteem generated by conforming, and why we would spend considerable resources on conforming. So we have an uncontroversial fact that explains most of the relevant phenomena associated with epistemic normativity. Non-instrumentalists need a powerful motivation for postulating some additional source of normativity that in the overwhelming majority of cases is otiose.

The non-instrumentalist's position is a little too convenient. If you hold that there is a value, independent of usefulness, in conforming to the epistemic norms, then you are suggesting an extremely fortunate coincidence took place. We learned to value conforming because conforming is so useful, but we also happened to find that conforming

15. Strictly speaking, the embedded goal model is compatible with there being *sui generis* epistemic normativity. But the explanatory power of the model tells against this possibility (see Section 4).

16. There are many philosophers who make such claims. Cowie (2014, p. 4004) cites Kelly (2003), Parfit (2011), Scanlon (1998), and Wedgwood (2007). The classic *sui generis* view of normativity is in Moore (1903).

is valuable in an unrelated way. This is suspiciously convenient. As such, making the too-few-reasons objection, suggesting that we need a non-instrumental theory of epistemic normativity, commits one to controversial claims from the outset. The too-few-reasons objection isn't an objection from a neutral position. That doesn't mean, however, that instrumentalists should ignore it. Too-few-reasons cases are supposed to provide the extra motivation required for positing an additional source of epistemic normativity. They are commonly cited as the motivation for rejecting instrumentalism. So the too-few-reasons objection is worth considering.

Côté-Bouchard (2015, p.340) describes a too-few-reasons case, paraphrasing Kelly (2003):

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!'

In Spoiler Alert, supposedly, Max has no goals that are promoted by conforming to the norms and believing the testimony of his colleague. Therefore, he has no instrumental reasons to conform to the norms. And yet, philosophers judge¹⁷ that he still has epistemic reasons to conform. So instrumentalism alone cannot provide a complete account of epistemic normativity. This is a typical too-few-reasons case. It is hard to underestimate the influence these cases, and the objection based on them, has had on the literature.

In response, this type of everyday too-few-reasons case isn't a problem for embedded epistemic instrumentalism. If Max is anything like a

17. Again, let us suppose. See n. 4.

typical person, he has a wide range of goals that are promoted by conforming to the epistemic norms. For example, if he were to not believe his colleague, then his colleague would feel very awkward. Imagine a conversation in which Max maintains that he doesn't believe character X died. His colleague would feel distrusted, perhaps feel bad about themselves, or perhaps feel anger towards Max. The colleague might even lose respect for Max. For a lot of people, they don't want their colleagues to feel distrusted or bad about themselves or angry. Most of us care greatly about how much we are respected. And this is just one among many social goals that are promoted by conforming. If Max has the usual complex web of goals, there are plenty of instrumental reasons to conform. In an everyday case, Max has the embedded epistemic goal.

The above response is available to every instrumentalist. Once the embedded model is described, it is easy to see that Max will have many relevant goals. But nothing in this response actually depends on the embedded model. And yet this response isn't commonly offered. I suspect that philosophers can anticipate the next epicycle of the dialectic, idealized too-few-reasons cases. Suppose Max no longer has the normal range of goals, but rather only one: to believe character X didn't die. Place Idealized Max in the same circumstances, that is in the circumstances where all the epistemic norms point to the fact that character X died. Philosophers still judge that Idealized Max has epistemic reasons. But this Idealized Max doesn't have any goals that are promoted by conforming. So I can't claim Max has relevant instrumental reasons.

The standard response is to propose universal epistemic instrumentalism (Côté-Bouchard, 2015; Kornblith, 1993; Leite, 2007; Schroeder, 2007). Supposedly, some way or another, conforming to the epistemic norms *always* promotes at least one goal you happen to have, whether that's because we all share a common goal that is promoted, or because conforming promotes every goal. This is where standard developments of instrumentalism start focusing on the way that conforming promotes specific individual goals.

There are various reasons to develop the theory in this universal direction. One such reason is that the orthodoxy holds that the vast majority of philosophers judge that Idealized Max has epistemic reasons. Perhaps partially explaining why I am not drawn to universalism, I don't judge that way. To make my judgment a bit more vivid, imagine a **Zen Boltzmann brain**, that is a brain that spontaneously came into existence in a void, lacking any goals we have and lacking a body, so it is essentially causally isolated. Is there some normative force compelling the Zen Boltzmann brain to conform to the epistemic norms? I don't have a strong intuition one way or another but if I were to bet, I would say there isn't. I have a similar response to Idealized Max. Perhaps there are more philosophers who judge similarly to me than commonly supposed. But I am not questioning what judgments philosophers make in this paper. Suppose the orthodoxy is right, that the vast majority of philosophers judge that Idealized Max has epistemic reasons. That is a motivation to develop a universal solution but it is not a strong one. These judgments do not have much epistemic weight.

In response to idealized cases, I offer an error theory regarding too-few-reasons judgments.¹⁸ To warm you up to the idea of an error theory, allow me to make the following three observations. First, perhaps no actual agent has ever had just one goal such as "believe *p*". Indeed, I am not sure it is nomically possible for such an agent to exist. Second, even if there were such an agent, they would have a psychology very unlike the vast majority of humans. Third, the judgment we are considering is that the idealized agent has epistemic reasons. This is not a Moorean proposition like "I have hands", which we arguably have more reasons to believe than any philosophical argument to the contrary could provide. The judgment we are considering concerns a technical concept in an extremely unusual context. Even if you take

18. A similar response is offered by Sharadin (2018), although he does not have the resources of embedded epistemic instrumentalism to support this response. Also, Sharadin's positive view has at least one serious issue: see Section 5.

philosophical judgments seriously, this is not a strong candidate for an epistemically superb judgment.

The embedded epistemic goal model offers an explanation for why we would make such judgments. The epistemic goal is so well-embedded that it is an efficient heuristic to treat epistemic reasons as if they applied always and everywhere (just as in the health case). We can talk and act as if epistemic reasons apply always and everywhere, because they in fact apply to almost all actual agents almost all the time. The cases where there are no epistemic reasons are so unusual that the heuristic rarely provides the wrong judgment. And even in cases where we do over generalize, such as in response to *Idealized Spoiler Alert*, the mistake has little downside. The only cost is that our abstract epistemic theorizing is led astray. But if we need a place to sleep, or we are hungry, or we don't know where our kids are, some damage to our abstract epistemic theorizing can be safely ignored. But as we are now fat and warm and safe, we may stop to appreciate how this heuristic is misfiring in very unusual cases and correct our abstract epistemic theorizing.

This response was partly anticipated by Thomas Kelly (2003, p. 623, also see pp. 622, 632–633), who rejects it:

The viability of this instrumentalist response, of course, presupposes that there is some shared cognitive goal which might underwrite the existence and intersubjectivity of epistemic reasons. In fact, it is here, I believe, where the instrumentalist conception of epistemic rationality founders: there is simply no cognitive goal or goals, which it is plausible to attribute to people generally, which is sufficient to account for the relevant phenomena. Individuals do not typically have this goal: believing the truth.

Kelly's rejection of the error theory response is plausible if the relevant goal in question is an explicit goal, or if we think about some specific goal people have in isolation. Not everyone explicitly wants to

conform for its own sake. Getting coffee doesn't require us to conform to the epistemic norms wholesale. Kelly's reply lands hard when the instrumentalist is focusing on the trees rather than the forest.

However, thinking about instrumentalism in this way is a mistake. The plurality of goals people have gives rise to the well-embedded epistemic goal. Many people have the embedded epistemic goal because they want to be reliable testifiers, because they want to achieve their long-term plans, because they don't want to be patronized, because they don't want to make people uncomfortable — to present just a few of the long list of goals any typical person will have. As in the health case, more than enough people have the embedded epistemic goal to cause the practice of treating epistemic norms as if they gave rise to reasons that apply always and everywhere.

Too-few-reasons objectors like Kelly are not on firm ground. The proponents of the embedded epistemic goal model are committed to the claim that philosophers are overgeneralizing from a useful heuristic in response to idealized too-few-reasons cases. However, we know that humans use heuristics and we know they over-generalize from them, so it is no surprise that we would be doing this in response to cases that involve very unusual agents. Meanwhile, too-few-reasons objectors are committed to the claim that our judgments in response to idealized too-few-reasons cases are accurate. Considering these judgments are the product of messy processes — like biological and cultural evolution as well as on-the-fly idiosyncratic learning processes — and considering the cases involve very unusual agents, this is a more significant commitment than the overgeneralizing one. So the error theory here is not some last-ditch effort to save instrumentalism. Rather it is at least as plausible as (if not more so than) maintaining that judgments in response to idealized too-few-reasons cases are uniformly accurate.

To sum up the response of the too-few-reasons objection: first, the objection is not made from an uncontroversial position. Usefulness is how we learnt to appreciate conforming, so positing an additional source of normativity is suspiciously epistemically convenient.

Second, actual too-few-reasons cases that are put to instrumentalists, like Spoiler Alert, are easily replied to: there are plenty of relevant goals. Third, the cases that are not easily replied to, such as Idealized Spoiler Alert, are plausibly replied to with an error theory — an error theory that is at least as likely as, if not more so than, the claim that our judgments in response to idealized cases are all strictly accurate. So the too-few-reasons objector is not in a strong position.

Now, I want to get precise about exactly what the embedded epistemic instrumentalist is claiming. Suppose in *Idealized Spoiler Alert*, that Idealized Max believes “I have no evidence that character X dies”, flouting the epistemic norms. In response, embedded instrumentalists hold that Idealized Max has no epistemic reasons against believing “I have no evidence that character X dies”. The idealized agent lacks the rich set of goals that give rise to epistemic reasons. Importantly, embedded instrumentalists do not say that Idealized Max *has* epistemic reasons to believe, either.¹⁹ Idealized Max requires the embedded epistemic goal, or some other relevant goal, to have epistemic reasons. But he lacks any epistemic reasons. In terms of “rationality”, Idealized Max is neither epistemically rational nor epistemically irrational for believing “I have no evidence that character X dies”. I consider Idealized Max pragmatically rational. I would consider him pragmatically irrational if he failed to believe. Although your view will turn on your view of pragmatic rationality. I also consider Idealized Max all-things-considered rational, though again, your view will turn on your view of all-things-considered rationality. Of course, embedded instrumentalists can still evaluate Idealized Max as if he had the embedded epistemic goal. They can point out that Idealized Max didn't believe in accordance with the epistemic norms. They just deny that this fact has normative significance on its own.

19. This is one of the differences between an instrumental position and a purely pragmatic position. The pure pragmatist holds that all reasons are of the same kind. They do not draw a distinction between epistemic and non-epistemic reasons. Epistemic instrumentalists draw a distinction, holding that there is a special class of epistemic reasons.

5. The Advantages of Embedded Epistemic Instrumentalism

Instrumentalism has many virtues. I articulated three such virtues in the introduction: it avoids metaphysical and epistemic difficulties associated with other accounts; it is compatible with a naturalistic world view; and it is not a theory about mere semantics. To this we can add that embedded instrumentalism does not require some convenient epistemic coincidence nor does it rest on epistemically dubious judgments (see Section 4). However, many theories in the vicinity of embedded epistemic instrumentalism share these virtues. Embedded epistemic instrumentalism has other advantages over these accounts. In this section, I will describe some of these advantages.

Contrast embedded epistemic instrumentalism with the standard instrumental position, universal epistemic instrumentalism. (Côté-Bouchard, 2015; Foley, 1987; Kornblith, 1993; Leite, 2007; Schroeder, 2007). Universal instrumentalists hold that we always and everywhere have epistemic reasons to conform to the epistemic norms. They hold that in every too-few-reasons case, no matter how idealized, there are instrumental reasons to conform to the epistemic norms. As such, standard instrumentalism also rests on those epistemically dubious judgments about highly idealized cases. Moreover, universal instrumentalism has resulted in some complicated positions, which have provoked many criticisms.²⁰ The problems with universal instrumentalism stem from the fact that universal instrumentalists are holding on to two ideas:

1. We always have epistemic reasons to conform to the epistemic norms, no matter the situation (this is the universal part).
2. Epistemic reasons boil down to facts about usefulness (this is the instrumental part).

20. For criticism of various universal positions, see Côté-Bouchard (2015), Lockard (2013), Sharadin (2018), and Willoughby (2022).

Combining these parts, the universalist must hold that it is always useful to conform to the epistemic norms. But consider our inductive evidence against this. In the vast majority of cases, whether or not something is useful depends on the situation you are in. A screwdriver is useful if you need to screw a screw but terrible for making an espresso. Having a lot of money is useful for buying a house but terrible in a revolution. Consider how surprising it would be to learn that a tool is always useful, to learn that there are no possible situations where that tool wouldn't be at least a little helpful — extraordinarily surprising. The fickle nature of usefulness suggests we should have a low credence for the idea that conforming to the epistemic norms is always useful. This is why the central fact that motivates instrumentalism, that conforming is widely useful, is a poor foundation for trying to explain the purported universal truth, that the epistemic norms have normative force always and everywhere. Usefulness is a capricious property, shifting wildly with slight changes to context. It is ill-suited for the universal project. This is why attempts to argue for the universal position require a great deal of philosophical gymnastics and why the universal project is unlikely to succeed. Embedded epistemic instrumentalism, by contrast, does not have this internal complexity. It does not have to defend improbable claims. It is a much simpler theory with a stronger reply to the too-few-reasons objection. It is a more internally coherent version of epistemic instrumentalism than the standard.

Moreover, the idea of an embedded goal allows embedded instrumentalism to address another concern about standard developments of instrumentalism. Mathew Lockard (2013, p. 1706) describes this additional concern:

One serious (and, in my view, underappreciated) difficulty facing such forms of epistemic instrumentalism concerns the level of conceptual sophistication they demand of individuals whose beliefs may properly be regarded as epistemically rational or irrational. Desiring to achieve an epistemic goal presumably requires one to have a

conceptual repertoire that includes such concepts as proposition, belief, and truth. Young children and higher animals seem capable of having rational and irrational beliefs despite lacking these concepts—hence despite lacking desires to achieve any epistemic ends whatsoever.

But embedded goals can have a normative grip on us while not being desired or even being unconceived. As a result, this overly demanding objection is not a problem for embedded epistemic instrumentalism. Indeed, as I mention above, even creatures that have no sophisticated cognition such as coral have something resembling an embedded epistemic goal. So embedded instrumentalism has at least two serious advantages over standard instrumentalism: it offers a stronger response to the central objection, and it is plausibly undemanding of epistemic agents.

There are some other non-standard, non-universal instrumental accounts of epistemic normativity, however, and embedded epistemic instrumentalism has advantages over them, too. One such advantage is that it offers an account of epistemic normativity. Contrast this with ecumenical epistemic instrumentalism, as defended by Nathaniel Sharadin (2019).²¹ Sharadin's position is that if there is an epistemic reason to believe something, then there is also an instrumental reason to believe that same thing but that we should be silent about whether the epistemic reason is identical to the instrumental one. This position is ecumenical in the sense that philosophers of all metaphysical stripes—naturalists and non-naturalists alike—could accept it. A central motivation for holding this position, according to Sharadin, is that the literature already implicitly identifies the instrumental position with the ecumenical position (2019, Section 5). Much of the literature on epistemic instrumentalism is concerned with whether there are instrumental reasons in all the purported cases of epistemic reasons; essentially, a lot of the literature is concerned with the too-few-reasons

21. Sharadin's position can be considered universal or non-universal, though he is a non-universalist in other work (Sharadin, 2018).

objection. Sharadin holds that this implies that both instrumentalists and non-instrumentalists already identify the instrumental position with the ecumenical position; the thought being that because much of the disagreement in the literature concerns whether there are instrumental reasons in certain cases, and ecumenical instrumentalism involves the claim that there always is, then the literature is implicitly just debating ecumenical instrumentalism. I reject this supposed implication. Rather, I hold that philosophers believe that answering whether there are instrumental reasons in certain cases will provide evidence for or against non-ecumenical instrumentalism (the position that epistemic reasons are instrumental reasons). If there are instrumental reasons in all cases of epistemic reasons, then that is evidence that epistemic reasons are a sub-species of instrumental reasons, as this is the obvious explanation for why epistemic reasons always come in tandem with instrumental ones. Alternatively, if there are epistemic reasons without instrumental reasons, then that is evidence for epistemic reasons being something else. Moreover, the question of whether there are instrumental reasons in certain cases loses much of its importance once it is divorced from the question of whether epistemic reasons are instrumental reasons. I am (and I suspect many other philosophers in the literature also are) trying to understand the nature of epistemic normativity. Ecumenical instrumentalism refuses to engage with that concern. While staying silent on that issue may result in a position that is easier to defend, it undermines a key objective of an investigation into epistemic normativity. Embedded instrumentalism engages with the core concerns of many philosophers who are interested in epistemic normativity in a way that ecumenical instrumentalism does not.

Dyke also offers a non-universal instrumentalism. As I outlined above, she holds that communities have epistemic goals and that epistemic normativity arises because people are a part of these communities (Dyke, 2020). A problem with this position is that she does not identify how being part of a community with some goal gives an individual reasons to achieve that goal. She does not require people

to endorse the community or personally share the community's goal, making it unclear what the reason-giving-connection is between individuals and communities. Furthermore, there is evidence that community goals alone do not give rise to individual reasons. I might be part of a community that pursues sexist ends, but that does not give me reasons to pursue those sexist ends. Similarly, I might be part of an epistemic community, but that does not give me reasons to pursue those epistemic ends. So Dyke offers no account of how community goals give rise to individual reasons, and there is evidence that community goals alone don't give rise to individual reasons. Contrast this with embedded epistemic instrumentalism, which offers reasons that bear on individuals. It is your complex web of goals — including those shared with a community — that give rise to the embedded epistemic goal.

There are also non-instrumental accounts that have many of the virtues of instrumental accounts. Consider Kate Nolfi's (2018) account of epistemic normativity, which holds that the proper function of beliefs gives rise to epistemic normativity. Here, the "proper function" of an object means the function that was selected-for when the object evolved. We might say the heart has the proper function to pump blood, for example. Suppose the proper function of beliefs is to provide a map to guide your behavior. Roughly, Nolfi holds that beliefs should conform with epistemic norms because that best achieves their proper function. This shares many of the virtues of instrumentalism. But a problem with this approach, much like Dyke's approach, is that it has nothing to do with us. For example, at the time and location of writing, I am in a very calorie-rich environment. My body evolved to store a lot of calories. I would prefer it if my body didn't store them. I like eating a lot and I want to live a long time. I am in a very different environment to the one in which my ancestors evolved. The proper function of my fat storage system doesn't match my current environment. The proper function of my fat storage system isn't about me in a relevant sense. Similarly, the proper function of my beliefs is not an important factor when considering what I should believe, because it

isn't about me. Contrast this with embedded epistemic instrumentalism. The goals or ends that I have are about me. They can even explain what might be appealing about a proper function theory. If by fulfilling its proper function, my fat storage system promotes the goal of me living longer, its proper function becomes important. If the proper function doesn't promote (or hinder) my goals or ends, I can discard consideration of it. Embedded epistemic instrumentalism provides reasons for individuals to have certain beliefs, in a way that Nolfi's and Dyke's accounts do not.

Of course, this is not an exhaustive list of advantages. Some philosophers may even find disadvantages in holding the embedded position. But this is nonetheless a strong candidate for an account of epistemic normativity. It has the virtues of instrumentalism. It has advantages over both standard and non-standard developments of instrumentalism. It explains how epistemic reasons bear on individuals — on you. So embedded epistemic instrumentalism has a lot going for it.

Acknowledgments

Noga Gratvol, Alan Hájek, Frank Jackson, Thomas Kelly, Mathew Kopeck, Hilary Kornblith, Katie Steele, Julia Stifle, Jeremy Strasser, Timothy L. Williamson, Shang Long Yeo, three anonymous reviewers, as well as audiences at the Centre for Human Values, University of Colorado Boulder, the Australasian Association of Philosophy Conference 2019 and the ANU morning tea crowd. Thanks to Thom Fulton and Anya Tchoupakov for editorial services. This research was supported by the Australian Government Research Training Program (RTP) Scholarship.

References

- Berker, S. (2013). The rejection of epistemic consequentialism. *Philosophical Issues*, 23(1), 363–387. doi: 10.1111/phils.12019
- Clifford, W.K. (1877 [1999]). The ethics of belief. In T. Madigan (Ed.), *The Ethics of Belief and Other Essays* (pp. 70–96): Prometheus.
- Conee, E., & Feldman, R. (2004). *Evidentialism: Essays in Epistemology*: Oxford University Press.

- Côté-Bouchard, C. (2015). Epistemic instrumentalism and the too few reasons objection. *International Journal of Philosophical Studies*, 23(3), 337–355. doi: 10.1080/09672559.2015.1042007
- Cottingham, J., Stoothoff, R., & Murdoch, D. (1985). *The Philosophical Writings of Descartes: Volume 2*: Cambridge University Press.
- Cowie, C. (2014). In defence of instrumentalism about epistemic normativity. *Synthese*, 191, 4003–4017. doi: 10.1007/s11229-014-0510-6
- Dyke, M. M. (2020). Could our epistemic reasons be collective practical reasons? *Noûs*, 55(4), 842–862. doi: 10.1111/nous.12335
- Elgin, C. (2005). Non-foundationalist epistemology: Holism, coherence, and tenability. In M. Steup & E. Sosa (Eds.), *Contemporary Debates in Epistemology* (pp. 156–167): Blackwell.
- Enoch, D. (2011). *Taking Morality Seriously: A Defense of Robust Realism*: Oxford University Press.
- Foley, R. (1987). *The Theory of Epistemic Rationality*: Harvard University Press.
- Foley, R. (1993). *Working Without a Net: A Study of Egocentric Epistemology*: Oxford University Press.
- Friedman, J. (2018). Teleological epistemology. *Philosophical Studies*, 176, 673–691. doi: 10.1007/s11098-018-1033-7
- Giere, R. N. (1989). Scientific rationality as instrument rationality. *Studies in History and Philosophy of Science Part A*, 20(3), 377–384. doi: 10.1016/0039-3681(89)90013-7
- Goldman, A. I. (1967). A causal theory of knowing. *Journal of Philosophy*, 64(12), 357–372.
- Kelly, T. (2003). Epistemic rationality as instrumental rationality: A critique. *Philosophy and Phenomenological Research*, 66(3), 612–640. doi: 10.1111/j.1933-1592.2003.tb00281.x
- Kelly, T. (2007). Evidence and normativity: Reply to Leite. *Philosophy and Phenomenological Research*, 75(2), 465–474.
- Kitcher, P. (1992). The naturalists return. *The Philosophical Review*, 101(1), 53–114. doi: 10.2307/2185044
- Kornblith, H. (1993). Epistemic normativity. *Synthese*, 94(3), 357–376.
- Laudan, L. (1986). *Science and Values: An Essay on the Aims of Science and Their Role in Scientific Debate*. University of California Press.
- Laudan, L. (1990a). Aim-less epistemology? *Studies in History and Philosophy of Science Part A*, 21(2), 315–322. doi: 10.1016/0039-3681(90)90029-8
- Laudan, L. (1990b). Normative naturalism. *Philosophy of Science*, 57(1), 44–59.
- 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), 456–464. doi: 10.1111/j.1933-1592.2007.00084.x
- Lockard, M. (2013). Epistemic instrumentalism. *Synthese*, 190, 1701–1718. doi: 10.1007/s11229-011-9932-6
- Moore, G. E. (1903). *Principia Ethica*: Cambridge University Press.
- Nolfi, K. (2018). Why only evidential considerations can justify belief. In C. McHugh, J. Way, & D. Whiting (Eds.), *Normativity: Epistemic and Practical* (pp. 179–199): Oxford University Press.
- Nozick, R. (1993). *The Nature of Rationality*. Princeton University Press.
- Parfit, D. (2011). *On What Matters: Two-Volume Set*. Oxford University Press.
- Quine, W. V. (1969). *Ontological Relativity & Other Essays*. Columbia University Press.
- Quine, W. V., & Ullian, J. S. (1970). *The Web of Belief*: Random House.
- Schroeder, M. A. (2007). *Slaves of the Passions*: Oxford University Press.
- Sharadin, N. (2018). Epistemic instrumentalism and the reason to believe in accord with the evidence. *Synthese*, 195, 3791–3809. doi: 10.1007/s11229-016-1245-3
- Sharadin, N. (2019). Ecumenical epistemic instrumentalism. *Synthese*, 198: 2613–2639. doi: 10.1007/s11229-019-02232-7
- Siegel, H. (1989). Philosophy of science naturalized? Some problems with Giere's naturalism. *Studies in History and Philosophy of Science Part A*, 20(3), 365–375. doi: 10.1016/0039-3681(89)90012-5
- Siegel, H. (2019). Epistemic rationality: Not (just) instrumental. *Metaphilosophy*, 50(5), 608–630. doi: 10.1111/meta.12392

- Sosa, E. (2003). The place of truth in epistemology. In M. DePaul & L. Zagzebski (Eds.), *Intellectual Virtue: Perspectives from Ethics and Epistemology*: Oxford University Press.
- Wedgwood, R. (2007). *The Nature of Normativity*. Oxford University Press.
- Williamson, T. (2000). *Knowledge and its Limits*: Oxford University Press.
- Willoughby, J.B. (2022). *Against Universal Epistemic Instrumentalism* [Unpublished Manuscript].