

In Defense of Constitutivism About Epistemic Normativity

David Horst

(penultimate draft; final version forthcoming in *Pacific Philosophical Quarterly*)

Epistemic Constitutivism (*EC*) holds that the nature of believing is such that it gives rise to a standard of correctness, and that other epistemic normative notions (e.g., reasons for belief) can be explained in terms of this standard. If defensible, this view promises an attractive and unifying account of epistemic normativity. However, *EC* faces a forceful objection: that constitutive standards of correctness are never enough for generating normative reasons. This paper aims to defend *EC* in the face of this objection. I do so in two steps. First, I dispute a crucial assumption underlying the case against *EC*: that constitutive standards of correctness in general are “reason-giving” only if and because there is also a *prior* reason to comply with them. Second, I outline a strategy of how *EC* can meet the challenge of explaining what’s special about the activity of believing such that, unlike other standard-governed activities, it is capable of generating normative reasons.

1. Introduction

Epistemic Constitutivism (*EC*) is the attempt to ground epistemic normativity in the nature of belief. In this, *EC* is the epistemic version of a corresponding constitutivist strategy in metaethics: the attempt to ground practical normativity in the nature of action.¹ More specifically, the basic idea behind *EC* is that the nature of believing is such that it gives rise to a standard of correctness, and that other epistemic normative notions can be explained in terms of this standard.² In particular, proponents of *EC* use belief’s constitutive standard of correctness to explain what it is for some consideration to be an epistemic normative reason for belief. On this view, roughly, what makes a consideration an epistemic normative reason for believing *P* is that it bears positively on whether believing *P* is correct. It’s because believing is governed by a constitutive standard of correctness that some considerations are epistemic normative reasons for belief: namely those that count in favor of a belief by bearing on its correctness.³

¹ For versions of practical constitutivism, see, e.g., Velleman (2000), Foot (2001), Korsgaard (2008), Ferrero (2009), Smith (2013), Lavin (2017), and Schafer (2018). Many have emphasized the parallel between the constitutivist projects in epistemology and metaethics, see, e.g., Velleman (2000), Railton (1997), Tubert (2010), and Wiland (2012: ch. 6).

² With some qualifications, versions of *EC* can be found in, e.g., Railton (1997), Velleman (2000), Boghossian (2003), Wedgwood (2002), Shah (2003, 2006), and Lord/Sylvan (2019).

³ In what follows, I will mostly drop the qualifier “normative” when discussing normative reasons. Unless otherwise specified, by “reason” I mean normative reason. I understand normative reasons in the usual way: as considerations that *count in favor of* acts and attitudes. As such, normative reasons differ both from motivating reasons (reasons *for which* one acts or holds an attitude) and from explanatory reasons (reasons *why* one acts

If defensible, *EC* would offer a straightforward and unifying account of epistemic normativity. But just like the corresponding constitutivist strategy in metaethics, *EC* is subject to a forceful objection: constitutive standards seem insufficient to ground genuine reasons. There are many human activities—such as playing chess, dancing tango, or cooking risotto—which are arguably governed by constitutive standards of correctness. But none of these activities gives rise to reasons. That is, in none of these activities do considerations amount to reasons for performing a particular chess-move, dance-move, or cooking step *just because* they bear on the correctness of the relevant chess-move, dance-move, or cooking step. In light of this objection, at the very least, proponents of *EC* owe us an explanation of what’s special about the activity of believing such that, unlike other human activities, it does give rise to reasons. However, many philosophers are convinced that *EC* cannot meet this challenge on its own terms. This is because they think that, in relevant respects, all standard-governed activities are the same: in all such cases, correctness-relevant considerations amount to reasons only if and because there is a *prior* reason to comply with the relevant standard. If this were true, *EC* would be doomed indeed. For, then, belief’s correctness-standard would not be explanatorily fundamental in the relevant sense: any normative authority or reason-giving force it may have would in turn derive from an independently given (set of) reason(s).

The aim of this paper is to defend *EC* in the face of this objection. I do so in two steps. First, I shall dispute the crucial assumption underlying this objection: that there is never a reason to do what it is correct to do by the constitutive standard of some activity unless there is also a *prior* reason to correctly engage in that activity. In particular, I argue that applying this principle to the activity of believing threatens to seriously undergenerate epistemic reasons (sections 4 and 5). However, dismissing the requirement for prior reasons doesn’t yet show what, if anything, is special about the activity of believing. That’s why, secondly, I outline a strategy of how *EC* might be able to meet the challenge of distinguishing believing from other standard-governed activities in a way that helps explain why considerations bearing on the correctness of some belief amount to reasons for that belief (sections 6 and 7). I will begin with a more detailed presentation of *EC* and the challenge facing the view (sections 2 and 3).

2. Epistemic Constitutivism

As noted above, *EC* can be defined by two central claims:

- (1) It’s part of the nature of the activity of believing that instances of this activity—particular beliefs—are subject to a standard of correctness (i.e., truth).

or holds an attitude). I assume that epistemic reasons are a species of *normative* reasons. For a defense of this claim, see Paakkunainen (2019) and Kiesewetter (2021).

- (2) Other epistemic normative notions (e.g., the notion of an epistemic reason) can be explained in terms of belief's constitutive standard of correctness.

To illustrate the first claim, proponents of *EC* often turn to other standard-governed activities.⁴ Consider playing chess. Typically, when we say that a certain chess-move is the *correct* move to make, we evaluate the move *qua* chess-move. And the standard for being correct *qua* chess-move is plausibly one that's internal to the activity itself. This means that it is a standard to which particular chess-moves are subject simply because they are instances of playing chess, not (say) because of any contingent end one might pursue in making this move. What counts as the correct move to make relative to this standard is, arguably, the permissible move that best serves the objective of checkmating one's opponent.⁵ When we assess chess-moves *qua* chess-moves, we evaluate them in terms of their relation to this standard of correctness. Since chess-moves are subject to this standard simply because they are instances of playing chess, the standard deserves to be called a *constitutive* standard.

Proponents of *EC* claim that something similar holds for believing. Typically, when we say that believing a certain proposition is *correct*, we evaluate the belief *qua* belief. Again, the standard for being correct *qua* belief is plausibly one that's internal to the activity of believing itself. It is a standard to which particular beliefs are subject simply because they are instances of believing, not (say) because of some contingent end that one might happen to have. The standard of correctness for belief is commonly taken to be truth. Thus, when we assess beliefs *qua* beliefs, we evaluate them in terms of their relation to the standard of truth. Since an attitude is subject to this standard simply because it is an instance of believing, the standard deserves to be called a *constitutive* standard of belief.

There is, of course, much more to say about what makes an activity such that its instances are subject to an evaluative standard simply because they are instances of that activity.⁶ However, my focus in this paper will be on *EC*'s second defining claim: that other epistemic normative notions can be explained in terms of belief's constitutive standard of correctness. In particular, as mentioned above, proponents of *EC* seek to explain the notion of an epistemic reason for belief in terms of this standard—that's the claim I'll focus on.⁷ According to a simple version of this view,

⁴ See, e.g., Velleman (2000: 187f.).

⁵ So, importantly, as I use the term, a correct chess-move isn't *just* one that conforms to the rules of chess, it's one that *promotes success*, relative to the defining objective of chess (i.e., checkmating one's opponent).

⁶ According to a prominent view, activities that ground evaluative standards are functional or teleological kinds: i.e., activities that are partly defined by a certain function, aim, or point. See, e.g., Velleman (2000), Thomson (2008), and Korsgaard (2008).

⁷ See Lord/Sylvan (2021) for an extension of the constitutivist story to epistemic reasons for suspension.

a consideration is an epistemic reason to believe P just in case and because it's a consideration that bears positively on whether believing P is correct (i.e., on whether P is true).⁸ The prime example of considerations that bear positively on whether believing P is correct is evidence for P's truth. So, on this view, evidential considerations are epistemic reasons for beliefs simply because they are considerations that bear on the correctness of these beliefs. This distinguishes *EC* both from views that seek to ground epistemic reasons in something other than belief's constitutive standard (e.g., in the *value* of believing correctly) and from views that take the fact that some considerations are epistemic reasons to be *primitive*.⁹

One should note that *EC* isn't committed to the claim that epistemic reasons are *only* constituted by evidence. There might be considerations that, in some way, bear on whether believing P would be correct without being evidence for or against P. For example, the fact that I lack any evidence for or against P arguably bears on whether believing P would be correct without being (strictly speaking) evidence for or against P.¹⁰ Nor does *EC* imply that there can't be *practical* reasons for belief: i.e., considerations that count in favor of a belief, not by bearing on its correctness, but by bearing on (say) the benefits of so believing.¹¹ At least primarily, *EC* is just a claim about *epistemic* reasons for belief; as such, the view leaves open whether there might be reasons for belief other than epistemic ones. To facilitate the following discussion, I will focus on evidence as a paradigm example of epistemic reasons for belief, setting aside questions as to whether all epistemic reasons are evidence, or whether all normative reasons for belief are epistemic. The core dispute between proponents of *EC* and their opponents is about what, if anything, makes some consideration—paradigmatically: evidence—an epistemic reason for a belief. Proponents of *EC* maintain, roughly, that it is some relation in which that consideration stands to belief's constitutive standard of correctness; opponents of *EC* deny this because, as we will see, they hold that constitutive standards of correctness are never enough to ground genuine normativity.

3. *The Objection from Games*

One of the most prominent objections to *EC* is what I call the *objection from games*, which specifically targets *EC*'s second defining claim. In short, the problem is that there seem to be many activities which are governed by a constitutive standard, but which do not give rise to reasons. If so,

⁸ See, e.g., Velleman (2000: 15) and Shah (2006: 489).

⁹ I briefly discuss value- and goal-based views of epistemic reasons in section 5. For primitivism about normative reasons in general, see, e.g., Scanlon (1998) and Parfit (2011).

¹⁰ See Littlejohn (2018: 9) for this point.

¹¹ For a defense of practical reasons for belief, see, e.g., Leary (2017) and Rinard (2015).

standard-governed activities by themselves seem insufficient to ground reasons. Here is how David Enoch puts the crucial point of this objection:

[E]ven if you find yourself engaging in a kind of activity [...] and even if that activity is constitutively governed by some norm [...] this does not suffice for you to have a reason to obey that norm [...]. Rather what is also needed is that you *have a reason* to engage in that activity. [...] Even if you somehow find yourself playing chess, and even if checkmating your opponent is a constitutive aim of playing chess, still you may not have a reason to (try to) checkmate your opponent. You may lack such a reason if you lack a reason to play chess.’ (Enoch 2011: 210f.)¹²

To make this more vivid, suppose you are playing chess. It’s your turn to move and you know that by advancing your knight to c5 you could take out your opponent’s queen. This, we may assume, is the *correct* move to make in the situation of the game, relative to the game’s constitutive standard: i.e., it’s the available move that best serves the objective of checkmating your opponent. Now, does this fact provide you with a reason to move your knight to c5? Well, that depends. It depends on whether you have a *prior* reason to play chess and win. If you do, then considerations that bear positively on the correctness of particular chess-moves—like the fact that moving your knight to c5 would take out your opponent’s queen—will provide you with reasons to perform these moves. But if you lack such a prior reason, then considerations relevant to the correctness of particular chess-moves seem to lack any reason-giving force as well. Suppose, for example, you have been offered a lot of money to let your opponent win. In that case, the fact that moving your knight to c5 will take out your opponent’s queen may very well *not* be a reason to perform that move. (To the contrary, it may be a reason to *refrain* from making that move.) What this suggests is that, in the case of chess, considerations that bear positively on the correctness of some chess-move do not *necessarily* amount to reasons to perform that move. Whether they do amount to such reasons depends entirely on whether you have a prior reason to comply with the correctness-standard of chess (e.g., a reason to play chess and win).

Some may want to insist that you do not really count as “playing chess” unless you intend to win and that intending to win gives you (at least *pro tanto*) reasons to perform correct chess-moves.¹³ I don’t think it’s true that you can’t play chess without intending to win. But even if it were true, the mere fact that you intend to win is hardly sufficient to give you reasons to perform correct chess-moves (even *pro tanto* reasons). Suppose your opponent credibly threatens to kill your entire family if you don’t let him win. If so, surely you don’t have reasons to make correct chess-moves (i.e.,

¹² Enoch’s primary target is constitutivism about *practical* normativity, but he states explicitly that he intends his criticism to apply to *epistemic* constitutivism as well (see Enoch 2006: 171). Others make similar points, drawing on similar examples. See, e.g., Rosen (2001: 622).

¹³ See, e.g., Katsafanas (2018).

moves conducive to your opponent's defeat). And this is true even if you do in fact intend to checkmate your opponent, and it remains true even if you are somehow forced to play chess and, presumably, can't help but intend to win. At best, if you have no choice in the matter, this may exculpate you, but it doesn't give you good reasons to do what you know will lead to the demise of your entire family.

Playing chess, then, is an activity which is governed by a constitutive standard of correctness, but which, all by itself, doesn't give rise to reasons. Plausibly, moreover, playing chess is far from being the only example of such an activity. Many other human activities—such as dancing tango, cooking risotto, or building a house—seem to be exactly alike in this respect. In case of each such activity, considerations that bear positively on whether certain performances would meet their standard of correctness do not necessarily amount to reasons for these performances. Such cases spell trouble for proponents of *EC* because they constitute counterexamples to *EC*'s second defining claim: they suggest that standard-governed activities alone are not enough for generating reasons. That is, the mere fact that beliefs are subject to a standard of correctness seems insufficient to explain why considerations that bear positively on the correctness of particular beliefs amount to reasons for these beliefs. At the very least, then, proponents of *EC* owe us an explanation of what's special about the activity of believing such that, unlike other standard-governed activities, it does give rise to reasons.

I think meeting this challenge involves two tasks: (a) identifying a difference between believing and other standard-governed activities, and (b) showing that this difference is normatively relevant—that it makes a difference as to whether considerations that bear on the correctness of particular beliefs amount to reasons for that belief. Moreover, proponents of *EC* would need to accomplish these tasks without giving up on their commitment to the explanatory priority of activities and their constitutive standards. That is, they would have to say what's special about the activity of believing without appeal to some independent normative notion—like the *value* of believing correctly or some such. But before we can explore how *EC* might be able to meet this challenge, we need to address a prominent line of thought according to which *any* such attempt is doomed to fail.

4. The Analogy Between Believing and Playing Chess

Many think that *EC* *must* fail to meet the challenge. This is because many think that what holds for playing chess holds for standard-governed activities *in general*: namely that no correctness-relevant consideration amounts to a reason unless there is also a *prior* reason to engage in the relevant activity in a way that complies with its standard of correctness. On this view, given *any* standard of correctness, you may always sensibly ask whether there is good reason to comply with it in the first

place. If no such prior reason exists, then there is also no reason to do any of the more specific things that would (seem to) meet the relevant standard. This can be captured in the following claim:

Prior Reasons: For any particular ϕ -ing that is subject to a constitutive standard of correctness, a consideration that bears positively on whether ϕ -ing would meet its standard of correctness is a reason for you to ϕ only if and because there is a *prior* reason for you to comply with the relevant standard of correctness.

The idea here is that we can appeal to prior reasons to comply with specific standards of correctness to explain *when* and *why* correctness-relevant considerations themselves acquire the status of reasons. Thus, for example, it's only when and because you have a reason to play chess and win (e.g., that you enjoy competitive games) that you also have a reason to do what it is correct to do relative to the standard constitutive of chess (e.g., move your knight to c5). As we have seen, this looks plausible for activities like playing chess, cooking risotto, or dancing tango. The crucial assumption of *Prior Reasons*, however, is that we can generalize from these examples to *all* standard-governed activities: they are all alike in that there's only ever a reason to do what it is correct to do by their respective standards if and because there is also a prior reason to comply with these standards. For instance, something like this generalizing assumption is clearly at work in Enoch's well-known objection to constitutivism. Consider how he continues the passage quoted above:

'Even if you somehow find yourself playing chess, and even if checkmating your opponent is a constitutive aim of playing chess, still you may not have a reason to (try to) checkmate your opponent. You may lack such a reason if you lack a reason to play chess. *The analogy is clear enough:* even if you find yourself playing the agency game, and even if agency has a constitutive aim, still you may not have a reason to be an agent'. (Enoch 2011: 210f., my emphasis)

In moving from "playing chess" to "agency", Enoch assumes that, in all normatively relevant respects, being an agent is exactly analogous to being a chess player. Hence, what goes for chess goes for potentially all standard-governed activities. In each case, constitutive standards—or "constitutive aims", as Enoch puts it—are reason-giving only in conjunction with some prior reason to engage in the relevant activity in a way that complies with its standard. Others rely on something like *Prior Reasons* specifically in discussing belief's standard of correctness:

'[T]o posit a constitutive norm for belief is not yet to say that there is good reason to do as that norm says. As an analogy, writing the same number twice in one of the columns of a Sudoku grid is incorrect relative to the constitutive norms or rules of Sudoku. Yet there might be no good reason for me to avoid that incorrect Sudoku move. Similarly, we can very well accept that [...] a belief is correct if and only if it is true, but still ask whether there is any good reason to have correct beliefs.' (Côté-Bouchard 2016: 3193f.)

Again, the underlying thought here is that, quite generally, standards of correctness are reason-giving only if and because there is a prior reason to comply with them. Hence, just like in the case of games, a consideration that bears positively on the correctness of believing P doesn't automatically provide you with a reason to believe P. It does so only if and because there is a reason for you to have *correct* beliefs—as opposed to *incorrect* ones—in the first place. Thus, when we spell out what *Prior Reasons* amounts to in the case of believing, we get something like the following claim:

Prior Reasons_B: A consideration that bears positively on whether believing P would meet its standard of correctness (i.e., truth) is a reason for you to believe P only if and because there is a *prior* reason for you to comply with belief's standard of correctness (i.e., a reason to have correct/true beliefs).

If *Prior Reasons_B* were true, *EC* would be hopeless. For, then, what's explanatorily fundamental for epistemic normativity isn't the activity of believing, but prior reasons to engage in that activity—or whatever it is that grounds such prior reasons (e.g., the value of believing correctly). Whatever normative authority belief's standard of correctness might have, such authority would be inherited from something else. But is *Prior Reasons_B* true?

In the next section, I will argue that there is strong reason to doubt *Prior Reasons_B*. This is because the analogy between believing and activities like playing chess—on which *Prior Reasons_B* rests—breaks down at the crucial point. If so, this makes room to explore ways in which proponents of *EC* might be able to *positively* meet the challenge facing their view: to identify a difference between believing and playing chess which can help explain why considerations that bear on the correctness of particular beliefs necessarily amount to reasons for these beliefs.

5. Why the Analogy Breaks Down

Prior Reasons_B claims that a consideration that bears positively on the correctness of some belief doesn't amount to a reason for that belief unless there is also a *prior* reason to comply with belief's standard of correctness. What could such a prior reason be? Obviously, such reasons cannot themselves be considerations that bear on the correctness (i.e., truth) of particular beliefs. What we are looking for are considerations that count in favor of going in for *correct* beliefs—as opposed to *incorrect* ones—in the first place. It is only once we have established that there are such prior reasons for complying with belief's standard of correctness that truth-related considerations themselves acquire the status of reasons.

And, in fact, it's not hard to find examples of such prior reasons for having true beliefs. For instance, having true beliefs is often *useful* to us. Without true beliefs, we would be utterly lost in the world; we wouldn't be able to carry out the most basic tasks—like preparing food—unless we can rely on true beliefs about our surroundings. Arguably, moreover, believing the truth isn't only of practical importance to us; it can also be something we pursue for its *own sake*, simply because we are curious creatures. Thus, having true beliefs can be instrumentally or intrinsically *valuable* and/or can help us achieve our *goals* (including our cognitive goals).¹⁴ And, plausibly, facts about the value or goal-conduciveness of having true beliefs can provide you with reasons to comply with belief's standard of correctness. Accordingly, given *Prior Reasons_B*, when finding out whether P is true is in some sense valuable or serves a goal that you have, then considerations that bear positively on the correctness of believing P—e.g., evidence for P's truth—amount to a reason for you to believe P, as conforming to such considerations helps you track P's truth-value.¹⁵

However, even though we certainly often have a reason to have true beliefs—e.g., when having such a belief is useful—it seems equally clear that this isn't *always* the case. Some well-known examples from the literature on value- and goal-based accounts of epistemic normativity can help illustrate this point.¹⁶ Consider that many truths out there are completely trivial—like the truth about what is the 323rd entry in the Wichita, Kansas, phone directory, or the truth about whether Bertrand Russell was right- or left-handed.¹⁷ Plausibly, you have no reason to find out the truth about such trivial matters, doing so might certainly be neither valuable nor conducive to any of your goals. Nonetheless, you might stumble upon excellent evidence for one of these truths. If so, it seems that there is also a reason for you to believe that truth, one that's provided by your evidence. But that's not what *Prior Reasons_B* predicts. Since, by hypothesis, you lack any prior reason to have true beliefs about these matters, *Prior Reasons_B* predicts that there is also no reason for you to believe the particular trivial truth—despite the fact that you possess excellent evidence for it. Clearly, however, this looks like the wrong verdict.

So, trivial truths seem to constitute counterexamples to *Prior Reasons_B*. Moreover, trivial truths are hardly the only examples to this effect. Imagine a case where believing some truth would have *amful*

¹⁴ Such claims are the basis for value- and goal-based accounts of epistemic normativity. For value-based views, see, e.g., Goldman (1999), Lynch (2004), Alston (2005), and Steglich-Peterson (2011). For goal-based views, see, e.g., Foley (1987), Kornblith (1993), Cowie (2014), and Sharadin (2016).

¹⁵ To be sure, value- and goal-based views aren't the *only* options for understanding the relevant prior reasons (e.g., one might also adopt a fundamentalist approach to these reasons). But I don't think alternative approaches will help with the problems of extensional adequacy I am going to discuss. So, to keep things manageable, I will mainly focus on value- and goal-based approaches.

¹⁶ For critical discussion of value- and/or goal-based accounts of epistemic normativity, see, e.g., Grimm (2009), Kelly (2003), and Lockard (2013).

¹⁷ The examples are from Goldman (1999: 88) and Kelly (2003: 625).

consequences—say, an evil demon threatens to inflict great pain on you and your loved ones if you were to believe the truth as to whether P. Still, if you stumble upon decisive evidence for P’s truth, you seem to have an excellent epistemic reason to believe P, even though you lack any prior reason to have a true belief as to whether P—in fact, in such a case, you plausibly have a very strong prior reason to *avoid* having a true belief as to whether P.¹⁸

What these examples bring out, I think, is that the analogy between believing and games like chess—on which critics base their objection to *EC*—breaks down at the crucial point. For, intuitively, in cases of trivial and harmful truths, there is a reason for you to believe P—provided by your evidence for P’s truth—even though you *lack* any prior reason for having a true belief with respect to whether P. Hence, unlike in the case of playing chess, no prior reasons to comply with belief’s standard of correctness seem to be required for there to be reasons to hold particular beliefs, provided by considerations bearing on the correctness of these beliefs. If this is right, *Prior Reasons_B* looks seriously flawed, which in turn critically undermines the case against *EC* that relies on this claim.

To appreciate this point, recall the dialectical situation. We started with an important challenge to *EC*: to explain what’s special about the activity of believing such that, unlike other standard-governed activities, it can give rise to reasons. According to many critics, however, *EC* must fail to meet this challenge. This is because believing is said to be *just like* playing chess in relevant respects—in both cases, correctness-relevant considerations amount to reasons only if and because there is a prior reason to comply with the relevant standard of correctness. Yet, it’s precisely this analogy between believing and playing chess which the examples of trivial and harmful truths call into question. For, intuitively, in these examples, considerations that bear positively on the correctness of some belief amount to reasons for this belief despite the fact that you *lack* any prior reason to comply with belief’s standard of correctness.

To be sure, some philosophers have argued that, even in cases of seemingly useless or harmful truths, it’s always possible to identify some sense in which beliefs in these truths can be seen as valuable or goal-conducive.¹⁹ For example, some have argued that, despite appearances to the contrary, believing the truth is *always* intrinsically valuable, others have claimed that going in for true beliefs is the best strategy because it serves us best *overall* or *in the long run*.²⁰ But it should be clear that adopting one of these strategies to salvage the analogy between believing and playing

¹⁸ For a more mundane example to the same effect, see Kelly (2003: 626).

¹⁹ Another response to examples of trivial and harmful truths is to deny that, in these cases, evidence is a *genuinely normative* reason to believe what the evidence supports. See, e.g., Leite (2007) and Steglich-Petersen (2011). To my mind, Paakkunainen (2018) makes a compelling case against this type of response.

²⁰ See Lynch (2004) for the first proposal, Kornblith (1993) and Leite (2007) for versions of the second.

chess would lack any dialectical force within the context of a debate with *EC*. The theoretical assumptions underlying these strategies—e.g., the claim that *all* true beliefs are intrinsically valuable—tend to be highly controversial, even among those with no sympathies for *EC*.²¹ So it's hard to see why anyone would accept them unless they are already on board with a broadly value- or goal-based view of epistemic normativity. Certainly, no defender of *EC* would be moved by an argument premised on these assumptions. Moreover, from the point of view of *EC*, responding to the examples of trivial and harmful truths by trying to find, say, some value that *all* true beliefs possess is to miss the crucial lesson of these examples: namely that *no such value is needed* for the existence of epistemic reasons. So, even if all true beliefs possessed some—perhaps minimal—value, or even if all true beliefs served some—perhaps very remote—goal, it would still be unclear why proponents of *EC* should accept the *further* assumption that the existence of epistemic reasons *depends* on the presence of these values or goals (or *any* prior reasons to have true beliefs, for that matter). After all, the examples of trivial and harmful truths certainly suggest the opposite: i.e., that no such dependence exists. So why go looking for prior reasons when, apparently, there's no need for them?

In sum, then, proponents of *EC* seem well justified in rejecting the requirement for prior reasons to comply with belief's standard of correctness. This requirement rests on a flawed analogy with games like chess. What should also be clear, however, is that the initial challenge to *EC* still stands—the challenge to identify a distinctive feature of believing which can explain why considerations that bear on the correctness of particular beliefs necessarily amount to reasons for these beliefs. Dismissing *Prior Reasons_B* removes a crucial reason for thinking that *EC* is in principle incapable of meeting this challenge. This is an important first step in defending *EC* against the objection from games. But it doesn't yet show how *EC* can *positively* meet the relevant challenge.

Another way to make this point is to observe that accepting *Prior Reasons_B* isn't necessary to get the challenge going. All that's needed for the challenge to arise is the fact that there are standard-governed activities—like playing chess—which do *not* give rise to genuine reasons. If so, one might legitimately wonder why this should be any different in the case of believing. Pushing that question doesn't require one to embrace a view like *Prior Reasons_B*. Instead, for example, one might take the apparent difficulties facing attempts to explain epistemic reasons in terms of values, goals, and

²¹ For critical discussion of the idea that believing the truth is always intrinsically valuable, see, e.g., Grimm (2009). For objections to Kornblith's (1993) and Leite's (2007) "rule-based" versions of epistemic instrumentalism, see, e.g., Lockard (2013).

constitutive standards to favor a *primitivist* approach to such reasons, according to which it's simply a brute fact that some considerations (such as evidence) constitute reasons for belief.²²

My aim in the remainder of this paper is to sketch a strategy of how *EC* might be able to meet the challenge. Thus, in the next section, I will draw attention to an important difference between believing and playing chess; subsequently, in section 7, I will argue that this difference is indeed normatively relevant.

6. *How Believing Differs from Playing Chess*

According to *EC*'s first defining claim, when we evaluate beliefs *qua* beliefs, we do so in terms of their relation to belief's constitutive standard of correctness. This is analogous to the way in which we evaluate chess-moves. To bring out the contrast between believing and playing chess, it will be helpful to first distinguish more carefully between different positive relations in which beliefs and chess-moves can stand to their respective constitutive standards, giving rise to different evaluative dimensions—different dimensions of performing well *qua* believer or *qua* chess-player.

Borrowing some well-known machinery from Ernest Sosa (2007, 2010, 2015), we can distinguish between three dimensions of performing well: correctness, competence, and aptness.²³ Suppose you have formed a belief; we may then ask:

- (i) whether the belief is *correct*: i.e., whether it's true,
- (ii) whether it's *competent*: i.e., whether it manifests epistemic competence on your part, and
- (iii) whether it's *apt*: i.e., whether its correctness manifests your epistemic competence.

So far, this is analogous to the way in which we evaluate the performance of chess-players. Suppose you have moved your bishop to f4; we may then ask:

- (i) whether the move is *correct*: i.e., whether it best serves the objective of checkmating your opponent,
- (ii) whether it's *competent*: i.e., whether it manifests chess competence on your part, and
- (iii) whether it's *apt*: i.e., whether its correctness manifests your chess competence.

²² For primitivist approaches to normative reasons—presumably, including epistemic normative reasons—see, e.g., Scanlon (1998) and Parfit (2011).

²³ I am not suggesting that Sosa subscribes to the constitutivist project in epistemology as described here; I am only borrowing some machinery that, I think, is helpful to refine this project. Moreover, as we will see, my way of fleshing out “competence” differs significantly from Sosa's view. For more on how my account of epistemic competence differs from Sosa's, see Horst (forthcoming).

Each of these dimensions—correctness, competence, aptness—specifies a way in which a subject may perform *well* by the lights of the relevant activity’s constitutive standard (e.g., truth, checkmating). So, each of these dimensions specifies a different positive relation to the relevant constitutive standard:

- (i) a *correct* performance (belief, chess-move) is one that meets the standard,
- (ii) a *competent* performance—one that manifests a relevant competence—is one that’s done in response to considerations that bear on whether the performance would meet the standard (more on this below), and
- (iii) an *apt* performance is one that meets the standard because it’s done competently.

To bring out the difference between believing well and playing chess well, I suggest we take a closer look at the second evaluative dimension: competence.

In general, I take it that competent performances—competent beliefs, competent chess-moves—involve responding to considerations relevant to the correctness of one’s performance. Competent believers and competent chess-players aren’t blind: in trying to meet the relevant standard of correctness, they take into account aspects of their situation that bear on how to meet these standards. Thus, as I use the term, a competent chess-move is one that you perform in response to considerations that bear on the correctness of the move. And a competent belief is one that you hold in response to considerations that bear on the correctness of your belief. But, as I want to argue now, the *way in which* competent chess players and competent believers are responsive to their respective standards of correctness is very different, giving rise to an important evaluative difference between both kinds of activities.

To bring out the difference, suppose again that you are playing chess, this time against your own son. It’s your turn to move and you know that by advancing your bishop to f4 you could checkmate your son. However, not wanting to demoralize your son—who is only just getting a hang of the game—you refrain from advancing your bishop to f4, performing an inferior move instead. Does this call into question your skillfulness *qua* chess player? Does it show a *lack* or *deficit* of chess competence on your part? Surely not. In fact, for all it matters, you might be the world’s most skilled chess-player; it’s just that you don’t want to jeopardize your son’s blossoming interest in the game.²⁴

²⁴ What if you don’t count as “playing chess” unless you intend to win? Even if true, this wouldn’t undermine the present point. We could easily stipulate that you *do* intend to win. It’s just that you want to do so in a way that doesn’t completely demoralize your son, which is why you deliberately make the game more competitive than it would otherwise be.

What this suggests is that competence in chess is a matter of being merely *conditionally* responsive to considerations relevant to the correctness of one's chess-moves, conditional on an independent motivation to comply with chess' standard of correctness—"independent" in the sense that the relevant motivation isn't implied by the competence itself. Possessing chess competence and being motivated to exercise it on a particular occasion are two independent conditions. That is, we don't expect a competent chess player to respond to correctness-relevant considerations by performing the move they recommend, unless she has, say, a desire to play well. Hence, if you don't respond to the fact that advancing your bishop to f4 would be the correct move to make by performing that move, this doesn't necessarily reveal a *deficit* of chess competence on your part—not if you lack motivation to so respond in the first place (as in our example). For the same reason, the degree of competence exhibited by a chess player's performance is in no way diminished by the fact that she wouldn't have performed in that way, had she not had an independent motivation to do so. For example, we wouldn't retract our judgment that someone's chess game exhibited supreme skillfulness, if we found out that, had it not been for the prize money, she wouldn't have competed at all. Her performance might still be a masterclass in chess.

Moreover, this is clearly not an idiosyncratic characteristic of competence in chess. Suppose you are a French teacher, teaching your students French orthography. As you write French verbs on the blackboard, you deliberately include some spelling mistakes in order to quiz your students, asking them to identify the incorrect spellings. Does this call into question your competence at French orthography? Again, surely not. Being a competent speller is perfectly compatible with knowingly violating the standards of correct orthography when quizzing one's students. Thus, just like with competence in chess, orthographical competence is a matter of being merely *conditionally* responsive to the demands of correct orthography—conditional on an independent motivation to spell correctly on a given occasion. And, plausibly, the same holds for competent performance in many other standard-governed activities, such as dancing tango, cooking risotto, or building a house.

But now contrast this with competence in belief. Plausibly, what we expect from a competent believer is that she is, not just conditionally, but *unconditionally* responsive to considerations relevant to the correctness of her beliefs. Take a classic example: if you refrain from believing that your partner is cheating on you despite clear evidence that they are, then this reflects *negatively* on you *qua* believer, regardless of the fact that you might lack any motivation to comply with belief's standard of correctness on that occasion (or with respect to the topic of your partner's fidelity, more generally). This is a paradigm case of wishful thinking, and such cases do demonstrate a *deficit* of epistemic competence. Thus, unlike in the case of competent chess players, once a competent

believer recognizes the relevance of some fact to the correctness of her belief, we expect her to adjust the belief accordingly, and to do so *unconditionally*—not just when and as long as she happens to have an independent motivation to do so. By the same token, if a subject responds to her evidence for P by believing P but does so only because she likes where the evidence points to, then her belief still falls short of being competent. For, again, what we expect from a competent believer is that she responds to her evidence, not because of some distinct motivation to do so, but simply because it's the evidence.²⁵

What is epistemic competence? A paradigm example of such competence, I take it, is a subject's *theoretical reasoning* competence. In fact, if we understand "theoretical reasoning" broadly enough—including, e.g., transitions from non-doxastic mental states to beliefs—we might even identify epistemic competence with theoretical reasoning competence. At any rate, someone who manifests reasoning competence in forming a belief—i.e., someone who is reasoning *well*—is plausibly someone who is *unconditionally* responsive to her possessed sufficient evidence (in appropriate conditions). That is, for a competent reasoner, there is typically no gap between recognizing that some consideration positively settles the question as to whether P and believing P—"no gap" in the sense that there is neither need nor room for the motivational work of an independent desire (a desire to believe the truth, say) to take one from the recognition that the case for P's truth is settled to actually believing P.²⁶ Plausibly, then, a belief that manifests epistemic competence is—at least paradigmatically—one that has been formed through *good* reasoning.

I propose to capture the distinctive character of believing well (i.e., competently) as follows:

Believing Competently (BC): A competent belief is one that manifests epistemic competence, where epistemic competence is (at least in part) a matter of *unconditional* responsiveness to possessed considerations that bear (in a sufficiently weighty manner) on the correctness of one's (potential) beliefs, such that if one possessed (say) sufficient evidence E for P but didn't respond to E by believing P, then one would exhibit a *deficit* of epistemic competence (in appropriate conditions).²⁷

²⁵ In Horst (forthcoming) I argue that the distinctive character of epistemic competence undermines a prominent tendency in virtue epistemology to construe such competence on the model of skill, i.e., the sort of competence familiar from the domain of skillful action such as archery, baseball, or chess.

²⁶ Compare Raz (2011: 28) on what he calls the "No Gap Principle". For similar points, see also Hieronymi (2006: 51), Scanlon (1998: 33), and Velleman (2000: 278).

²⁷ I don't mean *BC* to suggest that we can *analyze* epistemic competence in terms of responsiveness to possessed truth-related considerations. Whether this is possible is—I think, rightly—controversial. See, e.g., Sosa/Sylvan (2018). *BC* only claims that there is a metaphysical connection between epistemic competence and the relevant sort of responsiveness; it doesn't claim that the latter takes explanatory priority over the former.

Let me add some comments and clarifications. First of all, it bears emphasizing that what matters for the question of whether someone possesses a certain competence is her performance in *appropriate* conditions. For example, the fact that someone lost control over his car on an icy road needn't count against his driving skills. Similarly, failure to appropriately respond to one's evidence needn't reveal a deficit of epistemic competence if one is drunk or distracted.²⁸

Second, it's only one's *possessed* evidence that's relevant to believing competently, not just *any* evidence (e.g., undiscovered pieces of evidence). That matters because one doesn't lack in epistemic competence for not responding to evidence one isn't even aware of. Plausibly, there are two conditions on possessing (or "having") evidence E for P: one must be *aware* of E (e.g., one must see that E) and one must (in some way) *recognize* that E is relevant to whether believing P would be correct.²⁹ Of course, how these conditions are to be fleshed out in any detail is a matter of much controversy.³⁰ For our purposes, what's important to stress is just that, as others have argued, possession of evidence need not be understood in terms of justified or reasons-based beliefs. Alternatively, for example, one might understand possession in terms of non-doxastic mental states (Schmidt 2019), animal knowledge (Sosa 2007), or competence (Sosa/Sylvan 2018). If so, proponents of *EC* can appeal to *BC*—and, thus, to the notion of possession—without compromising their attempt to account for epistemic reasons in terms of the activity of believing and its constitutive standard.

Third, it's only *sufficiently weighty* considerations which competent believers are expected to respond to by forming outright beliefs.³¹ This is important because one surely doesn't lack in epistemic competence for *not* believing P in the face of *outweighed* or *defeated* evidence for P. (I consider how the present account might be extended to *pro tanto* evidence in the next section.) Epistemic permissivists might insist on an even stricter formulation of *BC*, perhaps replacing "sufficient evidence" with "conclusive evidence". However, this wouldn't undermine the contrast with chess

²⁸ Compare Sosa (2015: 94-106).

²⁹ Both Lord (2018) and Sylvan (2016) stress the importance of such a recognition condition on possession. It's important to note that "recognition" need not be understood in terms of belief (e.g., believing that E is evidence for P). Less demanding forms of recognition are possible (e.g., recognition might be understood in *dispositional* terms). See, again, Lord (2018) and Sylvan (2016) for defending non-doxastic versions of the recognition condition.

³⁰ For relevant discussion, see, e.g., Schroeder (2011), Sylvan (2016), Lord (2018: chs. 3 and 4), Sylvan/Sosa (2018).

³¹ Another question I have to set aside here is whether—and if so, how—practical considerations can bear on when one's evidence counts as "sufficiently weighty", as proponents of pragmatic encroachment claim. This is legitimate because the core dispute between constitutivists and their opponents is somewhat orthogonal to the issues surrounding pragmatic encroachment. As noted before, the core dispute is about what, if anything, makes some consideration (paradigmatically, but perhaps not exclusively: evidence) an epistemic reason for believing some proposition. The dispute is not—at least not primarily—about whether only evidence can be an epistemic reason for belief, or about whether practical considerations can alter the strength of one's evidence.

competence. For, in the domain of chess, no matter how strong the considerations that recommend a certain chess-move, you may still disregard them without displaying a deficit of chess competence. But I take it that no epistemic permissivist will want to be *that* permissive with respect to the connection between evidence and epistemic competence.

Fourth, insofar as correctness-relevant considerations have weights and compete with one another, they possess the functional profile of reasons. But this by itself doesn't *make* them reasons (in the fully normative sense of "reason" at issue here). For, considerations that bear on the correctness of chess-moves, dance-moves, or cooking steps have weights and compete with one another as well. Yet, as we know, this isn't enough to turn them into reasons.³² So, I don't think *BC* is begging any questions by assuming that considerations which bear on the correctness of beliefs have weights and compete with one another.

One might still find *BC* overly demanding. Every day we encounter huge amounts of information. But, surely, we are not lacking in epistemic competence for not forming all the beliefs that would be supported by the daily flow of information. If so, one might wonder if *BC* isn't imposing unrealistic demands on what it takes to be a competent believer.³³ In response, it's important to emphasize again that what matters for assessments of epistemic competence is our responsiveness to the evidence we *possess*. And, plausibly, much of the information we encounter every day doesn't rise to the level of *possessed* evidence. For, recall that, in order to possess evidence E for P, one need not only be aware of E, one must also recognize E's relevance to whether believing P would be correct. However, if one *does* recognize E's relevance to whether believing P would be correct but fails to believe accordingly, then a negative appraisal of one's epistemic competence seems perfectly apt. So, I think, once we appreciate that it's *possessed* evidence, not just any evidence, to which epistemic competence is responsive, this should alleviate worries that *BC* is overly demanding.

Another worry is that the emphasis on *possessed* evidence might ultimately conflict with the purported *unconditional* character of a competent believer's responsiveness to her evidence. To see why, consider that, often enough, you come into possession of evidence only because you went looking for it, and you went looking for it only because you are inquiring into a question to which that evidence is relevant. Yet, which questions you are trying to answer is surely something that depends on your practical interests and goals. But if so, doesn't that mean that whether or not you are lacking in epistemic competence for not forming a certain belief in a particular situation is ultimately conditional on your practical interests and goals (at least on many occasions)?

³² On this point, see also Lord/Sylvan (2019).

³³ Thanks to Hille Paakkunainen for urging me to address this issue.

Here it's important to separate two issues. Whether you are *in a position* to competently answer the question of whether P—i.e., whether you possess sufficient evidence to decide that question—will often depend on your interests in finding out whether P. But once you *do* possess sufficient evidence to decide that question—sufficient evidence for P's truth, say—it's no longer up to your interests whether or not failure to respond to that evidence would reveal a deficit of epistemic competence on your part. That's precisely where epistemic competence differs from chess competence. If you refrain from making a certain chess-move in the face of considerations that clearly show this to be the correct move to make, then this needn't reflect badly on your chess competence—not if, for example, you are no longer interested in continuing your game of chess. By contrast, if you refrain from believing P in the face of considerations that clearly show P to be true, then this normally does reveal a deficit of epistemic competence on your part, regardless of whether or not you are still interested in finding out whether P.

Now, if these observations are on the right track, then it looks as if there is an important difference between the activities of believing and playing chess after all. Both activities give rise to crucially different patterns of evaluation. What counts as performing well (i.e., competently) in the one activity is *structurally* different from what counts as performing well in the other: while believing competently requires *unconditional* responsiveness to correctness-relevant considerations, playing chess competently requires only *conditional* responsiveness to such considerations. What we have to consider next, then, is why this difference matters to answering the challenge facing *EC*—why it makes a difference as to whether correctness-relevant considerations amount to epistemic reasons.

Before doing so, however, it will be instructive to briefly compare the present proposal with a popular alternative view on how to distinguish believing from other standard-governed activities. On this alternative view, the distinctive feature of the activity of believing is that engagement in believing is *constitutive* for creatures like us. Errol Lord and Kurt Sylvan suggest a view of this sort:

“The standard of correctness for belief is [...] surely not on all fours with norms of etiquette or some specific set of instructions for cooking *cacio e pepe*. But what could distinguish the standard of correctness for belief from these other standards other than the fact that ‘playing the game of belief’ is a constitutive feature of agency?” (Lord/Sylvan 2019: 64)

Presumably, on this view, believing is constitutive of agency in the sense that forming and having beliefs is *inescapable* for creatures like us. Unlike playing the game of chess, we can't opt out of “playing the game of belief” without ceasing to be the sort of creature that we are.³⁴ This might be true; but I doubt that putting the contrast between believing and playing chess in this way gets to

³⁴ For similar claims, see also Railton (1997: 59), Feldman (2001: 88), and McHugh (2012: 23). Compare Ferrero (2009) for the general strategy of defending constitutivism on the basis of the (alleged) inescapability of agency.

the heart of things. That's because *inescapability* doesn't explain why there should be anything wrong with a believer (*qua* believer) who was only *conditionally* responsive to belief's constitutive standard. After all, from the fact that some activity A is inescapable for you, it doesn't follow that by disregarding A's constitutive standard you necessarily demonstrate a *deficit* of competence at A. To see this, imagine a creature that's condemned to a life-long series of chess games—playing chess for that creature is inescapable. Still, at times, she might deliberately make a *bad* chess-move (e.g., when she has been bribed to do so).³⁵ Such occasional disregard for chess' standard of correctness needn't reveal any deficit of chess competence on her part—there need be nothing amiss with her *qua* chess player. Hence, I doubt that the notion of inescapability is well suited to capture what strikes me as the most salient difference between believing and playing chess: namely that performing well *qua* believer requires unconditional, not just conditional, responsiveness to belief's standard of correctness. Thus, contrary to what Lord and Sylvan's remark suggests, believing isn't just a *special kind of game*—one that, unlike the game of chess, you can't stop playing without ceasing to be an agent. Rather, the activities of believing and playing chess differ in terms of their *internal* evaluative structure. That is, as we saw, belief's standard of correctness grounds evaluations—i.e., assessments of competence—that have no *structural* analogue in the domain of chess and other games.

Furthermore, a common complaint against constitutivist views which invoke inescapability is that this feature is normatively irrelevant.³⁶ And I tend to agree—at least, it's not obvious why the fact that I can't stop engaging in some standard-governed activity should make a difference as to whether considerations that bear on how to meet that activity's standard amount to genuine reasons. By contrast, as I will argue now, distinguishing believing from other standard-governed activities in terms of unconditional responsiveness *is* normatively relevant.

7. Why the Difference Matters

Why should the fact that competent believers are unconditionally responsive to correctness-relevant considerations suggest that such considerations amount to reasons for belief? The answer I want to suggest is based on an independently plausible thought about epistemic reasons: i.e., that epistemic reasons just are the sort of considerations that we expect competent believers to be unconditionally responsive to. If so, then the fact that competent believers are unconditionally responsive to correctness-relevant considerations—which we captured in *BC*—is indeed

³⁵ As noted earlier (footnote 24), deliberately performing *bad* chess-moves is fully compatible with the (alleged) requirement that, in order to count as “playing chess”, one must intend to win.

³⁶ See, e.g., Côté-Bouchard (2016: 3189).

normatively relevant: it supports the claim that such considerations amount to reasons for belief. Let me elaborate on this suggestion.

Start with the popular idea that, just like other normative reasons, epistemic reasons must be capable of *guiding* us to the response they favor.³⁷ Plausibly, to be guided by an epistemic reason is a matter of responding to the relevant reason-giving consideration by way of an exercise of epistemic competence (e.g., by reasoning *well* from this consideration to the belief it supports). This means that, for every epistemic reason, there must be a possible exercise of epistemic competence (e.g., a possible instance of reasoning *well* from this reason to the belief it supports). Thus, suppose consideration C is a sufficient epistemic reason for you to believe P. If so, responding to the possession of C by believing P is a possible exercise of epistemic competence, such that if you were suitably epistemically competent and you possessed C, we can expect you to respond to C by forming the belief that P (absent interfering factors such as drunkenness or distraction).³⁸ In other words, a consideration C is a sufficient epistemic reason for you to believe P only if you would exhibit a *deficit* of epistemic competence for not responding to your possession of C by believing P.

Furthermore, there is the similarly intuitive thought that, in appropriate conditions, epistemic competence enables its subject to form beliefs in ways that are responsive to the demands of epistemic reasons. After all, it would be quite mysterious why we cared about competent belief-formation if forming beliefs in that way would systematically lead us astray with respect to what there is epistemic reason to believe. But then, if you form a belief through an exercise of epistemic competence (e.g., by reasoning *well*) we can expect that you normally do so in response to considerations that amount to epistemic reasons for that belief—at least if the relevant considerations are *true* or *facts*. (The restriction to facts or true considerations is needed if we conceive of epistemic reasons as facts, as some philosophers do. For, plausibly, there are situations in which you would exhibit a *deficit* of epistemic competence for not responding to some *false* consideration. If so, and if epistemic reasons are indeed facts, we need the restriction in order to avoid incorrect predictions.) Another way of putting this second intuitive thought, then, is to say that a (true) consideration C is a sufficient epistemic reason for you to believe P if you would

³⁷ Something like this idea is endorsed by Way (2017), Silverstein (2016), Setiya (2014), Raz (2011), Shah (2006), Hieronymi (2005), Kolodny (2005), and many others.

³⁸ Note that this doesn't tie epistemic reasons to *ideal* or *perfect* epistemic competence. This is important because, arguably, what we have reason to believe will sometimes depend on our *imperfections* as believers. For discussion of the analogous—more prominent—problem facing views that tie practical reasons to *ideal* rationality, virtue, or the like, see, e.g., Johnson (1999), Markovits (2014), and Setiya (2014).

exhibit a *deficit* of epistemic competence for not responding to your possession of C by believing P.

We can use these two points to formulate the following account of what it is for some consideration to be a sufficient epistemic reason for belief:

Epistemic Reason (ER): For a (true) consideration C to be a sufficient epistemic reason for you to believe P is for C to be such that, if you possessed C but didn't respond to C by believing P, then you would exhibit a *deficit* of epistemic competence (in appropriate conditions).

ER—or something close enough—strikes me as a plausible claim about the connection between epistemic reasons and epistemic competence, one that draws on familiar and widely shared intuitions about epistemic reasons and normative reasons more generally.³⁹ Certainly, as noted, the intuitions underlying *ER* aren't unique to proponents of *EC*. However, this is not the place to launch a full-fledged defense of this claim.⁴⁰ My aim here is just to point out that, if something like *ER* is indeed correct, then the normative relevance of *BC* should be clear. According to *BC*, the sort of considerations which are such that, if you failed to respond to their possession, you would exhibit a *deficit* of epistemic competence, are considerations that bear on the correctness of your beliefs. And so, together with *ER*—which identifies sufficient epistemic reasons with the sort of considerations which are such that, if you failed to respond to their possession, you would exhibit a *deficit* of epistemic competence—*BC* entails that considerations that bear positively on the correctness of your beliefs are sufficient epistemic reasons for that belief. In other words, assuming *ER*, *BC* is normatively relevant because it tells us something about which considerations are epistemic reasons.

Importantly, this argument rests on the assumption that the notion of epistemic competence at work in *ER* is plausibly the same as the one at work in *BC*. But why think that? Why think that the notion of epistemic competence relevant to determining which considerations are sufficient

³⁹ At least in spirit, *ER* resembles the so-called “reasoning view”, according to which normative reasons in general are premises of possible instances of good reasoning. See, e.g., Hieronymi (2005), Setiya (2014), Silverstein (2016), Way (2017), and Asarnow (2017). To be sure, however, the reasoning view and normative constitutivism are distinct projects: one might accept a conception of normative reasons as premises of good reasoning, without endorsing a constitutivist account of good reasoning (i.e., one where the goodness of reasoning is understood in terms of the standards constitutive of agents and believers).

⁴⁰ Some have challenged claims in the vicinity of *ER* by presenting counterexamples: i.e., purported cases where epistemic reasons and competent/good theoretical reasoning come apart. See Markovits (2014:41f.) and Schmidt (2020). Another important challenge is to say how, if at all, an account like *ER* can deal with the *defeasibility* of epistemic reasons. For relevant discussion, see, e.g., Setiya (2014), Silverstein (2016), Way (2017), and Asarnow (2017). I will briefly come back to this point below.

epistemic reasons should be understood along the lines of *BC*—i.e., in terms of performing well *qua* believer? I think one way to support this assumption is to highlight the intuitively correct verdicts that we get from *ER* if we conceive of epistemic competence along the lines of *BC*. To bring this out, consider again the case where you acquire clear evidence that your partner is cheating on you. Intuitively, such evidence is a sufficient epistemic reason for you to believe that your partner is cheating on you. And that’s exactly the verdict we get if we combine *ER* with a construal of epistemic competence along the lines of *BC*. For, according to *BC*, if you refrain from believing *P* despite possessing clear evidence for *P*’s truth, then this reveals a *deficit* or *failure* of epistemic competence on your part. In other words, if we combine *ER* with a construal of epistemic competence along the lines of *BC*, cases in which you have clear evidence for *P*’s truth will be cases in which you possess sufficient epistemic reason to believe *P*. And this is surely the correct verdict.

To support this further, consider how things would be if we *didn’t* construe epistemic competence along the lines of *BC*, but instead used the sort of competence familiar from the domain of skillful action—such as playing chess, making risotto, or dancing tango—as a model. On this construal, it would be unclear why it should reveal any *deficit* or *failure* of epistemic competence if you refrained from believing that your partner is cheating on you in the face of clear evidence that they are. After all, as we know, it doesn’t necessarily reveal any deficit of chess competence if you refrain from doing what you know would be the correct thing to do, relative to the standard of good chess. Consequently, if we were to combine *ER* with a construal of epistemic competence on the model of chess competence, cases in which you possess clear evidence for *P*’s truth wouldn’t necessarily be cases in which you possess epistemic reason to believe *P*. And this is clearly the wrong verdict.

What this suggests, then, is that *ER* delivers intuitively correct verdicts only if we understand the notion of epistemic competence at work in that claim along the lines of *BC*. In other words, it’s only because of the *distinctive* character of competent epistemic performance—i.e., the fact that such performance is *unconditionally*, not just conditionally, responsive to its own correctness—that we can extract conclusions about what there is reason to believe from competent belief-formation. As long as we are in the dark about how believing competently differs from other kinds of competent performance (like playing chess competently), it’s hard to see how *ER* can be true.

Before concluding, let me highlight one other important question: i.e., the question of how (if at all) the present view can be extended to *pro tanto* epistemic reasons. Suppose you have *some*, but not *sufficient*, reason to believe *P*. If so, responding to that reason by believing *P* would surely *not* be a competent response. This means that there are epistemic reasons—namely, *pro tanto* ones—that competent believers will *not* respond to by forming the beliefs they support. That’s why *ER* is a claim specifically about *sufficient* epistemic reasons, not about epistemic reasons in general. But

that's a limitation, and one might wonder whether there is a way to amend *ER* (and *BC*) so as to accommodate the existence of *pro tanto* reasons.

Perhaps there is. To see how this might work, note that we can distinguish between *competent* and *inept* ways to respond to *pro tanto* reasons as well. It's just that these responses will not consist in one's forming—or failing to form—an outright belief. Thus, when you have a *pro tanto* reason to believe P, the competent response to his reason may consist in an increase of confidence in P.⁴¹ If so, you would exhibit a deficit of epistemic competence if you didn't respond to this reason in that way. This suggests that we can improve on *ER* by distinguishing more carefully between *different* manifestations of epistemic competence, corresponding to the different weights of epistemic reasons. We could then amend *ER* by adding conditions for specifically *pro tanto* epistemic reasons. Very roughly, this would yield an amendment along the following lines: a *pro tanto* epistemic reason for you to believe P is a (true) consideration such that, if you possessed that consideration but didn't respond to it by an increase of confidence in P, you would exhibit a deficit of epistemic competence. *BC* could be complemented accordingly.

8. Conclusion

Proponents of *EC* seek to explain epistemic reasons in terms of the activity of believing and its constitutive standard of correctness. I have argued that this view faces an important challenge: to explain what's special about believing such that, unlike other standard-governed activities, it can give rise to reasons. According to many critics, *EC* must fail to meet this challenge. This is because, in relevant respects, believing is said to be just like playing chess: considerations bearing on how to meet belief's standard of correctness amount to reasons for particular beliefs only if and because there is a *prior* reason to comply with that standard. Yet, as I have argued, this assumption is flawed. For, intuitively, there are many cases in which you have a reason to believe P—provided by considerations that bear on the correctness of believing P—even though you *lack* a prior reason to comply with belief's standard of correctness.

However, this by itself doesn't show how, if at all, proponents of *EC* can positively meet the challenge facing their view. Doing so requires more than the dismissal of the requirement for prior reasons; it requires the identification of a normatively relevant difference between believing and activities like playing chess—one that makes a difference as to whether considerations that bear on the correctness of particular beliefs amount to reasons for that belief. In the last two sections, I have argued for a proposal of how *EC* can accomplish this task. Roughly, on this view, the activity

⁴¹ See Setiya (2014: 234) for a suggestion along these lines.

of believing is special in that competent participants in that activity—those who perform well *qua* believers—are *unconditionally* responsive to considerations bearing on how to meet its standard of correctness. This, moreover, is a difference that makes a difference as to whether these considerations amount to epistemic reasons. For, assuming *ER*, epistemic reasons just are the sort of considerations in response to which competent believers form their beliefs.

In sum, then, on the resulting view, epistemic reasons are considerations that figure in competent belief-formation, where epistemic competence is understood as performing well *qua* believer—i.e., in terms of unconditional responsiveness to considerations bearing on how to meet belief's constitutive standard. Of course, this is but a sketch of a view, and there is significantly more work to do for proponents of *EC* to turn this into a full-fledged constitutivist account of epistemic reasons for belief. But again, my aim here was not to develop such an account—something which would surely require more than a single paper—but rather to defend *EC* in the face of the objection that no such account is to be had. What I hope to have accomplished, then, is a vindication of *EC* as a serious contender for an account of epistemic reasons for belief—at least one that can't be dismissed on the basis of an analogy with games like chess.⁴²

References

- Alston, W. (2005). *Beyond "Justification": Dimensions of Epistemic Evaluation*. Ithaca: Cornell University Press.
- Asarnow, S. (2017). 'The Reasoning View and Defeasible Practical Reasoning,' *Philosophy and Phenomenological Research* 95, pp. 614-636.
- Boghossian, P. (2003). 'The normativity of content,' *Philosophical Issues*, 13, pp. 31-45.
- Côte-Bouchard, C. (2016). 'Can the aim of belief ground epistemic normativity?' *Philosophical Studies*, 173, pp. 3181-3198.
- Cowie, C. (2014). 'In defence of instrumentalism about epistemic normativity,' *Synthese* 191(16), pp. 4003-4017.
- Enoch, D. (2006). 'Agency, shmagency: Why normativity won't come from what is constitutive of action,' *The Philosophical Review*, 115, pp. 169-198.
- Enoch, D. (2011). 'Shmagency revisited,' in M. Brady (ed.), *New Waves in Metaethics*. London: Palgrave-Macmillan.
- Feldman, R. (2001). 'Voluntary belief and epistemic evaluation,' in M. Steup (ed.), *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue*. Oxford: Oxford University Press.

⁴² For comments on previous drafts and/or helpful discussions, I am thankful to David Hunter, Luca Ferrero, Hille Paakkunainen, Karl Schafer, Sergio Tenenbaum, Josh Thorpe, and two anonymous referees for this journal who provided exceptionally detailed and constructive feedback. Thanks also to audiences at the 2021 APA Central Division meeting, the Brazilian Society for Analytic Philosophy Conference at the Federal University of Bahia in Salvador, and the University of Campinas.

- Ferrero, L. (2009). 'Constitutivism and the inescapability of agency,' *Oxford Studies in Metaethics*, 4, pp. 303–333.
- Foley, R. (1987). *The Theory of Epistemic Rationality*. Cambridge: Cambridge University Press.
- Goldman, A. I. (1999). *Knowledge in a Social World*. Oxford: Oxford University Press.
- Grimm, S. R. (2009). 'Epistemic normativity,' in A. Haddock, A. Millar, & D. Pritchard (eds.), *Epistemic value*. Oxford: Oxford University Press.
- Hieronymi, P. (2005). 'The Wrong Kind of Reason,' *Journal of Philosophy* 102, pp. 437–57.
- Hieronymi, P. (2006). 'Controlling Attitudes,' *Pacific Philosophical Quarterly* 87(1), pp. 45-74.
- Horst, D. (forthcoming). 'Is Epistemic Competence a Skill?' *Australasian Journal of Philosophy*, early view: DOI: [10.1080/00048402.2021.1912125](https://doi.org/10.1080/00048402.2021.1912125)
- Johnson, R. N. (1999). 'Internal Reasons and the Conditional Fallacy,' *Philosophical Quarterly* 49, pp. 53-72.
- Katsafanas, P. (2018). 'Constitutivism and Practical Reason,' in D. Star (ed.), *The Oxford Handbook of Reasons and Normativity*. Oxford: University Press.
- Kelly, T. (2003). 'Epistemic rationality as instrumental rationality: A critique,' *Philosophy and Phenomenological Research* 66, pp. 612–640.
- Kiesewetter, B. (2021). 'Are epistemic reasons normative?' *Noûs* early view, pp. 1-26.
- Kolodny, N. (2005). 'Why be rational?' *Mind* 114, pp. 509–563.
- Kornblith, H. (1993). 'Epistemic Normativity,' *Synthese* 94, pp. 357-376.
- Korsgaard, C. (2008). *The Constitution of Agency*. Oxford: Oxford University Press.
- Lavin, D. (2017). 'Forms of Rational Agency,' *Royal Institute of Philosophy Supplement* 80, pp. 171-93.
- Leary, S. (2017). 'In Defense of Practical Reasons for Belief,' *Australasian Journal of Philosophy* 95, pp. 529–542.
- Leite, A. (2007). 'Epistemic Instrumentalism and Reasons for Belief: A Reply to Tom Kelly,' *Philosophy and Phenomenological Research* 75, pp. 456-464.
- Littlejohn, C. (2018). 'Reasons and Theoretical Rationality,' in D. Star (ed.), *Oxford Handbook of Reasons and Normativity*. Oxford: Oxford University Press.
- Lockard, M. (2013). 'Epistemic instrumentalism,' *Synthese* 190, pp. 1701-18.
- Lord, E. (2018). *The Importance of Being Rational*. Oxford: Oxford University Press.
- Lord, E. and Sylvan, K. (2019). 'Reasons: Wrong, Right, Normative, Fundamental,' *Journal of Ethics and Social Philosophy* 15, pp. 43-74.
- Lord, E. and Sylvan, K. (2021). 'Suspension, Higher-Order Evidence, and Defeat,' in M. Simion & J. Brown (eds.), *Reasons, Justification, and Defeat*. Oxford: Oxford University Press.
- Lynch, M. P. (2004). *True to Life: Why Truth Matters*. Cambridge, Mass.: MIT Press.
- Markovits, J. (2014). *Moral Reasons*. Oxford: Oxford University Press.
- McHugh, C. (2012). 'The truth norm of belief,' *Pacific Philosophical Quarterly*, 93, pp. 8–30.
- Paakkunainen, H. (2018). 'Doubts about 'genuinely normative' epistemic reasons,' in C. McHugh, J. Way & D. Whiting (eds.), *Metaepistemology*. Oxford: Oxford University Press.
- Parfit, D. (2011). *On what matters*. OUP.

- Railton, P. (1997). 'On the hypothetical and non-hypothetical in reasoning about belief and action,' in G. Cullity & B. N. Gaut (eds.), *Ethics and practical reason*. Oxford: Oxford University Press.
- Raz, J. (2011). *From normativity to responsibility*. Oxford: Oxford University Press.
- Rinard, S. (2015). 'Against the New Evidentialists,' *Philosophical Issues* 25, pp. 208–223.
- Rosen, G. (2001). 'Brandom on modality, normativity, and intentionality,' *Philosophy and Phenomenological Research* 63, pp. 611–23.
- Scanlon, T. (1998). *What we owe to each other*. Cambridge, Mass.: Harvard University Press.
- Schafer, K. (2018). 'Constitutivism about Reasons,' in K. Jones and F. Schroeter (eds.), *The Many Moral Rationalisms*. Oxford: Oxford University Press.
- Schmidt, E. (2020). 'Where reasons and reasoning come apart,' *Noûs*, early view, pp. 1-20.
- Schmidt, E. (2019). 'Possessing epistemic reasons: the role of rational capacities,' *Philosophical Studies* 176, pp. 483-501.
- Schroeder, M. (2011). 'What Does It Take to 'Have' a Reason?' in A. Reisner & A. Steglich-Petersen (eds.), *Reasons for Belief*. Cambridge: Cambridge University Press.
- Setiya, K. (2014). 'What is a Reason to Act?' *Philosophical Studies*, 167, pp. 221–235.
- Shah, N. (2003). 'How Truth Governs Belief,' *The Philosophical Review*, 112, pp. 447-482.
- Shah, N. (2006). 'A New Argument for Evidentialism,' *Philosophical Quarterly*, 56, pp. 481–498.
- Sharadin, N. (2018). 'Epistemic instrumentalism and the reason to believe in accord with the evidence,' *Synthese* 195, pp. 3791–3809.
- Silverstein, M. (2016). 'Reducing Reasons,' *Journal of Ethics and Social Philosophy*, 10, pp. 1–22.
- Smith, M. (2013). 'A constitutivist theory of reasons: Its promises and parts,' *Law, Ethics, and Philosophy*, 1, pp. 9-30.
- Sosa, E. (2007). *A Virtue Epistemology: Apt Belief and Reflective Knowledge*. Oxford: Oxford University Press.
- Sosa, E. (2010). *Knowing Full Well*. Princeton: Princeton University Press.
- Sosa, E. (2015). *Judgment and Agency*. Oxford: Oxford University Press.
- Sosa, E. & Sylvan, K. (2018). 'The place of reasons in epistemology,' in D. Star (ed.), *The Oxford Handbook of Reasons and Normativity*. Oxford: Oxford University Press.
- Steglich-Petersen, A. (2011). 'How to be a Teleologist about Epistemic Reasons,' in A. Reisner & A. Steglich-Petersen (eds.), *Reasons for Belief*. Cambridge: Cambridge University Press.
- Sylvan, K. (2016). 'Epistemic Reasons I: Normativity,' *Philosophy Compass*, 11, pp. 364-376.
- Thomson, J. (2008). *Normativity*. Chicago: Open Court.
- Tubert, A. (2010). 'Constitutive arguments,' *Philosophy Compass*, 5, pp. 656–666.
- Velleman, D. (2000). *The Possibility of Practical Reason*. Oxford: Carendon Press.
- Way, J. (2017). 'Reasons as Premises of Good Reasoning,' *Pacific Philosophical Quarterly* 98, pp. 251-270.
- Wedgwood, R. (2002). 'The aim of belief,' *Philosophical Perspectives* 16, pp. 267–297.
- Wiland, E. (2012). *Reasons*. London: Continuum.