

Epistemic Normativity Without Epistemic Teleology

Benjamin Kiesewetter, Bielefeld University

This is the pre-copyedited version of an article whose final and definitive form has been published in *Philosophical Issues* 34 (2024), 349-370. The final article is available (open access) at:

<http://dx.doi.org/10.1111/phis.12274>

Abstract: This article is concerned with a puzzle that arises from three initially plausible assumptions that form an inconsistent triad: (1) Epistemic reasons are normative reasons (normativism); (2) reasons are normative only if conformity with them is good (the reasons/value-link); (3) conformity with epistemic reasons need not be good (the nihilist assumption). I start by defending the reasons/value-link, arguing that normativists need to reject the nihilist assumption. I then argue that the most familiar view that denies the nihilist assumption – epistemic teleology – is untenable. Finally, I consider two alternative ways of accounting for the goodness of conformity with epistemic reasons: it may be good because it accords with the virtue of reasons-responsiveness, and it may be good because it is good to conform with normative reasons as such. I argue that both of these conceptions avoid the problems of epistemic teleology and merit serious consideration as potential solutions to the puzzle.

§0 INTRODUCTION

Suppose that you are waiting for an appointment with your physician, and in order to kill time, you grab a celebrity magazine that is offered in the waiting room. The magazine contains a voyeuristic report about the skin problems of Stella Starlet, a B-List celebrity you haven't heard of before and will likely never hear of again. The report and the accompanying pictures in the magazine provide you with excellent evidence for believing that Starlet has blemished skin, evidence that constitutes sufficient epistemic reason and thus epistemic justification for having the belief.

Call this case *Magazine*. It seems that, special circumstances aside, there need not be any value in having the epistemically supported belief in *Magazine*. Whether it's good to have a true or epistemically supported belief depends on whether we have reason to care (either instrumentally or

non-instrumentally) about its *content*, and other things being equal, we don't have reason to care about whether Stella Starlet has blemished skin. By contrast, our epistemic reasons don't seem to depend on whether we have reason to care about a certain content; they only depend on the presence and salience of evidence and the strength of support that it provides.

Cases like this provide a puzzle for *normativism about epistemic reasons* – the widely held view that epistemic reasons are normative reasons.¹ How could we have normative reasons for believing contents that do not matter, reasons that normatively support beliefs we have no reason to care about? The puzzle arises from a plausible link between normative reasons and value, according to which conformity with normative reasons must be good, at least to some extent. On the assumption that epistemic reasons are normative, this *reasons/value-link* entails that conformity with epistemic reasons must be good to some extent. But this implication seems hard to defend in cases in which our epistemic reasons support beliefs about entirely uninteresting matters. We thus have an inconsistent triad:

- (i) *Normativism about epistemic reasons*: Epistemic reasons are normative reasons.
- (ii) *The reasons/value-link*: Reasons are normative only if conformity with them is good to some extent.
- (iii) *The nihilist assumption*: Conformity with epistemic reasons need not be good to any extent.

Normativism is supported by a number of considerations. One argument that I find particularly forceful is that normativism promises to explain why epistemic reasons have so many features in common with normative reasons. Epistemic reasons provide a partial justification for what they are reasons for, they make the responses they are reasons for rational, they provide premises for good reasoning to, and they are good bases (or motivating reasons) for the responses they are reasons for. All of this is true for normative reasons in general. Normativism thus provides an excellent explanation of why epistemic reasons have these features, and it is difficult to see how there could be a similarly good explanation that doesn't entail normativism.

¹ Normativists include, among many others, Scanlon (1998, 18–19); Kelly (2007a); Grimm (2009); Raz (2009); Skorupski (2010, Ch. 2); Parfit (2011, esp. Chs. 1–5); Berker (2013); Talbot (2014); Gregory (2016); Schroeder (2021); McHugh and Way (2022). Normativism is often taken for granted by its proponents, but it has recently been called into question; see e.g. Rinard (2015); Glüer and Wikforss (2018); Mantel (2019); Maguire and Woods (2020). For a defence of normativism against these attacks, see Kiesewetter (2022b).

I have defended this and other arguments for normativism elsewhere (Kiesewetter, 2022b) and will here take this view for granted. From the normativist perspective, the inconsistency poses a dilemma: normativists either have to reject the reasons/value-link or the nihilist assumption. In Section §1, I argue that denying the reasons/value-link is unattractive, as it leaves us with the perplexing conclusion that we might have no reason to care about our conformity with normative reasons. Section §2 discusses *epistemic teleology*, the most familiar view that denies the nihilist assumption, which claims that epistemic reasons are based on (either the instrumental or the intrinsic) value of having true or epistemically supported beliefs. I show that epistemic teleology is subject to a variety of problems and argue that it is ultimately indefensible. This leaves normativists with the challenge to account for the reasons/value-link without assuming epistemic teleology or any of the claims that make it untenable. I address this challenge in section §3, by way of discussing two alternative accounts of the goodness of conforming to epistemic reasons. On the first view, conformity with epistemic reasons is good because it accords with the virtue of reasons-responsiveness. On the second view, it is good because conformity with normative reasons is good as such. While both of these views come with substantive commitments, I argue that they avoid the problems of epistemic teleology and should be taken seriously as ways of resolving the dilemma.

§1 THE REASONS/VALUE LINK

Let me start by explaining in some more detail the three central notions that I will be concerned with in this essay: epistemic reasons, normative reasons, and value. I use the technical term ‘epistemic reason’ to refer to the ordinary notion of a reason for (or against) belief, as it is used in contexts in which pragmatic and other so-called ‘wrong kind of reasons’ for (or against) belief are bracketed.² I assume that epistemic reasons in this sense are essentially connected to the constitutive function of belief, which involves representing certain contents as true, and at least paradigmatically constituted by evidence for (or against) the truth of the belief’s content.³ For simplicity, I shall disregard ‘wrong kind of reasons’ for attitudes in this essay – unless explicitly stated otherwise, the term ‘reason’ refers to a reason ‘of the right kind’.

By a ‘normative reason’, I mean a reason that *counts in favour* (or against) a response, such as an action or attitude. Normativists claim that epistemic reasons are a subset of the more general

² See Gertken and Kiesewetter (2017) for an overview of the right-kind/wrong-kind distinction.

³ For reasons to think that there are at least some non-evidential epistemic reasons, see Schroeder (2012).

class of reasons that support responses (or count against them) – a class commonly taken to include prudential and moral reasons for action. Normative reasons need to be distinguished from explanatory reasons (reasons that explain *why* something is the case), motivating reasons (reasons *for which* we respond), and also from formal standards like game rules or etiquette norms.

Further, I use ‘value’ as shorthand for ‘positive value’ and I treat ‘having value’ as equivalent to ‘being good in some way and to some extent’. The kinds of values I am interested in are the kinds that we have *unconditional reason to care about*; something is good insofar as there are reasons (for suitably related agents) to have pro-attitudes about it, such as desire, admiration, or appreciation.⁴ I take this to include predicative or ‘impersonal’ value (good *simpliciter*), prudential or ‘personal’ value (good *for*), as well as person-relative value (good from a person’s point of view), but not the kinds of goodness that are relative to purposes, kinds, or arbitrary standpoints. That something is good for a certain purpose, for example, does not correspond to unconditional reasons to care about it; at best, it corresponds to reasons that are conditional on a further reason to pursue the purpose in the first place.⁵ By contrast, it seems natural to think that impersonal value corresponds to agent-neutral reasons to care, person-relative value corresponds to agent-relative reasons to care, and personal value corresponds to reasons to care *for the sake of* a person (where it is a debatable, first-order normative question whether these reasons are agent-neutral or agent-relative).

With these clarifications at hand, let me come back to my opening example *Magazine*, in which you end up with epistemic reasons for believing that Stella Starlet has blemished skin, while intuitively, there is nothing good about having a true belief about that matter. If there is nothing good about a certain response that you could give, then (unless we assume that there is something bad about it, which we can set aside for now) you have *no reason to care* about whether you give that response. This much follows from the assumption that reasons to care correspond with facts about what’s good (or bad) in the relevant sense. Thus, in cases like *Magazine*, it looks like you might have no reason to care about whether you conform with your epistemic reasons. If normativism is true, it follows that there are cases in which you have no reason to care about whether you conform with your normative reasons. Call this the *perplexing conclusion*.

⁴ Note that this assumption is neutral on the questions of (i) whether reasons for pro-attitudes are normative, and (ii) whether reasons for pro-attitudes are explanatorily prior to goodness/value or not. It therefore seems to me an unproblematic assumption.

⁵ Compare Scanlon’s suggestion that judgments about attributive and purpose-relative goodness can be understood in terms of “*hypothetical practical reasoning*” (2011, 447).

Perhaps the perplexing conclusion can be avoided if we adopt epistemic permissivism, the view that epistemic reasons never require but merely permit the adoption of beliefs.⁶ Permissivists might maintain that since lack of belief never amounts to a failure to conform with epistemic reasons, the fact that there is nothing good about forming an evidentially supported belief does not entail that we have no reason to care about conformity with epistemic reasons. This would follow only if there was nothing good about refraining from insufficiently supported beliefs, as this is all that conformity with epistemic reasons actually requires.

This reply is based on two substantial assumptions, both of which seem questionable, however. The first of these assumptions is that there is disvalue in having false (or unwarranted) beliefs, independently of their content. This seems somewhat more plausible than assuming that true beliefs are valuable independently of their content but is, on reflection, hard to sustain. Even though false beliefs are often instrumentally bad, they need not be: there can be isolated, uninteresting false beliefs that do not have any negative effects, and there can even be heuristically useful false beliefs. And the assumption that false beliefs are independently of their content *intrinsically* bad seems to entail that we have practical reasons to invest resources in correcting our false beliefs even if we have no reason to care about their contents and even if they are heuristically useful and otherwise harmless, which seems doubtful.⁷

The second assumption is that it is epistemically permissible to refrain from adopting a justified belief even if one explicitly attends to the relevant question. This seems doubtful as well. If you are explicitly focusing on the question of whether Stella has blemished skin, and if what you know from reading the magazine clearly supports the belief that she has blemished skin, it seems plausible to think that conformity with epistemic reasons requires you to believe that Stella has blemished skin. But this doesn't show that there is any good in believing this. Thus, in such a case, you would seem to have no reason to care about your conformity with normative reasons.

How bad is the perplexing conclusion? Some normativists might hold that it is a confusion to think that just because we have reason to ϕ , we also need to have reason to *care* about whether we ϕ – the latter being an *additional evaluative assumption* that is plausible for some but not all

⁶ See e.g. Nelson (2010). I raise some doubts about this view in Kiesewetter (2019, 246–48).

⁷ The assumption that false beliefs have disvalue is also subject to the direction of fit problem that I will discuss later on in §2 (cf. Raz, 2009, 45). If false beliefs are not intrinsically bad, what explains why it's more plausible to think that false beliefs have content-independent disvalue than to think that true beliefs have content-independent value? My conjecture is that this is explained by the fact that uninteresting false beliefs are more likely to have bad effects than uninteresting true beliefs are likely to have good effects.

responses that are favoured by reasons. By extension, one might also argue that it is a confusion to think that we must have reason to care about whether we ϕ if conformity with reasons *requires* us to ϕ .

But is this response tenable? Reasons to ϕ and reasons to care about whether one ϕ -s are logically independent, but this doesn't show that they are normatively independent. And on the face of it, it's a plausible principle that we always have reasons to care whether we conform to our normative reasons.⁸ Indeed, one might think that to deny this principle would rob normative reasons of their significance. Surely if conformity with normative reasons requires us to ϕ , it must matter whether we ϕ , but, as Parfit has it, "something matters only if we ... have ... reason to care about this thing".⁹ If that is so, normativists cannot rest content with the perplexing conclusion.

One can care about something in virtue of having positive attitudes (like desire or appreciation) or negative attitudes (like aversion or dismissal) towards it. If we have reason to care about whether we conform with normative reasons, then we surely have reason to care about it in the positive rather than negative way – in other words, we have reasons for pro-attitudes (rather than con-attitudes) towards our conformity with normative reasons. On the basis of this assumption, we get the following argument for the reasons/value-link:

1. If conformity with normative reasons requires A to ϕ , then there are reasons for A to have pro-attitudes towards A's ϕ -ing.
2. If there are reasons for A to have pro-attitudes towards A's ϕ -ing, then A's ϕ -ing would be good to some extent.

The reasons/value-link: If conformity with normative reasons requires A to ϕ , then A's ϕ -ing would be good to some extent.

Conformity with normative reasons requires A to ϕ if ϕ -ing is supported by decisive reasons, but not necessarily if ϕ -ing is supported by sufficient or merely pro tanto reasons. Stronger versions of

⁸ Note that this assumption does not say that we always have *decisive* reasons for pro-attitudes about conformity with normative reasons, which would give rise to a regress.

⁹ Parfit (2011, 148). See also Kornblith (2002, 145): "If you tell me that a belief of mine is unjustified, this gives me reason to give up that belief. The epistemic claim is something about which I should care, and an account of the source of epistemic norms must explain why it is that I should care about such things."

the reasons/value-link would maintain that conformity with sufficient or pro tanto normative reasons must also be good.¹⁰ For all I argue in this essay, these stronger claims might also be true, but they don't seem to be equally well-supported by the principle that we always have reason to care about whether we conform to normative reasons. On the weakest and most plausible reading of this principle, we have reasons to care about conforming to decisive reasons, while we may be indifferent to our conformity with defeated or merely sufficient reasons. I will not assume anything stronger than that. Unless explicitly mentioned otherwise, "conforming to reasons" refers to conforming to decisive reasons in what follows.

One might wonder whether premise (1) is plausible in cases in which ϕ -ing is the lesser evil (such as, e.g., killing in self-defence), or simply the omission of something bad (such as, e.g., the omission of punching someone in the face). In such cases, we don't seem to have reason to desire that we ϕ for its own sake, but we still seem to have reason to desire that we ϕ as a means or pre-condition. As the term is used here, desiring something as a means or pre-condition counts as a pro-attitude. On this reading, (1) remains unaffected by such cases.

However, it might be argued that such cases then undermine premise (2). To kill someone in self-defence or to omit punching someone in the face are ways of avoiding something bad, and the mere absence of something bad shouldn't be conflated with the presence of something good. If we have reasons for pro-attitudes towards these things, then such reasons do not entail the presence of positive value – the absence of negative value suffices. With respect to the reasons/value-link, this means that conformity with normative reasons need not have positive value in order for there to be reasons for pro-attitudes towards it, since the disvalue of non-conformity suffices for that.

In substance, I agree with this: conformity with normative reasons need not be good in a non-comparative sense; for all I want to claim, it's enough for non-conformity to be bad. However, in such cases conformity is still *comparatively* good, i.e. better than non-conformity. If something is comparatively better than all relevant alternatives in either the impersonal, the personal, or the person-relative senses of goodness distinguished above, there are reasons for pro-attitudes towards it. I will therefore count it as good in a sense relevant for the reasons/value-link.

Together with normativism, the reasons/value-link entails that conformity with epistemic reasons must be good to some extent, at least compared to relevant alternatives. But the nihilist

¹⁰ The previous formulation of the reasons/value-link, according to which "reasons are normative only if conformity with them is good to some extent" is ambiguous between these readings.

assumption should be taken to involve the claim that conformity with epistemic reasons need neither be non-comparatively good nor comparatively better than the relevant alternatives. For example, it should be taken to involve the claim that in *Magazine*, believing that Starlet has blemished skin is not better than having no belief about the matter. Therefore, normativism and the reasons/value-link are inconsistent with the nihilist assumption.

§2 EPISTEMIC TELEOLOGY

There is a family of views in meta-epistemology that deny the nihilist assumption: *teleological* views, which hold that epistemic norms are based on the value of beliefs. In this section, I will be concerned with the following view:

Epistemic teleology: There is an epistemic reason for A to believe p iff and because A's believing p has (a certain kind of) value.

The term 'epistemic teleology' is used in different ways in the literature and more frequently applied to teleological views about epistemic *requirements* and *permissions*. My distinctive aim here, however, is to investigate teleology about epistemic *reasons*, as I am interested in the normativity of such reasons. Epistemic teleology is also often spelled out in terms of *desires* rather than *values*.¹¹ Such views might be understood as instances of the present view on the assumption that the relevant value is conduciveness to the satisfaction of a desire of a certain kind.

Epistemic teleology preserves both normativism and the reasons/value-link. Moreover, it can be seen as an implication of normativism on the assumption of a more general value-based conception of normative reasons that some philosophers find attractive:

The value-based conception of normative reasons: There is a normative reason for A to ϕ iff and because A's ϕ -ing has value.

Two remarks are in order. Firstly, note that both of these views leave open whether the object of our reasons is derivatively or non-derivatively valuable, and also whether in case it is derivatively valuable, it is instrumentally valuable (i.e. in virtue of *promoting* some final value), or non-

¹¹ See Kelly (2007a) for an important critical discussion of such views.

instrumentally derivatively valuable (e.g. in virtue of *honouring* some final value).¹² While covering instrumentalist or consequentialist theories, according to which all normative reasons are reasons to promote valuable states of affairs, these views therefore do not entail such theories.¹³

Secondly, both of these views are *direct* value-based views, which require reasons for ϕ -ing to be grounded in the value of ϕ -ing itself. A broader conception of value-based reasons would allow a reason for ϕ -ing to be grounded in values that are only indirectly related to ϕ -ing, such as the value of adopting a rule that prescribes ϕ -ing. However, such indirect views cannot vindicate the reasons/value-link, which requires the goodness of *conformity*, and thus can be bracketed here. As I shall highlight below, however, some of the objections to direct value-based views of epistemic reasons apply to indirect views as well.

Epistemic teleology can helpfully be separated into the following three assumptions:

1. *The axiological assumption:* Whenever there is epistemic reason to believe p , believing p has a certain kind of value.
2. *The epistemic assumption:* Whenever believing p has this kind of value, there is epistemic reason to believe p .
3. *The explanatory assumption:* Whenever there is epistemic reason to believe p , this is *because* believing p has this kind of value.

In what follows, I will argue that each of these claims faces significant difficulties.

§2.1 *The axiological assumption*

Let's start with the axiological assumption. The view that epistemically supported beliefs are valuable can be spelled out in different ways. On the first conception, they are instrumentally good for the purpose of getting to know the truth, which in turn is instrumentally good for all sorts of other values. As Nozick puts it, "truth ... is useful for a very wide range of purposes – almost all – and hence will be desired and bring benefit (almost) no matter what our particular purposes might

¹² See Pettit (1989) for the distinction between honouring and promoting a value, and Hurka (2001, esp. Ch. 1) for the claim that honouring a value is valuable.

¹³ See e.g. Maguire (2016) for a consequentialist and Wedgwood (2018) for a non-consequentialist value-based conception. Some limit the value-based conception to practical reasons (e.g. Portmore, 2011, Ch. 3), but there is some pressure to extend it to other normative reasons (Portmore 2018), and both anti-normativists and epistemic teleologists often implicitly assume the view that *all* normative reasons have to be based on value (or desires).

be” (1993, 68). On the second conception, the value of epistemically supported beliefs derives from a specifically *epistemic* kind of value or goal, which is most commonly identified with true or correct belief (Foley, 1987; Alston, 1988; Goldman, 2001; Lynch, 2004; Wedgwood, 2017), but sometimes also with justified belief (Feldman, 2000), knowledge (Littlejohn, 2018), or understanding (Kvanvig, 2003). I will call the first view the *universal means conception* and the second view the *epistemic value conception*.

There can be little doubt that true belief and the like are often useful in the sense stressed by the universal means conception, and I agree with the proponents of the epistemic value conception that they can be valued for their own sake in many cases. The problem is that both their usefulness and intrinsic desirability is content-dependent, while epistemic reasons are content-independent. By the *content-independence* of epistemic reasons, I mean that evidence for *p* provides epistemic reasons for believing *p* independently of what *p* is about and whether it matters. In other words, one can have excellent epistemic reasons for beliefs that are foreseeably useless and of no intrinsic interest at all. This was the point of *Magazine*. Beliefs need neither be useful nor of intrinsic interest in order to be supported by epistemic reasons. Call this the *problem of content-independence*.¹⁴

One reply to this problem maintains that it isn’t a failure of epistemic rationality to refrain from forming evidence-supported beliefs that are trivial or uninteresting, and that it is therefore inappropriate to demand a teleologist explanation of a requirement to have such beliefs (Leite, 2007; Steglich-Petersen, 2011). But this reply is beside the point, given that we are here concerned with epistemic *reasons* rather than epistemic requirements. Let’s grant for the sake of the argument that one is required to believe a proposition supported by the evidence only if that proposition is of practical relevance.¹⁵ Even if you’re permitted in *Magazine* not to believe that Stella Starlet has blemished skin, you still have epistemic reasons for this belief, and this is what the teleologist cannot explain. For example, even if Starlet’s skin is of no practical relevance, and even if we grant that you are therefore permitted to refrain from forming any belief about that matter, you might nevertheless rationally form the belief that Starlet has blemished skin on the basis of your evidence,

¹⁴ This is sometimes called the “triviality problem”; see Grimm (2009) and Côté-Bouchard (2017) for helpful discussions.

¹⁵ See e.g. Nozick (1993, 86); Nelson (2010); Steglich-Petersen (2011, 26). As suggested above already, it seems to me more plausible to say that whether one ought to believe a proposition supported by one’s evidence depends on whether one *attends* to the relevant question rather than on the question’s practical relevance (cf. Kelly, 2007b, 468–69; Kiesewetter, 2017, 185–185).

citing that evidence as a good reason to believe what you believe – a reason that justifies your belief.¹⁶

As far as I can see, there are two ways of responding to this challenge on behalf of the teleologist, and both don't seem very appealing. The first is to insist that there really is intrinsic value in believing the truth or gaining knowledge or understanding, independently of the content of what is believed.¹⁷ The second response accepts that there is no content-independent intrinsic value in believing the truth, but stresses the role of believing the truth as a universal means. The idea is that believing the truth about some matter either directly helps to promote a value, or otherwise does so indirectly because beliefs are interconnected in such a way that believing some truth always promotes believing other truths and thus ultimately promotes believing a truth that is relevant for the attainment of a valuable end.¹⁸ As far as this second response is concerned, I agree with Sharadin (2018) that we have little reason to accept that believing the truth about just *any* matter does in fact promote the adoption of true beliefs about relevant matters. It's already difficult to see why this would need to hold in the case of a belief about Stella Starlet's skin problems, but it seems incredible to think that it is true of each and every one of the infinitely many trivial disjunctions that are supported by our evidence. Moreover, we can construct examples in which we exclude this possibility by way of stipulation. Consider a variant of *Magazine* in which you are reading about Stella Starlet's blemished skin immediately before you are given a general anaesthesia, and suppose that you are certain to forget everything that you have learned within the last minutes before the anaesthesia. In such a case, forming a true belief about the matter cannot possibly promote forming

¹⁶ Both Leite (2007, 462–63) and Steglich-Petersen (2011, 24–27) seem to agree. This is where Leite suggests to abandon a direct account in favour of an indirect one, and Steglich-Petersen – who originally set out to explain “how to be a teleologist about epistemic reasons” – ends up asserting that “it is doubtful whether epistemic reasons have normative ‘force’” (2011, 26), thus collapsing his teleological view into an anti-normativist one. More recently, Steglich-Petersen has questioned the content-independence of epistemic reasons, suggesting that evidence constitutes an epistemic reason only if it speaks to questions we have independent reason to inquire about (Steglich-Petersen and Skipper, 2020, 114). However, ordinary epistemic discourse clearly treats epistemic reasons as content-independent; this view thus avoids anti-normativism only at the cost of maintaining that ordinary epistemic discourse is fundamentally misguided in this respect. And while other (apparent) suppositions of epistemic discourse that are in tension with epistemic instrumentalism (such as the assumption that epistemic reasons are constituted by evidence) might be explained by the appeal to the assumption that epistemic discourse is elliptical (Steglich-Petersen and Skipper, 2020, 114–15), this assumption doesn't explain why we take evidence to justify belief even in contexts in which it is common ground that there is no reason to care about a certain matter.

¹⁷ See e.g. Kvanvig (2003, 41) and Lynch (2004, 46–48 and 55). Wedgwood does not seem to address the problem, but he also holds that correctness has content-independent value (2017, 231).

¹⁸ Schroeder (2007, 113–14) presents (though falls short of endorsing) a strategy along these lines on behalf of a desire-based theory of epistemic reasons. See Sharadin (2018, 3796–3801) for an instructive discussion.

other true beliefs that are useful for promoting values – but that doesn't affect your epistemic reasons at the moment of reading the magazine.

We are thus left with the assumption that true belief, knowledge, or something of that sort is intrinsically valuable, independently of its content. This would mean that we necessarily have reason to desire or care about believing the truth about any matter whatsoever. Along with many others – including Goldman (1991), Sosa (2003), Alston (2005), Grimm (2009), and Raz (2009) –, I find this view very hard to believe. While it can seem at least somewhat plausible to think that we always have reason to prefer having a true belief about a certain matter to having a false belief about it, it does not seem plausible at all to think that we always have reason to prefer having a true belief to having no belief at all, even if the matter in question is the number of sand grains in our hands (Sosa, 2003, 156) or the 323rd entry in the Wichita, Kansas, telephone directory (Goldman, 1991, 88).¹⁹

Moreover, the assumption of a content-independent value of truth gives rise to a further problem. If having true beliefs is intrinsically valuable independently of their content, it seems to follow that we have practical reasons to bring about states in which this value is realized, and one way to do this is to change the world in ways that make our beliefs true (Raz, 2009). In other words, epistemic teleologists are pressed to make axiological assumptions that pose a symmetry between normative pressure to adjust beliefs to the world and normative pressure to adjust the world to our beliefs, while in fact there should be a strong asymmetry. Call this the *direction of fit problem*.

§1.3 *The epistemic assumption*

The axiological assumption makes epistemic reasons conditional on the value of beliefs and thereby threatens to *undergenerate* epistemic reasons or avoid such an undergeneration only at the cost of an implausible axiology. The epistemic assumption, by contrast, runs the risk of *overgenerating* epistemic reasons because it claims that there is an epistemic reason for every belief that instantiates the relevant value. The problem is that beliefs might promote truth, knowledge and the like in cases

¹⁹ In my view, it is no consolation to hear that the reason in question can be outweighed by reasons for preferring more valuable outcomes (cf. Kvanvig 2003, 41; Lynch 2004, 47). Intuition suggests that independently of any antecedent value, there is not the slightest reason for preferring a true belief about the telephone book entry over having no belief at all. Schroeder (2007, 92–97) has famously argued that such intuitions about negative reasons existentials are systematically unreliable in cases of massively outweighed reasons. But this debunking strategy does not help with the cases at hand, in which the reasons in question are not massively outweighed (and it also relies on questionable assumptions, cf. Kiesewetter and Gertken 2021, 275–77).

in which we have no evidence for them, which is to say that the value-based account entails epistemic reasons for beliefs that are not supported by any evidence.

Some proponents are willing to accept this implication. Foley claims that “you can have an adequate epistemic reason to believe a proposition for which you lack sufficient evidence” if your believing this proposition itself “create[s] evidence” for it (1993, 30) – for example, if your belief that you will pass an exam itself makes it sufficiently likely that you will pass it. Indeed, this seems to follow if one thinks of epistemic reasons as reasons for promoting the goal of having true beliefs – but on the face of it, it does not strike me as a very plausible implication. Intuitively, we don’t have epistemic reason for believing a proposition that is not supported by any evidence just because it would be true *if* we believed it. Call this the *problem of self-supporting beliefs*.

Talbot, another epistemic teleologist, goes so far as to claim that a demon’s threat that one will be killed unless one believes *p* provides an *epistemic* reason for believing *p*, independently of one’s evidence regarding *p*, because doing so promotes forming epistemically valuable beliefs in the future. Again, while this follows from Talbot’s assumption that “epistemic reasons have their force because they promote the epistemic end” (2014, 603), it just seems to get the extension of epistemic reasons wrong. As Berker (2013) puts it aptly, it is difficult to see how epistemic teleologists can respect the “separateness of propositions” and avoid counterintuitive trade-offs, according to which we have epistemic reason to believe propositions that are highly unlikely on our evidence if only this promotes our having sufficiently many beliefs in propositions that are evidentially supported. Call this the *problem of epistemic trade-offs*.²⁰ Both the problem of self-supporting beliefs and the problem of epistemic trade-offs illustrate that in virtue of the epistemic assumption, epistemic teleology seems to generate epistemic reasons for beliefs in propositions that aren’t supported by any evidence – epistemic reasons that do not plausibly exist.

§2.3 *Non-promoting teleology*

Some of the problems for the axiological and the epistemic assumptions might be avoided on a view according to which the fundamental value that explains why epistemically supported beliefs are good is a value that does not give rise to reasons for promotion (Littlejohn, 2018; Wedgwood,

²⁰ Foley (1993, 20) and Portmore (2018, 766, n. 2) suggest that this problem can be avoided by a synchronic view, according to which the relevant value is believing the truth *now*. But why should time play such an important role for the relevant value? Moreover, Talbot (2014, 605) shows that a synchronic view allows for trade-offs as well: a person might lose a great number of beliefs with epistemic value instantaneously unless she believes *p* against the evidence, in which case the value of having true beliefs *now* supports believing against the evidence.

2018; Sylvan, 2020). Drawing on the distinction between promoting and honouring a value, epistemic teleologists might claim that epistemically supported beliefs are good not because they promote, but because they honour or respect truth or whatever it is that has fundamental epistemic value.²¹ If one is willing to deny that epistemic value gives rise to *any* reasons for promotion, one is thereby in a position to deny the existence of reasons to adjust the world to one's beliefs – thus solving the direction of fit problem – as well as reasons to believe insufficiently supported propositions in order to promote epistemically valuable beliefs in other propositions – thus solving the problem of epistemic trade-offs. It might also be argued that self-fulfilling prophecies do not show proper respect for epistemic value, which would solve the problem of self-supporting beliefs.

However, simply asserting that epistemic value does not give rise to reasons for promotion, without providing any explanation of why that is so, appears to be an *ad hoc* response to these problems. If true belief or knowledge is intrinsically valuable, why is it that we lack *any* reason to promote this value? The absence of reasons for promotion does not follow from the rejection of epistemic consequentialism or from appreciation of the fact that there are ways of appropriately responding to value that aren't ways of promoting that value.

Moreover, the problem of content-independence still persists. In order to explain the content-independence of epistemic reasons, non-promoting-teleologists still need to claim that true beliefs (or the like) are intrinsically valuable, no matter what they are about.

In response to these problems, non-promoting teleologists might now revert to the idea that epistemic value is purpose- or standpoint-relative value of the sort that is paradigmatically expressed by attributive uses of 'good' (Sosa, 2007; Littlejohn, 2018; Sylvan, 2020).²² If epistemic goodness is only relative to the kind 'belief', or relative to the epistemic standpoint, then it is perhaps not surprising that it is content-independent. Nor is it surprising that we lack reasons to promote this goodness. However, it is then unclear how epistemic value can ground normative reasons. Just as we lack reasons to *promote* attributive value as such, we also seem to lack reasons to *honour* or *respect*

²¹ See esp. Sylvan (2020). Similarly, Wedgwood claims that “relevant are not the values that [...] beliefs [...] *promote*, but the values that [beliefs] instantiate” (2018, 90). However, while focusing on instantiation might deal with the trade-off problem, it still seems subject to the problems of self-supporting beliefs and direction of fit, as in the latter two cases, the value seems instantiated.

²² According to Sosa, “epistemic evaluation is insulated from domain-independent evaluation in the same way attributive evaluations of an archer's shots as good shots are insulated” (Sylvan and Sosa 2018, 556, n. 4; see esp. Sosa 2007, Ch. 4). Littlejohn (2018) and Sylvan (2020) also explicitly embrace an attributive conception of epistemic value. In contrast to these authors, Wedgwood maintains that the values that ground epistemic norms are “non-relative or absolute, rather than relativized either to particular assumed standards or to arbitrary ends” (2017, 201).

attributive value as such. This strikes me as a huge gap in Sylvan’s “epistemic nonconsequentialism”.²³ Moreover, the fact that a response instantiates attributive value does not entail that there is a reason for it. For example, the fact that a certain act is a *good chess move* does not by itself entail a normative reason for this act; at best it entails a reason that is conditional on further assumptions, such as the assumption that you have a reason to play chess. Hence, the best we can expect from the attributive conception of epistemic value is the conditional claim that *if* we have reason to form a belief about a certain matter, *then* we have reason to do this in accordance with the evidence. But this doesn’t give us an account of epistemic reasons for beliefs in matters we have no independent reason to care about. In fact, it gives us no account of epistemic reasons at all. It is an essential truth about epistemic reasons that they do not depend on independent reasons for forming a belief about the matter. It follows that if one provides an account of reasons for belief that depend for their existence on antecedent reasons to form beliefs about the matter, the reasons one has accounted for aren’t epistemic reasons.²⁴

§2.4 *The explanatory assumption*

This leads to an even more fundamental worry with epistemic teleology, which finally concerns its *explanatory assumption*, i.e. the assumption that epistemic reasons are explained in terms of the value of beliefs. The worry is that this assumption entails that epistemic reasons are, in Parfit’s terms, “state-given” rather than “object-given” reasons (Parfit, 2011). State-given reasons for an attitude are reasons provided by properties of the attitude, while object-given reasons are reasons provided by properties of the attitude’s object. Since value-based reasons for belief are provided by the evaluative properties of beliefs, they are state-given. Indeed, it is worth noting that the value-based conception of normative reasons directly entails that all normative reasons for attitudes are state-given reasons, and that the only way to show that epistemic reasons are normative involves showing that they are state-given. The worry is that it is part of the nature of epistemic reasons for beliefs that they are provided by properties of the *object* of belief rather than by the properties of

²³ Similarly, it is unclear to me why it is “normatively good”, as Littlejohn (2018, 42) claims, to conform with the fundamental norm of belief if this norm is grounded in the attributive value of beliefs.

²⁴ This point also shows that pragmatists like Rinard and Sharadin are wrong to think that they can vindicate the claims made in epistemic discourse “for the most part” (Sharadin, 2018, 3807), by showing that “in most ordinary cases, evidence in favor of P constitutes a pragmatic reason to believe it” (Rinard, 2015, 219). The claims made in epistemic discourse are claims about epistemic reasons, which are by their nature content-independent. Hence, one doesn’t vindicate epistemic discourse *even in part* by showing that there are often pragmatic reasons for believing in accordance with the evidence.

the attitude of belief, and that an account of epistemic reasons as state-given reasons would make epistemic reasons, paradoxically, the wrong kind of reasons for belief.²⁵ Call this the *wrong kind of reasons problem* for epistemic teleology.²⁶ Since this worry is independent of the kind of axiology as well as the extensional story that an epistemic teleologist might offer, it is a worry with the very idea of a value-based conception of epistemic reasons. Moreover, since it applies independently of whether reasons are based on the value of the attitude or on some indirect relation that the attitude bears to a value, it affects not only direct but also indirect value-based views.²⁷

It might be replied that a value-based reason to believe in accordance with the evidence is at least partly object-given, because it at least partly depends on properties of the relevant proposition (such as the property that it is supported by the evidence). But this does not show that the reason is object-given in the relevant sense. Suppose that a demon threatens to kill you unless you form a doxastic attitude about whether *p* is true that is supported by your evidence, and suppose that your evidence supports *p*. A reason to believe *p* that depends on this threat is a state-given reason even though it partly depends on properties of the believed proposition. That the reason is explained by the threat is sufficient to show that it is not an object-given reason (even if it is coextensive with such a reason).

To further substantiate the problem at issue, recall the point – well-known from debates about pragmatic reasons for belief – that beliefs don't seem to be sensitive to considerations concerning their own value (see e.g. Parfit, 2011, App. A). Beliefs may be caused by beliefs about their own value in atypical cases, but they are never caused in this way by way of a competent exercise of our capacity of epistemic rationality. Epistemic reasons, in contrast, seem to be essentially the kind of thing a belief in which can cause another belief by way of a competent exercise of our capacity of epistemic rationality. The value-based reasons for belief that the epistemic teleologist appeals to therefore cannot be epistemic reasons.

In response, teleologists might argue that beliefs are sensitive only to considerations about *epistemic* value, such as truth or knowledge. But cases of “truth-promoting non-evidential reasons”

²⁵ With Parfit (2011, App. A) leading the way, many philosophers identify so-called wrong kind of reasons for attitudes with state-given reasons. But some have argued that there can be state-given reasons *against* beliefs that are nevertheless epistemic reasons or at least reasons of the right kind (see Schroeder, 2012; Wedgwood, 2017, 44–45). Whether or not these arguments are successful, they don't seem to call into question that state-given reasons *for* beliefs are reasons of the wrong kind, and thus not epistemic reasons.

²⁶ Not to be confused, of course, with the problem with the same name for the fitting attitude-analysis of value, put forward by Rabinowicz and Rønnow-Rasmussen (2004).

²⁷ Such as, for example, the one suggested in Rowland (2022).

(Talbot, 2014) show that our beliefs' sensitivity to evidence is misconstrued as a sensitivity to epistemic value. A person who learns that she could gain or maintain a great number of beliefs with final epistemic value (i.e. beliefs that are true, amount to knowledge, etc.) if only she believed a proposition that she knows to be ruled out by her evidence, cannot by way of competently exercising her capacity of epistemic rationality adopt this belief. Belief's insensitivity to value is thus not limited to non-epistemic values.

Non-promoting teleologists might reply that this only shows that beliefs couldn't be sensitive to the *instrumental* value they have in virtue of *promoting* final epistemic value, but not that they couldn't be sensitive to the *non-instrumental* value they have in virtue of *respecting* epistemic value. However, it also doesn't seem plausible to think that we could rationally form beliefs on the basis of considerations about whether those beliefs respect epistemic value (rather than considerations bearing on the belief's *content*). Suppose you believe that believing *p* respects the truth, but you do not in fact have any evidence for *p*, nor do you believe that you have such evidence (imagine a case of dogmatic religious belief, for example). In such a case, it does not seem plausible to think that you could base your belief in *p* on your judgement that believing *p* would respect epistemic value by way of exercising your capacity of epistemic rationality. Belief's insensitivity to value thus includes non-instrumental epistemic value as well.²⁸

To sum up, a value-based account of epistemic reasons faces a number of severe objections, both with respect to its axiological as well as with respect to its epistemic assumption. But even if there were an extensionally adequate account available that does not involve implausible axiological and epistemic assumptions, there is reason to doubt that this would be an account of epistemic reasons rather than a coextensive class of state-given reasons for belief.

If epistemic reasons are not value-based, then we have to make a choice between normativism about epistemic reasons and the value-based conception of normative reasons. Anti-normativists often presuppose a teleological framework, holding that epistemic reasons aren't normative since they cannot be value-based. Normativists have to draw the opposite conclusion and reject the value-based conception of normative reasons.²⁹ This, however, poses a challenge: How can normativists

²⁸ To consider an analogous case, it seems highly plausible that admiration can be rationally based only on considerations that bear on the admirability of the object of admiration, and not on considerations about whether the *admiration itself* has non-instrumental value in virtue of respecting the values that make the object admirable.

²⁹ Against the background of the arguments against epistemic teleology provided in this section, all arguments for normativism (e.g. Howard and Leary, 2022; Kiesewetter, 2022b; Schmidt, 2024) also provide arguments against the

account for the reasons/value-link without assuming epistemic teleology and without running into its problems?

§3 ALTERNATIVES

§3.1 *The aretaic value account*

The first response to this challenge that I wish to discuss is based on the idea that the disposition to respond correctly to normative reasons available to one is a *virtue*, by which I mean a disposition that is non-instrumentally valuable. This assumption seems supported by our ordinary practice of praise. We consider a disposition to be a virtue insofar as we think that manifesting it is, other things being equal, worthy of praise, and we do in fact praise people for manifesting their disposition to respond correctly to normative reasons.³⁰ For example, we find it commendable when a belief-forming process manifests a person's disposition to respond to epistemic reasons. More generally, responses that manifest the disposition to respond to normative reasons are often called 'rational', which seems to convey a sort of praise. (The term 'rational' can also be used with non-praising intent, but in such cases, it does not seem to refer to reasons-responsiveness).

If the disposition to respond correctly to normative reasons is a virtue, then conformity with normative reasons is a requirement of virtue: one cannot be (fully) virtuous without conforming to one's normative reasons. Insofar as we have reasons to care about whether we are virtuous, we also have reasons to care about whether we conform to normative reasons. If epistemic reasons are normative, the same holds for epistemic reasons. But on this aretaic value account, epistemic reasons for and against believing *p* are not *explained* by the fact that it is independently good to believe or refrain from believing *p*. Rather, what makes it good, in the relevant aretaic sense, to believe or to refrain from believing *p* is that our epistemic reasons require us to do so. The proposal is thus compatible with the rejection of epistemic teleology and the value-based conception of normative reasons.

A potential problem for the aretaic account is the possibility of non-virtuous conformity with reasons. If reasons-responsiveness is a virtue, then manifesting this virtue by conforming to reasons

value-based conception. For independent arguments against the value-based conception of normative reasons, see also Way (2013) and Kiesewetter (2022a; 2023).

³⁰ Why 'other things being equal'? Because manifestations of virtues can be no-brainers, in which case they are at best minimally praiseworthy (cf. Johnson King, ms).

is good. But the reasons/value-link seems to entail that conformity with reasons is good even if it does not manifest the virtue of reasons-responsiveness.

One reply to this challenge maintains that virtue-analogous responses – responses that merely accord with virtue without manifesting it – have aretaic value. Wedgwood’s account of virtue directly entails this claim:

In general [...], each virtue involves three related kinds of goodness:

- i. The goodness of a *disposition*;
- ii. The goodness of the performances that *manifest* this disposition;
- iii. The ‘*abstract*’ goodness of the performances that this disposition normally produces (a kind of goodness that could in principle be exemplified by performances that result by chance, rather than from the disposition).³¹

To illustrate, suppose that courage is a virtue. It follows that courage is a good disposition, but it also follows that certain actions are good. According to condition (ii), it follows that actions that manifest the disposition of courage, i.e. actions done from courage, are good. And according to condition (iii), it also follows that actions that are not performed out of courage, but which are *in accordance with* courage insofar as a courageous person would have performed them, are good. This latter implication is not intuitively evident, but it may be defended as follows. Aristotle famously distinguishes actions that are “done justly or moderately” (actions that manifest the virtue of justice or temperance) from “just and moderate actions” (actions that accord with the virtue of justice or temperance even if they do not manifest it) (E.N. II 4). Following him, we can say that the act that the courageous person would have performed is a courageous act no matter whether it is done courageously. If being courageous is a way of being good, such an act is good (in some way, to some extent). For example, suppose that some person is attacked and you come to her defence, but your action is the outcome of an accidental mood rather than a stable disposition to be courageous. We can still say that coming to the victim’s defence is good to the extent that it is a courageous act, i.e. an act that is in accordance with the virtue of courage.

One might object that the act of coming to a person’s defence is good not because it is in accordance with virtue, but because of other, non-aretaic values, such as the well-being of the victim.

³¹ Wedgwood (2017, 141).

But while coming to someone's defence may well *also* be good because of such non-aretaic values, it need not be so. Courage might require you to offer support even if you unknowingly witness a street art performance and no one is in real danger. In such cases, acts can be good for the mere reason that they are courageous.

Wedgwood's claim that virtue-analogous responses are valuable as such is contentious, however, and I don't think it is necessary in order to account for the reasons/value link in terms of virtue. Recall that the reasons/value-link was premised on the idea that we have reasons to care about whether we conform to normative reasons. It's not clear that this should be taken to entail that every token of conformity (whether virtuous or not) is good. Perhaps only the virtuous tokens of conformity are good (in the aretaic sense), and we have reason to care about whether we conform to reasons because such conformity is a prerequisite for instantiating this value. On this view, we should allow for reasons for pro-attitudes towards things that – although not in themselves good – are preconditions of goods, and consequently the reasons/value-link should be understood as allowing that conformity with normative reasons might merely be such a pre-condition. The essence of what needs to be accounted for – the fact that we have reasons to care about conformity with normative reasons – would still be preserved.

§3.2 Avoiding the problems of epistemic teleology

Let's consider how the aretaic value account avoids the problems for teleological views that I discussed above. To begin with, the aretaic value account isn't vulnerable to the epistemic worries concerning trade-offs and self-supporting beliefs. Since it doesn't say that value properties of beliefs ground epistemic reasons, there is no reason to suspect that considerations about the promotion of epistemic value could give rise to epistemic reasons for evidentially unsupported beliefs. For the same reason, the proposal isn't subject to a wrong kind of reasons problem. However, it involves the axiological assumption that belief (or abstention from belief) is good (or at least a precondition of a good) to the extent that it is required by conformity with epistemic reasons, and this assumption might give rise to worries analogous to those concerning the axiological assumption of epistemic teleology.

First and foremost, a pressing question is how the proposal deals with the problem of content-independence. Above, I argued that both the instrumental as well as the intrinsic value of having knowledge, true or likely beliefs is content-dependent, and that this value therefore cannot explain why we have epistemic reasons to believe (probable) truths or to refrain from believing (probable)

falsities about trivial or uninteresting matters. It is important to see that the present proposal is consistent with this plausible verdict, as it does not involve the claim that beliefs that are true or likely have instrumental or intrinsic value. Instead, it claims that epistemically supported beliefs have non-instrumental derivative value by way of manifesting or being in accord with the virtue of reasons-responsiveness, and that this value is content-independent does not strike me as objectionable.

For example, despite the fact that it doesn't matter whether Stella Starlet has blemished skin, and that there is no value in gaining knowledge about it, once you have the relevant evidence, it follows that you can manifest the virtue of epistemic rationality by forming the belief that Stella Starlet has blemished skin. Depending on our view about positive epistemic requirements, it need not be rationally required and need not be better, in terms of the relevant aretaic value, than refraining from belief. But at least in case you can conform to your epistemic reasons only by believing what you have justification to believe (for example because you attend to the relevant question), any reason to prefer to be virtuous entails a reason to prefer that you form that belief – and this reason will be independent of the content of the belief.

Consider an analogy. If answering a question sincerely is required by the virtue of honesty in a particular situation, then answering sincerely is good (or a precondition for a good) even if no non-aretaic value is promoted or respected by doing so. If honesty is a virtue, we have reason to care whether we answer sincerely in such a situation and have at least a *pro tanto* reason to prefer that we act in this way. Analogously, if believing what we have justification to believe is required by the virtue of epistemic rationality, then believing so is good (or a precondition for a good) even if no non-aretaic value is promoted or respected by doing so. If epistemic rationality is a virtue, we have reason to care whether we believe what we have justification to believe, and if we cannot be rational without believing so, we have at least a *pro tanto* reason to prefer that we do so. Once one allows that the value of justification derives from the aretaic value of reasons-responsiveness, it is no longer implausible that this value is content-independent.

However, the focus on aretaic value might also give rise to a worry about the proposal, namely that aretaic value is merely attributive or kind-relative, and that merely attributive value doesn't entail reasons to care. In response to this worry, I wish to say that it is independently plausible that we have reason to care about aretaic value, if only because we have reason to care about whether our behaviour satisfies conditions for praiseworthiness. This need not mean that praiseworthiness figures as an end in the foreground of practical deliberation – one can care about values in ways

other than treating them as aims to be promoted. But it is plausible to think that we have reason not to be indifferent towards the question of whether our behaviour is or fails conditions of praiseworthiness, and this entails that we have reason to care about aretaic value. So, either aretaic value involves (perhaps in addition to attributive value) predicative, personal or person-relative value, or it is a kind of attributive value that we have reason to care about as such. Either way, we can maintain that the aretaic defence of the reasons/value link delivers reasons to care about conformity with normative reasons without assuming that attributive value in general involves such reasons.

§3.3 Aretaic value, direction of fit and the promotion/respecting-distinction

Finally, let's consider the *direction of fit problem* for epistemic teleology. This is the problem that the assumption that true belief, knowledge, or the like are desirable as such seems to have the absurd implication that we have practical reasons to change the world in such a way that our beliefs come out true. It might be worried that my proposal has a similar implication. For suppose that you have an unjustified belief in *p* and that you could either give up this belief or alternatively create justification for it (e.g., by making *p* true or by creating new evidence for *p*). If justification has value independently of the benefits of getting things right about a particular matter, doesn't it follow that I have instrumental reasons to create justification for my unjustified beliefs? After all, this would seem to be a sufficient and often permissible means to achieve what is valuable.

As indicated above already, the conclusion follows only if the value in question is linked with reasons for *promotion* rather than merely with reasons for *honouring* or *respecting*. However, virtues are the prime examples of values that call for respect rather than promotion (see e.g. McNaughton and Rawling, 1992). More exactly, while virtues or good dispositions may be reasonably respected as well as promoted, it is doubtful that we have any reason to promote the goodness involved in responses that manifest or accord with virtue.

To see this point, consider virtues like courage, reliability, or honesty. While it's not implausible to think that we have reason to bring about states of affairs in which such dispositions are more pervasive, thereby promoting the value of these dispositions, it seems to be a mistake to think that we have reason to bring about states of affairs in which such dispositions are more often manifested or in which agents more often act in accordance with them. We might have such reasons in contexts of training someone to be virtuous, in which case we promote the responses in order to promote the dispositions. But we have no reasons to promote the responses as such, because of

their own aretaic value. For example, unless for the purpose of training, we have no reason to create dangerous situations so that courageous people more often manifest their courage, we have no reason to go around making promises in order to promote reliable action, and we have no reason to ask people questions in order to create opportunities for honest answers. It would be a mistake to conclude that virtuous action is merely instrumentally valuable (valuable only if and because it promotes virtuous dispositions), for we seem to value virtuous action independently of its effects on dispositions. Rather, what this shows is that the non-instrumental value of virtuous behaviour is not one that we have reason to promote.

Why is that? I said above that non-promoting teleologists owe us an explanation of why the value they appeal to does not give rise to reasons for promotion. So what is my explanation that the aretaic value I appeal to doesn't give rise to reasons for promotion? In my view, a plausible explanation is that this value is, despite being non-instrumental, nonetheless derivative: virtuous (and perhaps also virtue-analogous) responses are good because they manifest (or accord with) virtue; their value derives from the value of the disposition.³² This means that we can explain why we have no reason to promote virtuous responses by appealing to the more general fact that we never have reason to promote derivative value as such. When we have reason to promote derivative value, this is because doing so promotes non-derivative value, in which case the derivative value is instrumentally derivative. But if responses that manifest or accord with virtue are derivatively valuable, then their derivative value must be non-instrumental, and so there are no reasons to promote these responses as such.

To sum up my response to the direction of fit problem: creating justification for unjustified beliefs might *promote* the value of justification, but since this value is non-instrumentally derivative, there are no reasons to promote this value as such. There are reasons for caring about it, respecting, or honouring it, but creating justification for unjustified beliefs is not a way of caring about, respecting, or honouring the virtue of responsiveness to normative reasons for belief. For these reasons, the present explanation of the reason/value-link is not vulnerable to the problem. In my view, this is a remarkable advantage that this explanation has over other, non-aretaic ones.

³² Hurka (2001, Ch. 1) also claims that the value of virtuous responses is non-instrumentally derivative, although in his view this value derives from the value of its objects. This view is consistent with my approach only on the further assumption (discussed in the next subsection) that conformity with normative reasons is valuable as such.

§3.4 *The normative value account*

Finally, I would like to consider an alternative to the aretaic account that provides an even simpler explanation of the goodness of conforming with epistemic reasons. Consider a trivial promise, which requires you to perform an act that has no independent value. Keeping this promise might still plausibly be said to have *deontic value*, to be good in virtue of discharging an obligation that one has incurred by giving the promise (Kiesewetter, 2022a). Expanding on this idea, one could argue that conforming with normative reasons is generally valuable as such.³³ It would follow from this that conforming with epistemic reasons is valuable simply in virtue of the fact that epistemic reasons are normative.

What could be said in favour of the view that conformity with reasons is valuable as such? The argument for the reasons/value-link, together with the failure of the value-based conception of normative reasons, provides an abductive argument for this assumption. “My reasons most strongly support ϕ -ing, but I don’t have any reason to care whether I ϕ ” sounds paradoxical; coherent agents do in fact value conformity with normative reasons. The value-based conception of normative reasons could explain this, but as we have seen, it is ruled out by normativism. The assumption that conformity with normative reasons is finally valuable provides an attractive alternative explanation.

As in the case of aretaic value, we have independent reasons to think that normative value doesn’t call for promotion. If it did, we would have reasons to promise actions that we will perform anyway in order to promote the value of promise-keeping – a *reductio*.³⁴ In fact, the very idea that there are reasons to promote normative value is hard to make sense of. The normative value of a response already presupposes a reason for that response. If we had reasons to promote this value, these would have to be added to the overall balance of reasons, thereby strengthening the case in favour of the response. This not only seems like an implausible form of double-counting, it also renders the idea that there is normative value in conforming with the overall balance of reasons incoherent, as this idea requires that the overall balance can be determined independently of the

³³ This view might be spelled out in different ways: ϕ -ing might be good because it is supported by pro tanto, sufficient, or decisive normative reasons; or it might only be comparatively good because it is normatively bad not to conform to (pro tanto, sufficient, or decisive) normative reasons. The argument below supports no more than the assumption that ϕ -ing is (either comparatively or non-comparatively) good if and because ϕ -ing is supported by decisive reasons.

³⁴ Compare Smith (1997) and Kiesewetter (2022a, 39–40). As Smith notes, problems with promotion are also avoided if one holds that conformity is only *comparatively* good because of the badness of non-conformity.

normative value of conforming with it. Normative value thus cannot be understood as calling for promotion.³⁵

As normative value doesn't entail reasons for promotion, the normative value account doesn't face the problems of direction of fit, self-supporting beliefs and epistemic trade-offs. As it doesn't *base* epistemic reasons on value, it doesn't render epistemic reasons reasons of the wrong kind. It also avoids ascribing content-independent value to true or likely belief as such. However, as with the aretaic value account, it entails that there is content-independent value in conforming with epistemic reasons. Again, this doesn't strike me as objectionable. What is objectionable is the assumption that beliefs that are true, likely true, or amount to knowledge are *instrumentally* or *intrinsically* valuable independently of their content. Denying this is what motivates the nihilist assumption, and also what this assumption gets right. But if an independently supported evaluative assumption entails that there is *non-instrumental derivative* value in having such beliefs, we should not insist on the nihilist assumption. In fact, then, both the aretaic and the normative value account offer a solution to the puzzle I described at the outset that preserves not only normativism and the reasons/value-link, but also the core intuition behind the nihilist assumption.

§4. CONCLUSION

In light of the content-independence of epistemic reasons, normativists about epistemic reasons seem to face a dilemma. Either they accept the perplexing conclusion that there are normative reasons we have no reason to care about, or they have to defend the assumption that conformity with epistemic reasons is good even in cases in which the supported beliefs seem entirely uninteresting or trivial – cases like *Magazine*. In this essay, I have argued that they should take the second horn of this dilemma. However, I've also argued that the most common way of taking this horn is doomed to fail. Epistemic teleologists claim that epistemically supported beliefs are independently of their content (either instrumentally or intrinsically) valuable, and that this value explains why we have normative reasons to have such beliefs. But this assumption seems false, and the view poses a number of further problems that cannot satisfactorily be solved. Finally, I have outlined two non-teleological strategies for explaining the goodness of conformity with epistemic reasons. On the aretaic value account, conforming with epistemic reasons is good because it

³⁵ Note that this does not mean that we do not have reasons to promote conformity with final reasons – of course, we do. But these can be explained by instrumental transmission principles (cf. Kiesewetter, 2015; 2018; Kiesewetter and Gertken, 2021) rather than by normative value.

manifests or accords with the virtue of being responsive to normative reasons (or is at least a precondition of the goodness involved in having this virtue). On the normative value account, it is good simply because conformity with normative reasons is good as such. On both of these views, the value of epistemically supported beliefs is non-instrumentally derivative. This preserves the core intuition behind what I've called the nihilist assumption, namely that both the instrumental as well as the intrinsic value of epistemically supported beliefs depends on the subject matter. It thus solves the problem of content-independence, and I've argued that it solves the other problems of epistemic teleology as well.

Is there a reason to prefer one account over the other? The aretaic and the normative value account are not in conflict and could be argued to reinforce each other. On one possible view, normative value is more fundamental than aretaic value and explains why reasons-responsiveness is a virtue. This would allow us to maintain the traditional picture of virtues as dispositions "towards something good" (Aristotle, E.N. I 12). The aretaic goodness of conformity with reasons would be just a shadow of its normative goodness.³⁶ On another possible view, aretaic value is more fundamental, and normative value can be explained in terms of it. This would require us to understand virtues more broadly as dispositions to respond to normative reasons rather than values, an understanding that would then give us an explanation of why conformity with normative reasons is valuable. Both of these views carry substantive commitments and both have attractions. I won't try to adjudicate between them here. As I see it, they both offer tenable responses to the dilemma I have been concerned with and should be taken seriously as ways of resolving it.³⁷

³⁶ Note that on this picture there would also be no need to assume that non-virtuous conformity with virtue has aretaic value.

³⁷ Earlier versions of this article have been presented at the *Unity of Normativity* Conference at University of Vienna (2020); the LOGOS Epistemology Workshop at University of Barcelona (2020); the Dresden-Hamburg Workshop at TU Dresden and University of Hamburg (2021); as well as colloquium talks at University of St. Andrews (2021); the Human Abilities Centre at HU and FU Berlin (2022); Bielefeld University (2024); and the University of Hawai'i at Mānoa (2024). I would like thank the participants for their valuable feedback. For written comments, I am grateful to Dmitry Ananiev, Singa Behrens, Selim Berker, Jaakko Hirvelä, Felix Koch, Toni Rønnow-Rasmussen, Sebastian Schmidt, Shane Ward, and Jonathan Way. Work on this article was funded by the Deutsche Forschungsgemeinschaft (Centre for Advanced Studies "Human Abilities", HU and FU Berlin) and the European Union (ERC Grant 101040439, REASONS FIRST). Views and opinions expressed are however those of the author only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

References

- Alston, W. (1988). The Deontological Conception of Epistemic Justification. *Philosophical Perspectives*, 2, 257–99. doi: 10.2307/2214077
- . (2005). *Beyond “Justification”: Dimensions of Epistemic Evaluation*. Ithaca: Cornell University Press.
- Aristotle. *Nicomachean Ethics*. Edited by Christopher Rowe and Sarah Broadie. Oxford: Oxford University Press (2002).
- Berker, S. (2013). Epistemic Teleology and the Separateness of Propositions. *The Philosophical Review*, 122, 337–93. doi: 10.1215/00318108-2087645
- Côté-Bouchard, C. (2017). Is Epistemic Normativity Value-Based? *Dialogue*, 56, 407–30. doi: 10.1017/S0012217317000506
- Feldman, R. (2000). The Ethics of Belief. Reprinted in E. Conee & R. Feldman (Eds.), *Evidentialism: Essays in Epistemology* (pp. 166–95). Oxford: Oxford University Press (2004).
- Foley, R. (1987). *The Theory of Epistemic Rationality*. Cambridge, MA: Harvard University Press.
- . (1993). *Working without a Net: A Study of Egocentric Epistemology*. New York: Oxford University Press.
- Gertken, J. & Kiesewetter, B. (2017). The Right and the Wrong Kind of Reasons. *Philosophy Compass*, 12 (5), 1–14. doi: 10.1111/phc3.12412
- Glüer, K. & Wikforss, Å. (2018). Reasons for Belief and Normativity. In D. Star (Ed.), *The Oxford Handbook of Reasons and Normativity* (pp. 575–99). Oxford: Oxford University Press. doi: 10.1093/oxfordhb/9780199657889.013.26
- Goldman, A. I. (1991). *Knowledge in a Social World*. Oxford: Clarendon Press.
- . (2001). The Unity of the Epistemic Virtues. Reprinted in his *Pathways to Knowledge* (pp. 51–70). Oxford: Oxford University Press (2002).
- Gregory, A. (2016). Normative Reasons as Good Bases. *Philosophical Studies*, 173, 2291–2310. doi: 10.1007/s11098-015-0609-8
- Grimm, S. R. (2009). Epistemic Normativity. In A. Haddock et al. (Eds.), *Epistemic Value* (pp. 243–64). New York: Oxford University Press.
- Howard, C. & Leary, S. (2022). In Defence of the Right Kind of Reason. In C. Howard & R.A. Rowland (Eds.), *Fittingness: Essays in the Philosophy of Normativity* (pp. 221–41). Oxford: Oxford University Press. doi: 10.1093/oso/9780192895882.003.0010
- Hurka, T. (2001). *Virtue, Vice, and Value*. Oxford: Oxford University Press.
- Johnson King, Z. (ms). The Slow Clap Phenomenon. Unpublished manuscript. Harvard University (2023).
- Kelly, T. (2007a). Epistemic Rationality as Instrumental Rationality: A Critique. *Philosophy and Phenomenological Research*, 66, 612–40. doi: 10.1111/j.1933-1592.2003.tb00281.x
- . (2007b). Evidence and Normativity: Reply to Leite. *Philosophy and Phenomenological Research*, 75, 465–74. doi: 10.1111/j.1933-1592.2007.00085.x
- Kiesewetter, B. (2015). Instrumental Normativity: In Defense of the Transmission Principle. *Ethics* 125, 921–46. doi: 10.1086/680911
- . (2017). *The Normativity of Rationality*. Oxford: Oxford University Press. doi: 10.1093/oso/9780198754282.001.0001
- . (2018). Contrary-to-Duty Scenarios, Deontic Dilemmas, and Transmission Principles. *Ethics* 129, 98–115. doi: 10.1086/698734
- . (2019). Can the Lottery Paradox Be Solved by Identifying Epistemic Justification with Epistemic Permissibility? *Episteme* 16, 241–61. doi: 10.1017/epi.2018.4

- . (2022a). Are All Practical Reasons Based on Value? *Oxford Studies in Metaethics*, 17, 27–53.
- . (2022b). Are Epistemic Reasons Normative? *Noûs*, 56, 670–95. doi: 10.1111/nous.12379
- . (2023). Egalitarian Justice as a Challenge for the Value-Based Theory of Practical Reasons. In A. G. Garcia et al. (Eds.), *Value, Morality & Social Reality: Essays Dedicated to Dan Egonsson, Björn Petersson & Toni Rønnow-Rasmussen* (pp. 239–49). Lund: Department of Philosophy, Lund University.
- Kiesewetter, B & Gertken, J. (2021). How Do Reasons Transmit to Non-Necessary Means? *Australasian Journal of Philosophy*, 99, 271–85. doi: 10.1080/00048402.2020.1745252
- Kornblith, H. (2002). *Knowledge and its Place in Nature*. Oxford: Oxford University Press.
- Kvanvig, J. L. (2003). *The Value of Knowledge and the Pursuit of Understanding*. Cambridge: Cambridge University Press.
- 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, 456–64. doi: 10.1111/j.1933-1592.2007.00084.x
- Littlejohn, C. (2018). The Right in the Good: A Defense of Teleological Non-Consequentialism. In H. K. Ahlstrom-Vij & J. Dunn (Eds.), *Epistemic Consequentialism* (pp. 23–47). Oxford: Oxford University Press. doi: 10.1093/oso/9780198779681.003.0002
- Lynch, M. P. (2004). *True to Life: Why Truth Matters*. Cambridge, Mass: MIT Press.
- Maguire, B. (2016). The Value-Based Theory of Reasons. *Ergo*, 3, 233–62.
- Maguire, B. & Woods, J. (2020). The Game of Belief. *The Philosophical Review*, 129, 211–49. doi: 10.1215/00318108-8012843
- Mantel, S. (2019). Do Epistemic Reasons Bear on the Ought Simpliciter? *Philosophical Issues*, 29, 214–27. doi: 10.1111/phis.12149
- McHugh, C. & Way, J. (2022). *Getting Things Right: Fittingness, Reasons, and Value*. Oxford: Oxford University Press.
- McNaughton, D. & Rawling, P. (1992). Honoring and Promoting Values. *Ethics*, 102, 835–43. doi: 10.1086/293451
- Nelson, M. T. (2010). We Have No Positive Epistemic Duties. *Mind*, 119, 83–102.
- Nozick, R. (1993). *The Nature of Rationality*. Princeton: Princeton University Press (repr. 1995).
- Parfit, D. (2011). *On What Matters*. Vol. 1. Oxford: Oxford University Press.
- Pettit, P. (1989). Consequentialism and Respect for Persons. *Ethics*, 100, 116–26. doi: 10.1086/293149
- Portmore, D. W. (2011). *Commonsense Consequentialism: Wherein Morality Meets Rationality*. New York: Oxford University Press.
- . (2018). Teleological Reasons. In D. Star (Ed.), *The Oxford Handbook of Reasons and Normativity* (pp. 764–83). Oxford: Oxford University Press. doi: 10.1093/oxfordhb/9780199657889.013.33
- Rabinowicz, W. & Rønnow-Rasmussen, T. (2004). The Strike of the Demon: On Fitting Pro-Attitudes and Value. *Ethics*, 114, 391–423.
- Raz, J. (2009). Reasons: Practical and Adaptive. Reprinted in his *From Normativity to Responsibility*, 36–58. Oxford: Oxford University Press (2011).
- Rinard, S. (2015). Against the New Evidentialists. *Philosophical Issues*, 25, 208–23. doi: 10.1111/phis.12061
- Rowland, R. A. (2022). Value-First Accounts of Normativity. In C. Howard & R.A. Rowland (Eds.): *Fittingness: Essays in the Philosophy of Normativity* (pp. 176–99). Oxford: Oxford University Press.

- Scanlon, T.M. (1998). *What We Owe to Each Other*. Cambridge, MA: Belknap Press.
- . (2011). The Unity of the Normative. *Philosophical Studies*, 154, 443–50. doi: 10.1007/s11098-011-9736-z
- Schmidt, S. (2024). Epistemic Blame and the Normativity of Evidence. *Erkenntnis*, 89, 1–24. doi: 10.1007/s10670-021-00430-9
- Schroeder, M. (2007). *Slaves of the Passions*. Oxford: Oxford University Press.
- . (2012). The Ubiquity of State-Given Reasons. *Ethics*, 122, 457–88.
- . (2021). *Reasons First*. Oxford: 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
- Skorupski, J. (2010). *The Domain of Reasons*. Oxford: Oxford University Press.
- Smith, H. M. (1997). A Paradox of Promising. *The Philosophical Review*, 106, 153–96. doi: 10.2307/2998356
- Sosa, E. (2003). The Place of Truth in Epistemology. In M. DePaul & L. Zagzebski (Eds.), *Intellectual Virtue: Perspectives from Ethics and Epistemology* (pp. 155–79). Oxford: Clarendon Press.
- . (2007). *A Virtue Epistemology: Apt Belief and Reflective Knowledge*. 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* (pp. 13–33). Cambridge: Cambridge University Press.
- Steglich-Petersen, A. & Skipper, M. (2020). Instrumental Reasons for Belief: Elliptical Talk and Elusive Properties. In S. Schmidt & G. Ernst (Eds.), *The Ethics of Belief and Beyond: Understanding Mental Normativity* (pp. 109–25). New York: Routledge. doi: 10.4324/9780429284977
- Sylvan, K. (2020). An Epistemic Nonconsequentialism. *The Philosophical Review*, 129, 1–51. doi: 10.1215/00318108-7890455
- Sylvan, K. & Sosa, E. (2018). The Place of Reasons in Epistemology. In D. Star (Ed.), *The Oxford Handbook of Reasons and Normativity* (pp. 555–74). Oxford: Oxford University Press. doi: 10.1093/oxfordhb/9780199657889.013.25
- Talbot, B. (2014). Truth Promoting Non-Evidential Reasons for Belief. *Philosophical Studies*, 168, 599–618. doi: 10.1007/s11098-013-0139-1
- Way, J. (2013). Value and Reasons to Favour. *Oxford Studies in Metaethics*, 8, 27–49.
- Wedgwood, R. (2017). *The Value of Rationality*. Oxford: Oxford University Press.
- . (2018). Epistemic Teleology. In H. K. Ahlstrom-Vij & J. Dunn (Eds.), *Epistemic Consequentialism* (pp. 85–112). Oxford: Oxford University Press. doi: 10.1093/oso/9780198779681.003.0005