The true, the good and the justified

Essays on epistemic normativity and value

by

Félix-Antoine Gélineau

MA Philosophy, Université de Montréal, 2019
BA Philosophy, Université de Montréal, 2017

Submitted to the Department of Linguistics and Philosophy in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

at the

Massachusetts Institute of Technology

September 2025

Authored by:	Félix-Antoine Gélineau
	Department of Philosophy
	August 13, 2025
Certified by:	Roger White
	Professor of Philosophy
	Thesis Supervisor
Accepted by:	Brad Skow
	Laurance S. Rockefeller Professor of Philosophy
	Chair of the Committee on Graduate Students

The true, the good and the justified
Essays on epistemic normativity and value

by

Félix-Antoine Gélineau

Submitted to the Department of Linguistics and Philosophy in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

Abstract

This dissertation about the role that value—in particular, the value of truth—plays in the explanation of epistemic norms, the norms that should govern our belief-formation and belief-revision processes. It is a pervasive practice for us to assess one another's beliefs, decrying some and commending others: we find something amiss in a belief formed by wishful thinking, and deem proper a belief arrived at *via* meticulous consideration of one's evidence. It is natural to think that there is a close connection between what epistemic norms sanction, truth, and the fact that truth matters. In what sense is truth valuable? and what does that entail for how we should conceive of epistemic norms? These questions drive my dissertation.

In chapter 1, "The true, the good and the justified", I argue that the teleological conception of epistemic justification, the view that for a belief to be justified is for it to be formed in a way that is conducive to truth, is safe from the main objections it faces. These import assumptions about value that belong to ethical consequentialism; the epistemic teleologist need not and should not accept them. "Epistemic value" is best understood as a 'placeholder', not as a term denoting value in any substantive sense. The upshot is that endorsing a teleological explanation of epistemic norms does not commit one to the idea that true beliefs ought to be promoted in the way the good ought to be promoted according to ethical consequentialists.

In chapter 2, "Why be antisocial (about epistemic normativity)", I examine whether epistemic normativity is grounded in the usefulness of truth for communities and argue that it is not. If what we ought to believe were to be explained in terms of the usefulness of truth for communities, we should expect our epistemic norms to be wildly different from what they are: they could condone trade-offs between true and false beliefs across a community when doing so would be useful for it; they would often fail to prescribe believing over other doxastic attitudes when the epistemic aims of the community would be equally well-served by either; finally, there is no way to answer in a satisfying manner the question of what isolated individuals should believe.

In chapter 3, “What’s truth got to do with it?”, I develop an account of epistemic normativity that does justice to the idea that truth matters but avoids the shortcomings of other value-based accounts. I argue that, given any plausible account of the value of truth, the value of truth can explain some, but not all epistemic norms. The account of epistemic normativity that I present is pluralistic: I distinguish between substantive norms, which are explained by the value of truth, and basal norms, which stem from the good functioning of mechanisms of belief-formation and belief-revision, independently of the value of truth. This account is shown to be superior to other accounts discussed in the course of the dissertation.

Thesis supervisor: Roger White

Title: Professor of Philosophy

Acknowledgments

If this thesis has finally come to fruition, it is only because of the many people on whose guidance, generosity, and support I have been fortunate enough to rely throughout the years. I can only thank so many here, but I am deeply grateful to all those who have played a role in this journey. First and foremost, I offer my heartfelt thanks to my thesis supervisor, Roger White, and to my committee members, Kevin Dorst and Kieran Setiya, for their careful attention, patience, insight, and encouragement. I also want to thank Brad Skow, who, as my advisor for my first two years at MIT, supported me during an especially challenging time.

For daily camaraderie and philosophical companionship, I owe special thanks to my officemates and friends, Abe Mathew and Yonathan Fiat. My thanks also go to my cohort: Jessica Heine, Josh Pearson, Eliza Wells, and John Hill. Since my very first day in Cambridge, Alice Gelman and Andrew Caide have been there for me as the best roommates one could dream of and as exceptional friends. Life in Cambridge would have been unimaginable without you. Jessica Heine was a pillar for me throughout these years, always present, right up to the very last minute of the party. I would never have thought that signing up to be a first-year mentor and meeting up one arvo would lead to such a deep and lasting bond: Sakinah Munday, your friendship is an inseparable part of my life and a constant source of joy. My thanks also to Aaron Chevalier, Skyler Carlin and Gabriel Tobia.

À mes amis du Québec: Jordane Labrie, Félix Durand, Erika Olivaux-Marmignon, Rémi Charbonneau, Rémi Tison, Fannie Achard, Marc-Kevin Daoust, Jade Bergeron, Louis Félix Nadeau et Justine Falardeau : peu ont la chance d'avoir des amis et amies sur lesquels ils puissent compter en tout temps bien que plusieurs centaines de kilomètres et une frontière les séparent. Je me considère extrêmement privilégié d'être dans vos vies. Je tiens également à remercier trois membres du corps professoral de l'Université de Montréal dont les encouragements et le soutien indéfectible durant mon baccalauréat et ma maîtrise furent essentiels à mon succès et sans qui je n'aurais jamais pu réaliser le rêve de compléter un doctorat dans une institution telle que le MIT : Mme Aude Bandini, M. Daniel Laurier et M. Louis-André Dorion.

Il est impossible de sous-estimer l'impact qu'eurent et que continuent d'avoir l'amour et le soutien constant de mon frère, Laurent, ainsi que de mes parents, Andrée et Michel. Je vous aime de tout mon cœur.

Above all, my thanks go to my top person, Liz. “Words cannot do justice” may be a trite formulation, easy to reach for and often hollow, but here it is the very opposite: perfectly fitting and deeply heartfelt. What they cannot do justice to is how much your unshakable trust and confidence in my success, your restorative presence and your ability to keep me afloat, your profound patience, and the energizing power of your laughter have meant to me. You are a continual inspiration to me. I could not be luckier to start this new chapter at your side. I love you.

This thesis benefitted from the financial support of the Social Sciences and Humanities Research Council of Canada in the form of a Doctoral Fellowship (BESC-D).

Contents

Introduction

This dissertation is about the role that value—in particular, the value of truth—plays in the explanation of epistemic norms, the norms that should govern our belief-formation and belief-revision processes. It is a pervasive practice for us to assess one another's beliefs, decrying some and commending others: we find something amiss in a belief formed by wishful thinking, and deem proper a belief arrived at *via* meticulous consideration of one's evidence. It is natural to think that there is a close connection between what epistemic norms sanction, truth, and the fact that truth matters. In what sense is truth valuable? and what does that entail for how we should conceive of epistemic norms? These questions drive my dissertation.

In chapter 1, "The true, the good and the justified", I argue that the teleological conception of epistemic justification, the view that for a belief to be justified is for it to be formed in a way that is conducive to truth, is safe from the main objections it faces. These import assumptions about value that belong to ethical consequentialism; the epistemic teleologist need not and should not accept them. "Epistemic value" is best understood as a 'placeholder', not as a term denoting value in any substantive sense. The upshot is that endorsing a teleological explanation of epistemic norms does not commit one to the idea that true beliefs ought to be promoted in the way the good ought to be promoted according to ethical consequentialists.

In chapter 2, "Why be antisocial (about epistemic normativity)", I examine whether epistemic normativity is grounded in the usefulness of truth for communities and argue that it is not. If what we ought to believe were to be explained in terms of the usefulness of truth for communities, we should expect our epistemic norms to be wildly different from what they are: they could condone trade-offs between true and false beliefs across a community when doing so would be useful for it; they would often fail to prescribe believing over other doxastic attitudes when the epistemic aims of the community would be equally well-served by either; finally, there is no way to answer in a satisfying manner the question of what isolated individuals should believe.

In chapter 3, "What's truth got to do with it?", I develop an account of epistemic normativity that does justice to the idea that truth matters but avoids the shortcomings of other value-based accounts. I argue that, given any plausible account of the value of truth, the value of truth can explain some, but not all epistemic norms. The account of epistemic normativity that I present is pluralistic: I distinguish between substantive norms, which are explained by the value of truth, and basal norms, which stem from the good functioning of mechanisms of belief-formation and belief-revision, independently of the value of truth. This account is shown to be superior to other accounts discussed in the course of the dissertation.

The true, the good and the justified

On epistemic consequentialism and trade-offs

It is a pervasive practice for us to assess one another's beliefs, decrying some as unjustified, while commending others as justified. We find something amiss in a belief formed by wishful thinking, and deem proper a belief at which one arrived as a result of the meticulous consideration of one's evidence. Although beliefs that we call rational or justified in this way may not always be true, it is natural to think that there is a close connection between, on the one hand, justification, and, on the other, truth. At first pass, we might try to cash out this connection in the following way: *justification is conducive to, or is a guide to, truth.* Under normal circumstances, wherever we find justified beliefs, we should expect to find true beliefs, and the reason for this is that beliefs that are justified are beliefs formed in such a way that they are sufficiently likely to be true. A belief formed by wishful thinking deserves negative appraisal because it is unlikely that beliefs formed in this way will turn out to be true, whereas weighing carefully one's evidence is one's best guide to the truth, and thus, the resulting belief deserves a positive appraisal.

This account of the connection between justification and truth – call it the *teleological conception of epistemic justification* – has been a mainstay of contemporary epistemology, having been embraced by epistemologists who are otherwise of very different persuasions. In fact, it would be apt to call it a *common ground* on which debates about epistemic normativity and the nature of justification have been conducted. Yet, the teleological conception has faced vigorous challenges in recent years, with critics such as Selim Berker and Jane Friedman calling for its wholesale repudiation.

Why reject the teleological conception of justification? The main reason advanced by its opponents is that, just as consequentialist theories in ethics fail to respect the 'separateness of persons', so do teleological views in epistemology fail to respect 'the separateness of propositions'. Crudely, according to ethical consequentialists, one ought to do what brings about the best state of affairs, even when doing so would mean sacrificing the welfare of the few for the welfare of the many. Similarly, say the critics of teleological epistemology, the latter view entails that a belief is justified when it best promotes the end of having true beliefs and avoiding false ones, even if it comes at the cost of forming a belief in a patently irrational manner – as long as doing so is what would be the most conducive to the acquisition of true beliefs, such a belief would be justified. Such verdict, critics claim, is 'beyond the pale,' and thus, the teleological conception of epistemic normativity is false (Berker 2013a).

A further line of criticism, taken by Jane Friedman, is that a closer look at our ordinary epistemic norms and at what they prescribe, permit or forbid – what counts as justified – reveals that these norms are not teleological; if they were, they would look very different from the way they actually are. In particular, our ordinary norms don't require the maximization of true beliefs in the way that we should expect them to, were they teleological.

I will argue that these objections fail. They rest on taking the analogy with ethical consequentialism as a guide to the essence of the teleological conception of epistemic justification, whereas it only points to superficial similarities; objectors to the teleological conception import assumptions about how consequentialist views in ethics work and about what they are committed to and assume that teleological views in epistemology must countenance – and contend with – similar commitments and results. In doing so, critics of epistemic teleology superimpose a framework on the latter view that is alien to it. But once the teleologist grants this step to the critic, it may well be that the latter's objections succeed. I will argue that the teleologist should not let her adversary decide the grounds on which the debate will take place. She should instead reject the assumption that there are deep structural similarities between understanding justification in terms of conduciveness to truth and understanding right action in terms of promoting what is of value.

And she would have good reason to do so. I will argue that there is no epistemological correlate to the notion of value as it appears in consequentialist theories in ethics; attempts to understand 'epistemic value' as a substantive, weight-bearing notion that would play the same function in teleological epistemology as 'the good' does in ethical consequentialism do not hold water. Thus, the anti-teleologist's use of the notion of epistemic value as if it were a substantive notion is unwarranted, and objections that rest on this use fail. This point should also give pause to the epistemic teleologist who is tempted to construe their view as the epistemological equivalent to ethical consequentialism.

Here is the plan of the paper. First, I will explain in more detail the teleological conception of justification, while presenting considerations in its favor and illustrating its preponderance in contemporary epistemology (§1). I will then explore the analogy with ethical consequentialism in more detail (§2). In §3, I will argue that the notion of epistemic value should not be understood as a substantive, weight-bearing notion in the same way that the notion of the good is in ethical consequentialism. Finally, I will turn to the main objections to the teleological conception of epistemic normativity, and show how they rest on illicitly importing assumptions about how value behaves in ethics into epistemology (§4).

1.1 Teleological views of epistemic normativity

According to teleological views of epistemic justification, what makes a belief justified in to be explained in terms of truth-conduciveness.[1] In this section, I will explain in more detail what the teleological conception of epistemic justification consists in. While doing so, I will illustrate how widespread its acceptance has been among contemporary epistemologists and provide initial considerations in its favor.

A belief is *justified* if and only if it is formed in accordance with what epistemic norms prescribe, permit, or forbid, and *unjustified* otherwise.[2] By 'epistemic norms', I mean the set of norms of which includes, for instance, norms specifying how one ought to respond to one's evidence, or norms forbidding one to form one's beliefs on the basis of wishful thinking are usually taken to be paradigmatic members. Because a belief being justified and a belief being formed in accordance with epistemic norms are two sides of the same coin, I will henceforth use 'teleological theories of epistemic normativity' and 'teleological theories of justification' interchangeably.

It is natural to think that facts about the contents of our epistemic norms (that is, what they prescribe, permit, or forbid) are not brute facts, but that they can be explained; in other words, that there are *reasons why* our epistemic norms have the contents that they have. In addition, it seems very plausible that there is a close connection between justification, on the one hand, and truth, on the other hand. The teleological conception of epistemic normativity provides one explanation of these contents, one that seeks to capture the essential connection between justification and truth. In essence, according to the teleological conception of epistemic normativity, what the contents of our epistemic norms are to be explained in terms of conduciveness to true belief; what our epistemic norms condone are beliefs that are arrived at in such a way that they are (sufficiently) likely to be true, either from the agent's perspective or objectively, and what they forbid are beliefs formed in such a way that they are not (sufficiently) likely to be true. Consider, for instance, a norm prescribing that agents form their beliefs in accordance with their evidence: why is it the case that a belief formed in accordance with one's evidence is justified? The teleolo-

1. For simplification purposes, I will be mostly speaking of teleological theories in terms of conduciveness to true belief. However, what I will say also applies to theories that conceive of justification in terms of conduciveness to knowledge.
2. Some epistemologists, generally those of a more internalist bent, use 'rational' as equivalent to 'justified.' I do not want to take a stance on the question of the relation between justification and rationality; I will treat these terms as equivalent for the purposes of this paper. Furthermore, we will be mainly concerned with doxastic justification. Doxastic justification concerns the grounds on which a belief is held: in Ichikawa's and Steup's terms, a belief is doxastically justified if it is 'held appropriately.' In contrast, *propositional* justification concerns whether an agent is in a position to form a justified belief about *p* (Ichikawa and Steup 2018).

gist answer is, in essence, that it is because forming one's beliefs in accordance with the evidence is likely to yield true beliefs.[3]

So far, I have characterized teleological views of epistemic normativity as views according to which the content of epistemic norms is to be explained in terms of conduciveness to truth. However, teleological views of epistemic normativity are often construed as a conjunction of two claims. The first is a claim about epistemic value: true belief (or knowledge) is epistemically valuable. The second is a claim about the proper explanation of the contents of our epistemic norms: what epistemic norms prescribe, forbid and permit is to be explained in terms of what is conducive to, or promotes, these things that are of epistemic value.[4]

The teleological conception of epistemic normativity cuts across the traditional internalist/externalist divide. The best way to characterize internalism and externalism is controversial; for present purposes, I will call 'internalist' theories of justification according to which one is justified in believing p just in case the one either is aware of, or is in a position to recognize, the factors that make one justified in believing p and of why they are justification-conferring. In contrast, a theory of justification is externalist if the subject need not be aware of, or in a position to recognize, the factors that make her belief that p justified for that belief to be justified. Conduciveness to truth may have a broadly externalist flavor when it is cashed out as a matter of the objective reliability of a belief-forming process, but it may also have an internalist reading: on the latter, what matters is that, *by an agent's own lights*, one's belief is formed in a way that is sufficiently likely to be true.

As Selim Berker has pointed out, the teleological conception of epistemic normativity has been tacitly accepted as default for most of the past half-century in epistemology; from externalists about justification such as Alvin Goldman to proponents of thoroughly internalist accounts such as Laurence BonJour and Richard Foley, virtually every major epistemologist has operated within the framework of a teleological approach to epistemic normativity.[5] Alvin Goldman, who is responsible for introducing reliabilism about justification to contemporary epistemology, argued that what makes a belief justified is that it is the product of a cognitive process which is sufficiently likely to yield true beliefs. On the other side of the spectrum, we find BonJour, who, despite its qualms with externalist theories of justification, writes that

> The basic role of justification is that of a *means* to truth, a more directly attainable mediating link between our subjective starting point and our objective goal... If epistemic justification

3. 'Evidence' may itself be a normative notion which can't be explained unless we already have a theory of epistemic normativity. I meant to present this putative norm simply to illustrate the form that a teleological explanation of a norm would take.
4. See for instance Ahlstrom-Vij and Dunn (2018b) and Friedman (2019).
5. Selim Berker provides an overview of the pervasiveness of the teleological conception of normativity in contemporary epistemology in Berker (2013a, pp. 350–357).

> were not conducive to truth in this way, if finding epistemically justified beliefs did not substantially increase the likelihood of finding true ones, then epistemic justification would be irrelevant to our main cognitive goal and of dubious worth... Epistemic justification is therefore in the final analysis only an instrumental value, not an intrinsic one.[6]

Since he explicitly conceives of justification in terms of truth-conduciveness, BonJour is committed to the teleological view of epistemic normativity, and in this respect at least, finds himself on common ground with the reliabilist.

The point he raises in the passage quoted above is echoed in a recent paper by Miriam Schoenfield, who writes

> My own view is that rationality is important because truth is important, and principles of rationality are meant to help us in our pursuit of the truth. So I'm inclined to think that if a proposed principle of rationality can't be motivated from the truth-seeker's perspective, that should cast suspicion on the principle (Schoenfield 2022).

Finally, here is William Alston:

> We evaluate something epistemically ... when we judge it to be more or less good or bad from the epistemic point of view, that is, for the attainment of epistemic purposes (Alston 2006, p. 29).

What are these purposes? According to Alston, they are to form true beliefs about things that matter to us:

> I don't know how to prove that the *acquisition, retention, and use of true beliefs about matters that are of interest and/or importance* is the most basic and central goal of cognition (Alston 2006, p. 30).

On Alston's view, beliefs are deserving of a positive epistemic appraisal when they help us attain our cognitive goals, the primary one being that of having true beliefs regarding matters about which we care. This fits squarely within the teleological conception of epistemic normativity.

This brief historical survey is, of course, fragmentary. However, I hope that it gives the reader some idea of how widespread this conception of epistemic normativity has been in contemporary epistemology; as Berker has said, "somehow in the past few decades a teleological approach has become so uncontroversial that it is now not even up for debate" (Berker 2013a, p. 357).

The dominance of the teleological conception of epistemic normativity can partly be explained by the fact that it has been simply been *taken for granted* by many epistemologists, as Berker suggests. It is true that many of the epistemologists quoted above do not discuss the tele-

6. See BonJour (1985, pp. 7–8), quoted in Berker (2013a, p. 35).

ological conception as if they thought it needed to be defended. But there are good reasons to favor it.

a) *No brute fact.*

   First, as mentioned above, we usually don't take facts about what we should believe as brute facts – that is, as facts that are not grounded in any further, 'more fundamental' facts; we rather take it that our epistemic norms can be explained by appealing to other considerations, considerations in which *truth* or *knowledge* play a preeminent role. There is an air of fetishism to claiming, for instance, that one ought to form their beliefs in accordance with the evidence, while providing no explanation for this norm; what the evidence supports seems to matter and be a proper object of concern only derivately, because truth matters, when it comes to forming beliefs. The teleological conception of epistemic normativity is well-placed to capture this, as it provides us with a prima facie plausible explanation of our epistemic norms, and connects them appropriately with truth.

b) *The relevance of reliability.*

   Relatedly, when assessing a candidate belief-forming process, method, or rule, it is always relevant to ask whether this process, method, or rule is likely to yield true beliefs (or is likely to be conducive to some other epistemic end, such as knowledge or accuracy) (Audi 1988). The debate in formal epistemology about whether ideally rational agents should revise their beliefs in light of new information they receive by conditionalizing on that new information provides a good illustration of this fact. Arguments in favor of conditionalization have aimed to establish that conditionalizing on the new information one receives *maximizes expected accuracy*; those who have argued that conditionalization does not maximize expected accuracy have taken this to show that an ideally rational agent would not conditionalize on their evidence, but would instead adopt some other belief-revision procedure on the grounds that such a procedure does maximize expected accuracy (Greaves and Wallace 2006; Schoenfield 2017). That it is always relevant to ask, of any given belief-forming process or rule, whether it is likely to be conducive to true beliefs is easily explained by teleological accounts of epistemic normativity: since epistemic norms just are norms such that following them is sufficiently conducive to true beliefs, then, in order to determine whether a belief-forming process or rule is an appropriate one, what we have to do is simply to establish its truth-conduciveness.

c) *The function of epistemic evaluations.*

   A further reason that speaks in favor of the teleological view is that it dovetails nicely with recent accounts of the function of our epistemic evaluations. According to these accounts, which take inspiration from Edward Craig's *Knowledge and the State of Nature*, epistemic concepts such as *knowledge*, *rationality*, or *justification* can be explained, at least in part, in terms of their functional role in the epistemic life of groups.[7] More precisely, on these accounts, assessing someone's belief as justified serves to flag perceived reliable sources of testimony, whom members of the community may trust in forming their own beliefs.

d) *The non-accidentality condition on knowledge.*
Gettier cases have been taken to show that accidentally true belief cannot be knowledge. In Gettier cases, one has a justified belief, and that belief is true, but there's no relation between what makes one's belief justified and the truth of that belief (Schafer 2014). Understanding justification in terms of truth-conduciveness may offer a way to capture this non-accidentality condition while broadly keeping with the justified true belief account of knowledge: if a belief formed by the application of sufficiently truth-conducive processes is true, it isn't true by luck - there is a necessary connection between the process by which it is formed and truth, that of being *sufficiently conducive to it*.[8]

e) *No alternative.*
Finally, it is worth noting that, in the years that have passed since Berker launched his much-discussed critique of teleological epistemology, the number of attempts at developing alternative views of epistemic normativity has been minuscule, and none of these attempts seems to have gained any widespread support.[9]

Thus, there are good reasons to adopt a teleological conception of epistemic normativity. Yet, teleological accounts face serious challenges. Before discussing these challenges, however, I want to explore the parallel with ethical consequentialism in more detail.

1.2 Teleological theories, ethical and epistemological

So far, I have explained the teleological conception of epistemic normativity and illustrated how pervasive it has been in contemporary epistemology. I now turn to one crucial element of the discussion of teleological epistemology: the analogy with consequentialism in normative ethics. The teleological conception of epistemic normativity is frequently analogized to ethical consequentialism, and both proponents and opponents have taken this analogy to reveal deep features and commitments of the view.

Roughly, consequentialist theories in normative ethics, which include act- and rule-utilitarianism, attempt to explain what one ought to do in terms of what promotes things which are of value, be it happiness, pleasure, preference-satisfaction, etc. More precisely, consequentialist theories generally have the following structure: they identify some state of affairs that is deemed of

7. The *locus classicus* of this view is Craig (1990). For more recent developments, see Henderson and Greco (2015).
8. That the reliabilist's understanding of justification successfully captures the non-accidentality condition is contested; see Schafer (2014, pp. 378–380).
9. In fact, I only know of one full-fledged attempt at providing an explicitly non-teleological account of epistemic normativity, found in Sylvan (2020a).

final value, or, in other terms, good in and of itself. For instance, classical utilitarianism identifies the type of state of affairs <S is feeling pleasure> as being the sole type of state of affairs that is of final value, and identifies <S is feeling pain> as the sole type of state of affairs that is of final disvalue. What we ought to do, consequentialist think, is to be explained in terms of what best promotes these states of affairs which are of final value, what best promotes what is good. To continue with the example of classical utilitarianism: classical utilitarians think that what makes an action right is that it is the action which, among all possible actions available to subject S at time t, is the one that maximizes pleasure and minimizes pain. Because they derive a theory of right action from a theory of value, consequentialist views are often characterized as being ethical theories on which *the good* is prior to *the right*.

How, then, is epistemic teleology analogous to ethical consequentialism? On one way of construing the teleological conception of epistemic normativity presented above, such a view seems to have the same structure as consequentialist theories in ethics. The teleologist identifies something that is of final (epistemic) value, such as true belief or knowledge, and then attempts to explain what makes a belief justified or rational – justification being often understood as a deontic property – in terms of what is conducive to that which is of final (epistemic) value. Put another way: what we ought to believe is what is conducive to what is of epistemic value. We may push the analogy further by claiming that, as ethical consequentialists see the good as prior to the right, so do teleologists in epistemology see the epistemically good as prior to the epistemically right.[10]

To explore this parallel in more depth, I will present two teleological views of epistemic normativity: process reliabilism, and a version of accuracy-first epistemology in the spirit of Joyce (1998), Pettigrew (2016) and Schoenfield (2015). I will explain how each of them can be seen as counterparts to forms of ethical consequentialism. In its initial iteration, process reliabilism consists in the two following claims:

> (A) If S's belief in *p* at *t* results ('immediately') from a belief-independent process that is (unconditionally) reliable, then S's belief in *p* at t is justified;
>
> (B) If S's belief in *p* at *t* results ('immediately') from a belief-dependent process that is (at least) conditionally reliable, and if the beliefs (if any) on which this process operates in producing S's belief that *p* are themselves justified, then S's belief in *p* at *t* is justified (Goldman 1979).

A bit of terminology: a belief-forming process is a procedure that maps certain inputs (perceptions, beliefs, etc.) to outputs – and specifically, to states of belief. Examples of such processes include memory processes, which takes as inputs beliefs and experiences from earlier time and outputs beliefs at a later time, and perceptual processes, which takes as inputs perceptual expe-

10. See, for instance, Ahlstrom-Vij and Dunn (2014, p. 1); Ahlstrom-Vij and Dunn (2018b, p. 2).

riences. By 'immediately', Goldman means that the belief must be the immediate output of the belief-forming process. A process is *conditionally reliable* when the beliefs it yields are generally true *given that the beliefs that were used as an input were true*. In contrast, a process is *unconditionally reliable* when the beliefs it yields are generally true. For our present purposes, nothing hinges on this distinction, and we can operate using the following simplified version of the view:

> **Process reliabilism:** If S's belief in p at t results from a reliable belief-forming process (or set of processes), then S's belief in p at t is justified.

Briefly stated: according to the reliabilist, a belief is justified if it results from a belief-forming process that is sufficiently truth-conducive.

It has been customary, in discussions of epistemic teleology, to see process reliabilism as the epistemological counterpart to a form of objective satisficing indirect consequentialism.[11] A consequentialist theory is *objective* if it evaluates acts or rules in terms of their actual consequences; it is *subjective* if it focuses on the expected consequences of the acts or rules. In addition, a theory is *direct* if what is assessed in terms of whether it promotes the (expected or actual) good is the act itself; it is *indirect*, in contrast, if what is assessed in terms of the promotion of the good is something else, like rules upon which one may act. For instance, according to Bentham's classical utilitarianism, an action is right if and only if it is the one which, among all the alternatives, maximizes pleasure and the absence of pain. On this view, it is the action itself that is assessed in terms of its consequence. Thus, it is a form of direct consequentialism. By contrast, on a rule consequentialist picture, an action is right if and only if it was done in accordance with a rule whose general acceptance tends to maximize pleasure and the absence of pain. On this picture, it is not required for an action to be right that it be the one that maximizes pleasure; what matters is that *accepting the rule* does, overall, tends to maximize pleasure, even though it may not maximize pleasure on every occasion. What is assessed in terms of its consequences is the acceptance of the rule (or rather, the system of rules as whole), not the particular action. Thus, rule consequentialism is a form of indirect consequentialism.

A further distinction can be drawn between *satisficing* theories and *maximizing* theories. According to satisficing theories, an action is right only if it produces an outcome that is good enough (without necessarily being the best outcome), or if it is done in accordance with a rule that tends to bring about good enough outcomes (again, without necessarily being the best possible outcomes). By contrast, maximizing theories, as the name indicates, require that the action or rule that is assessed be the one which, among all alternatives, produces the best possible outcome.

In what sense is process reliabilism similar to satisficing indirect consequentialism? First, note that process reliabilism explains the content of our epistemic norms by appeal to how well

11. See for instance Percival (2002, pp. 132–135); Berker (2013a, pp. 349–350).

they conduce to true belief. True belief plays a structurally similar role to the good in consequentialist ethics. Additionally, according to process reliabilism, whether a belief is justified depends, not on the consequences of the belief itself, but on the consequences of the process by which it was formed. Thus, it is analogous with a form of indirect consequentialism. Furthermore, it is not required for a belief to be justified that the process by which it was formed be the most reliable of all processes available to one in given circumstances; process reliabilism simply requires this process to be sufficiently truth-conducive. Finally, process reliabilism is analogous to objective consequentialist theories, insofar as it focuses on whether processes actually lead in a reliable manner to true beliefs under normal circumstances, and not on whether they can be expected to lead to true beliefs in the circumstances. Thus, process reliabilism appears to be the epistemological counterpart of a form of satisficing indirect consequentialism.

If process reliabilism is the analog of satisficing indirect consequentialism, accuracy-first views in the spirit of Joyce, Pettigrew and Schoenfield may be the (imperfect) epistemological counterparts of maximizing act-consequentialist views. For accuracy-first theories, the locus of epistemic evaluation is an agent's credence function. A credence function is a function that maps the propositions an agent entertains onto numbers between 0 and 1, where 1 represents fully confident of the truth of the proposition and 0 represents being fully confident of the falsity of the proposition. Accuracy-firsters identify the *accuracy* of a credence function as being of final epistemic value. The greater the accuracy of a credence function, the higher its epistemic value. A credence function is perfectly accurate insofar as it assigns credence 1 to true propositions and credence 0 to false propositions. The distance of a given credence function from perfect accuracy is measured using a scoring rule, the most popular of which is the Brier score. Finally, the expected accuracy of a credence function c is the average of the accuracy of c at each world that might be actual, weighted by the agent's credence that this world obtains (Schoenfield 2017, p. 1159).

Roughly, accuracy-first epistemologists claim that one ought to have the credence function that maximizes expected accuracy. For instance, if a certain credence function c_1 scores better than other credence functions $c_2, ..., c_n$ regardless of whether the propositions to which it assigns credences are true or false, then it c_1 has a higher expected accuracy, and an agent ought to assign their credences according to c_1. In other words: if there is a credence function c that is guaranteed to be more accurate than any other credence function – if *c accuracy-dominates* other credence functions –, then the agent ought to adopt c. From this principle, accuracy-firsters derive various norms that one's credences should respect; for instance, if one can prove that a certain rule for updating one's credences maximizes expected accuracy, then agents ought to update their beliefs in conformity with what this rule prescribes.

Accuracy-first views are direct insofar as the credence function one ought to have is the primary locus of evaluation for value maximization. They are maximizing, since they require agents to have the credence function that is the most conducive to expected value, and not merely a credence function that is conducive enough, but not necessarily the most conducive, to expected value. Furthermore, accuracy-first views are subjective, since they see agents as having to maximize expected accuracy. In these respects, they are similar to subjective, maximizing, direct consequentialist views in ethics such as classical act-utilitarianism, according to which one ought to do the action which, among all alternatives, maximizes expected pleasure and minimizes expected displeasure.

1.3 'Epistemic value'

I have shown how teleological views of epistemic normativity share certain structural features with consequentialist theories in ethics. More often than not, epistemologists defending teleological views have willingly granted to their critics that the analogy with ethical consequentialism is not superficial, but reveals essential features of their view. They then have tried to argue that the unacceptable consequences highlighted by critics are nevertheless not entailed by their view.[12] Proponents of teleological views, I will argue, should not make this concession to their opponent.

A different comparison with ethics may here prove fruitful. Consequentialist detractors of deontological views have argued that deontologism is paradoxical in the following way. According to deontological ethical theories, there are 'agent-centered' restrictions, restrictions such that it is impermissible to violate them, even when doing so would prevent further violations of that same restrictions (Scheffler 1994, p. 409). For instance, a deontological view may entail that it is impermissible, under any circumstances, to kill someone. But what if that killing would prevent five other killings? Since the deontologist is committed to claiming that one killing is morally objectionable, they would have to say that it is impermissible even in the case where it would prevent five other killings. But, the critic say, how could killing one to minimize morally objectionable acts of the same kind be itself morally objectionable? There is an air of paradox here. As Foot (1985) puts it: 'If it is a bad state of affairs in which one of these actions is done it will presumably be a worse state of affairs in which several are. And must it not be irrational to prefer the worse to the better state of affairs?'

One move available to the deontologist is to do some fancy footwork and try to argue that in fact, a state of affairs in which one killing happens is as bad, or worse, than one in which five killings happen. I leave it to the reader to judge the viability of this reply.[13] What is noteworthy, however, is that there is another move available to the deontologist, as Philippa Foot (1985) has

12. See for instance Ahlstrom-Vij and Dunn (2014), Sharadin (2018), and Daniel J Singer (2018, 2019).
13. See, for instance, Setiya (2018).

argued. This move is to reject *the very grounds* that give rise to the paradox in the first place. For the paradox to arise, we must accept that there are better and worse states of affairs as the consequentialist conceives of them. The consequentialist *imposes her own framework*, one on which there are better or worse states of affairs (and a particular understanding of what those are) on deontologism, and then argues that the latter is paradoxical. But it is the apparently innocent initial move by the consequentialist, that of characterizing some states of affairs as good or bad in themselves, from an impersonal point of view, that the deontologist ought to reject, argues Foot.

We can think of trying to answer the detractor of epistemic teleology by arguing either that epistemic teleology does condone the trade-offs, but they are, despite initial appearances, permissible, or that it doesn't condone the trade-offs *while* still being committed to a substantive notion of epistemic value as akin to the fancy footwork the deontologist may do in response to the paradox of deontology. The proponent of this move embraces the claim that the teleological conception of epistemic normativity is, at its core, structurally identical to ethical consequentialism. But it is only by conceding this claim that the problems that Berker and Friedman raise for epistemic teleology arise; as Foot writes, quoting Wittgenstein: 'The decisive movement in the conjuring trick has been made, and it was the very one that we thought quite innocent.[14]'

It is, of course, not sufficient to point out that the teleologist can reject the idea that the teleological conception of epistemic normativity has essentially the same structure as consequentialist views in ethics; we must still provide reasons for her to do so, and show that doing so blocks the objections raised by Berker and Friedman. I think that these reasons can be found by taking a closer look at the notion of epistemic value and at the role it plays in the teleological conception of epistemic normativity. In particular, I think that teleologists should avoid thinking of epistemic value as a weight-bearing notion, one whose role in epistemological theorizing is structurally identical to that of 'the good' in ethical consequentialism. Rather, the notion of 'epistemic value' that epistemic teleologists use is largely stipulative; we can view it as a spandrel of the attempt to cash out the connection between justification and truth in terms of truth-conduciveness. The epistemic teleologist is free to claim that, since truth is not valuable in the substantive sense of the term, it need not behave in the way in which it behaves in ethics. The rest of this section will argue that attempts to put flesh on the idea of epistemic value fail.

I will focus on Timothy Perrine's recent paper, 'Conceptions of Epistemic Value' (Perrine 2021), as it provides both a useful classification and a survey of extant conceptions of epistemic value. Perrine himself argues against what he calls the 'formal' conception of epistemic value in favor of a substantive conception of epistemic value, but although his argument against what he calls the 'formal' conception of epistemic value is persuasive, I will argue that we should reject

14. The quote is from Wittgenstein (2009, §308) and is found in Foot (1985).

the substantive conception of epistemic value. There is significant difficulty in understanding 'epistemic value' as a kind of 'value' in any sense of the term.

Philosophers who have written about epistemic value have typically been concerned with final epistemic value, rather than instrumental value. That is, they have been concerned, not with what is epistemically good insofar as it leads to other things that are of epistemic value, but rather, about what is good 'in and of itself', from the epistemic point of view, if anything is. Typically, it is knowledge or true belief that have been thought to be the bearers of final epistemic value, with understanding occasionally included into the group.[15]

Setting aside for a moment worries about the viability of the notion of something being good 'in and of itself', let's follow Perrine and say that something is of final value if and only if it is 'appropriate or correct to value [i.e., to have a favorable pro-attitude towards it] it for its own sake' (Perrine 2021, p. 216). The question we are asking is the following: is epistemic value a kind of final value? If not, in what sense is it a kind of value?

1.3.1 Epistemic value is not value, *simpliciter*

It seems unlikely that epistemic value is a kind of final value, *simpliciter*. If epistemic value were a kind of final value in this sense, then it would be appropriate or correct to value what is of epistemic value for its own sake. But many true beliefs or instances of knowledge concern things that are of absolutely no importance; there seems to be nothing wrong or defective about not caring about believing or knowing the truth about these matters. Consider, for instance, a belief about the number of blades of grass on the lawn of the Parliament building in Ottawa, or the 323rd entry in the late Wichita, Kansas phone book. Neither of these subject matters are of any importance, and a belief about them, be it true or false, is unlikely to be of any import, be it practical or theoretical. If epistemic value were a kind of final value, *simpliciter*, then just anything that is of epistemic value would have to be of final value. The examples provided above show that this is not the case; thus, epistemic value is not a kind of final value, *simpliciter*.

One could reply that believing truths regarding phone books or blades of grass has final value, but that given our ordinary concerns, we always have overriding reasons not to care about them. For instance, since such truths are of no practical import, and there are other truths which would be useful for me to believe, then, despite the fact that *pro tanto* reasons to care about them, I do not have decisive reason to do so. If that were true, then it would be the case that truths about insignificant matters of fact are of final value – it is simply that they would be of extremely low value.

15. See for instance the debate between Marian David (2014) and Jonathan Kvanvig (2014).

This line of reasoning is ultimately not persuasive. Imagine that I were an immortal being, free of any practical concerns and possessing unbounded cognitive capacities. Should I seek out the truth about just *any* matter of fact? Should I care about the 323rd entry in the Wichita, Kansas phone book, or about the number of grains of sand on the beach in Ipswich, or just any similarly insignificant truth? If truths about even the most insignificant things were of final value, and I was facing eternity, then it would be that I ought to try to believe them; otherwise, I would be failing to respond properly to their value. I take it that there is little to envy about such a life; this should make us seriously doubt that true beliefs about insignificant matters of fact have final value, *simpliciter*.

Another line of reply would be to claim that the fact that I should prefer to be right than wrong about these matters of fact show that truths about blades of grass or phone books have final value. However, even if it were the case that I should prefer to be right than wrong when forming beliefs about things, this still wouldn't show that *I ought to have a favorable pro-attitude* towards just any truth. Consider the following situation: I have no use for a toaster; no interest of mine would be furthered by acquiring a new toaster. Yet, it seems plausible that if I actually were to acquire a new toaster, I should prefer one that fulfills its function well rather than one that doesn't. But from this fact, it doesn't follow that *I ought to care about* having a toaster that works well. Similarly, I have no interest that would be satisfied by having a true belief about the 323rd entry in the Wichita, Kansas phone book. It may be true that were I to form a belief about it, I should prefer it to be true. But that doesn't mean that I have reason to care about it for its own sake, as the proponent of the view that epistemic value is a kind of value, *simpliciter*, claims.

The preceding considerations make a strong case for thinking that being of epistemic value is not a kind of final value, *simpliciter*. One may disagree. However, importantly, this notion of epistemic value as a kind of final value, *simpliciter*, is simply not equivalent to the notion of epistemic value as it appears in teleological theories of epistemic normativity, despite the fact that final value, *simpliciter*, is the notion of value that figures in consequentialist theories in ethics. If the contents of our epistemic norms were explained by conduciveness to epistemic value and true beliefs had final value, *simpliciter*, then we should expect our norms to be sensitive to the degree to which true beliefs have final value. Justification would be easier or harder to achieve depending on the subject matter of the belief.

Here is why. Not all true beliefs have the same degree of final value: for instance, we should care more about having true beliefs about loved ones than about entries in a defunct phonebook. The former are undeniably more important to us. Thus, true beliefs about loved ones have much more final value than true beliefs about entries in the phonebook. Now, suppose that the teleologist were contending that the contents of epistemic norms is to be explained in terms of conduciveness to final value, *simpliciter*. We should expect our assessments aimed at establishing whether a given belief is justified to take into account how much final value the belief would have if true. It could, for instance, be easier to justify a belief that if true would have great final value than to justify a belief about a insignificant matter of fact, since believing the truth about the important matter could have much greater expected final value than believing the truth about the insignificant matter of fact.[16] However, as far as I know, no proponent of the teleological account of epistemic normativity embraces this claim. For instance, according to the process reliabilist, what determines whether a belief is justified is only a function of the reliability of the belief-forming process that produces it, and belief-forming processes are not individuated by subject-matter – they are, in Goldman (1979)'s term, 'content-neutral.' The conditions that must be met for beliefs formed through perceptual processes to be justified are the same whether the belief concerns loved ones or phone book entries. Thus, according to the process reliabilist, what determines whether a belief is justified is not sensitive to its degree of final value. In other words: it is a matter of complete indifference to the teleologist how much final value a belief has.

Summing up so far: (1) epistemic value is not a kind of final value, *simpliciter*, and (2) epistemic value as it appears in teleological theories is not taken to be a kind of final value, *simpliciter*.

1.3.2 Against the formal conception of epistemic value

Anticipating the problem raised by beliefs about insignificant matters of fact, many authors have opted for a *formal* conception of epistemic value. According to this conception, epistemic value is not a kind of final value, nor does being of epistemic value entail being of value, *simpliciter*. Here's Ernest Sosa on epistemic value:

> We humans are zestfully judgmental across the gamut of our experience: in art, literature, ..., and so on, across many domains... Any such domain ... admits values of two sorts: the derivative, and the fundamental – that is to say, the derivative or fundamental *for that domain*. A value might be irreducible to other values distinctive of a given domain, without being fundamental absolutely (Sosa 2007, p. 70).

Sosa goes on to claim that he won't settle the question of whether truth is of final value, *simpliciter*, but that this is not important anyway, since truth is finally value *from the epistemic point of view*, regardless of whether it is a kind of final value. Thus, he seems committed at least to this formal conception of epistemic value, as is Duncan Pritchard, who writes

> [w]hen I say that truth is the fundamental epistemic good, I mean that from a purely epistemic point of view it is ultimately only truth that we should care about... A key point to note about treating truth as a finally epistemically valuable good is that it does not follow

16. More precisely: this would be the case if the final value of the true belief about the important proposition is much greater than the final disvalue of believing a falsehood about it.

> that truth is finally valuable *simpliciter*... The point is just that from a purely epistemic point of view truth is rightly to be valued non-instrumentally. But this is entirely consistent with arguing that the general view of truth is merely instrumental (or indeed, that it has no value at all.[17]

Authors embracing the formal conception of epistemic value often express their view by saying that truth is valuable 'from the epistemic point of view', or that it is valuable 'in the epistemic domain'; they think that epistemic value is in some sense a kind of final value, despite it not being the case that we should value things that are of epistemic value for their own sake. But, as Perrine points out, it is very hard to make sense of this claim. Here are two different ways of trying to cash out the sense in which being of final epistemic value is being of final value, without being of final value, *simpliciter*:

A. If X is of final epistemic value, then it is appropriate to epistemically value X for its own sake.
B. If X is of final epistemic value, then it is epistemically appropriate to value X for its own sake.

Claim **A** is implausible: as Perrine remarks, it seems committed to the claim that there is a distinctive kind of attitude of valuing, of having a pro-attitude towards something, that is 'epistemic' – one that potentially stands in contrast with 'ethically' valuing or 'politically' valuing, etc. But there does not seem to be such a range of distinctive kinds of valuing (or at least, it is very difficult to make sense of this idea). Therefore, **A** is not a promising way to capture the idea that being of epistemic value is being of value without being of value, *simpliciter*.

What about claim **B**? At least, it is not committed to there being a distinctively epistemic kind of valuing. Yet, it suffers from another problem, which concerns how to understand an attitude's being 'epistemically appropriate'. Here's Perrine trying to flesh out this idea:

> A pro-attitude directed towards something is epistemically appropriate if and only if it is appropriate to direct pro-attitudes towards that thing for their own sake and the attitude has property P (where 'P' names some property that explains what makes the attitude epistemically appropriate) (Perrine 2021, p. 219).

The problem with this way of understanding 'epistemically appropriate' is that it is incompatible with the formal conception in the first place; it entails that being of final epistemic value is being valuable for its own sake, which the proponent of the formal conception denies, in addition to being implausible, as I have argued. A last-ditch attempt considered by Perrine is the following:

17. See Pritchard (2014, p. 113); quoted in Perrine (2021, p. 5). See also Ahlstrom-Vij (2013).

> If one is using the epistemic point of view, then it is appropriate to value for its own sake what the epistemic point of view identifies as being of final epistemic value (Perrine 2021, p. 219).

How should we understand the notion of an 'epistemic point of view'? Perrine thinks of the epistemic point of view as 'expressing the standards for epistemic evaluations' (Perrine 2021, p. 215). To occupy the epistemic point of view, then, would be to be assessing things using the standards for epistemic evaluation. A first problem with the idea of something being 'good from the epistemic point of view', one that is not identified by Perrine, is that we are not given any reason to think that the standards of epistemic evaluation contain any claims about things being of value in the first place. 'Good' does not seem to figure in our ordinary epistemic evaluations: we often evaluate beliefs as right or wrong, which more or less is equivalent with saying that they are true or false, and appraise beliefs using the terms 'rational' or 'justified' while criticizing others as irrational or unjustified. It seems strange to talk of beliefs being 'good', of 'excellent' beliefs, or of some beliefs being 'better' than others, in contrast with talking about good coffee or bad taste in music, for instance. Thus, it is far from obvious that the standards of epistemic evaluation contain claims about value.

Furthermore, such a view is not acceptable to a teleologist who wants to explain justification in terms of conduciveness to epistemic value. On this conception of epistemic value, the standards of epistemic evaluation – including the standards for justification – would come first in the order of explanation; we would be explaining what epistemic value is by appealing to *our evaluative standards*. But the teleologist wants to explain our evaluative standards by appeal to epistemic value. Therefore, she should reject such a conception of epistemic value.

A final worry about understanding epistemic value in that last sense, this one due to Perrine, is that, even if the epistemic point of view did identify some things as being of epistemic value and you were using that point of view, it wouldn't follow that it is appropriate to value for its own sake what is identified to be of final value by the epistemic point of view. There are a plurality of points of view that identify as being of final value things that it is plainly defective to value, whatever the point of view that one is using. From the point of view of the sadist, for instance, inflicting pain for one's personal gratification has final value. But there seems to be no sense in which it is appropriate to value inflicting pain in this way – and that holds even if we were occupying the sadist's point of view. Put another way: it does not follow from the fact that something is deemed to be of final value from a certain point of view that we really should care about it if we adopted this point of view, because in some instances, things that are deemed to be of final value from a certain perspective simply aren't proper objects of favorable pro-attitude. Therefore, from the fact that the epistemic point of view identifies something (e.g., true beliefs) as being of final value, it does not follow that it is appropriate to value this thing for its own sake, even if one is 'occupying the epistemic point of view.'

Thus, not only are there good reasons to reject the claim that epistemic value is a kind of final value, *simpliciter*, but there are also good reasons to reject understanding it as a weaker, more 'formal' notion of value.

1.3.3 What is epistemic value, then?

In the last section, I have argued that we should be doubtful that there is really such a thing as epistemic value in a substantive sense of the term 'value'. Neither a conception according to which epistemic value is a kind of final value, nor a conception according to which it is only value *from the epistemic point of view* hold water. One may think that this seals the coffin of the teleological conception of epistemic normativity, since it appeals to epistemic value. But this substantive notion of epistemic value is not essential to the teleological conception.

One indication that it is so can be found in how freely and fluently epistemologists who subscribe to the teleological account move between talk of epistemic value, epistemic goals, and epistemic ends.[18] None of these three terms, taken in the strictest sense, are close to be synonyms, let alone completely interchangeable. Yet, many epistemologists often talk as if nothing was really hanging on the choice of terms. Furthermore, the notion of epistemic value is completely absent from Alvin Goldman's early work on reliabilism, something that he himself has pointed out in response to criticism by Berker (it does, however, appear in later work in a way that could suggest that he subscribing to the picture of normativity that Berker attributes to him).[19] Finally, in contrast to the notion of a 'good state of affairs' as it appears in consequentialist ethics, we have no pre-theoretical grasp of the notion of epistemic value; this suggests that the use of the term is going to be somewhat stipulative.

Epistemic value, on the teleological conception of justification, is best understood as a 'placeholder' term, and not as a term denoting value in any substantive sense of the term. As stated above, it is a spandrel of the attempt to cash out the connection between justification and truth in terms of truth-conduciveness. Put another way: the use of the term 'value' is largely stipulative, and however epistemic teleologists conceive of 'epistemic value', they are free to think that it behaves in a very different way from value in ethics. The explanation of our epistemic norms

18. For instance, in "The Unity of the Epistemic Virtues," Goldman takes the aforequoted passage from BonJour to show that the latter is committed to veritism, according to which only truth is of final epistemic value, although the section of the passage that Goldman quotes only refers to the goal of having true beliefs. Goldman then goes on to restate veritism in terms of epistemic aims as if it were an innocent rephrasing (Goldman 2002, pp. 52–54).

19. For instance, '[t]he principal relation that epistemic virtues bear to the core epistemic value will be a teleological or consequentialist one. A process, trait, or action is an epistemic virtue to the extent that it tends to produce, generate, or promote (roughly) true belief.' (Goldman 2002, p. 52).

may very well be teleological in nature, but giving a teleological explanation of our epistemic norms does not commit one to the idea that true beliefs ought to be promoted or brought about in the way the good ought to be promoted according to ethical consequentialists – by bringing about as many of them as possible, or a certain quantity of them. I will offer an example of a teleological explanation of our epistemic norms that leaves no place for value aggregation while discussing Berker's objection.

Contrast this way of speaking about 'value' with a substantive or weight-bearing notion of value: when utilitarians speak of pleasure as being of final value, of pleasure being good in and of itself, they are using this notion of final value, or of the good, as a substantive notion. It is *because pleasure is good* that we ought to maximize it. In other words, in the ethical case, it is a normative property of pleasure (*being good*) that explains why we ought to promote it. In contrast, because of the difficulties involved in understanding epistemic value as a substantive kind of value, the epistemic teleologist should not be claiming that true beliefs are good in order to explain the contents of our epistemic norms. It is not because of the normative properties of true beliefs that our epistemic norms are explained in terms of truth-conduciveness[20] Justification is a guide to truth, but not a guide to the good.

Why, then, call it 'value'? Simply because it plays a structural role analogous to that played by value in consequentialist ethics: our epistemic norms are to be explained in terms of conduciveness to this thing. Likewise, the consequentialist takes our moral norms to be explained in terms of conduciveness to what is valuable. By noticing this feature, however, neither the teleologists nor her opponent has revealed some deep feature of teleological theories in epistemology. This role could be played equally well by using the terms 'epistemic goals' or 'epistemic ends', as long as we also understand them as placeholder terms.

One may worry that such an account of epistemic value is insufficient to ground epistemic normativity. The worry is that if truth is only valuable in a largely stipulative sense of the term 'value', then it is unclear why we ought to believe in accordance with what epistemic norms recommend; one may have hoped that the value of truth would be what gives us reason to conform to these norms. The teleological account of epistemic normativity, however, does not aim to answer the question of why we should believe in accordance with epistemic norms. It is an account of the *structure* of epistemic normativity: it states that this structure is teleological. This account

20. Goldman's process reliabilism was explicitly introduced as an attempt at explaining justification by appealing to non-normative notions (Goldman 1979); (Goldman 2015, p. 136); this point is also noted in Sylvan (2020b). The point I am making here goes beyond the claim that characterizing process reliabilism as explaining justification by appealing to a substantive notion of value is descriptively incorrect, although this is worth emphasizing. It is also that it is legitimate, for any teleologist, to use the term 'value' in a very different way from the way it behaves in ethics; that she, in fact, *should not* understand epistemic value as a substantive kind of value; and finally, that this move, which is at least partly independently motivated, blocks Berker and Friedman's objections.

is compatible with many different answers to the question of the source of the authority – of the normative 'oomph' – of epistemic norms. For instance, it may be that the fact that some true beliefs are valuable, theoretically and practically, is a reason for us to conform to epistemic norms in all cases. It could be that, as members of communities in which we can't escape relying on others' testimony, we owe it to others to be reliable testifiers, and that this places requirements on how we ought to go about forming and revising our beliefs. There is doubt that any of these proposals, even if true, would give us the result that we always have a reason to believe in accordance with epistemic norms. But it could partly explains why we take ourselves and others to be on the hook for following them.

1.4 Challenges to teleological conceptions of epistemic normativity

1.4.1 Objection 1: impermissible trade-offs

Perhaps the most influential objection to teleological accounts of epistemic normativity is that such accounts are committed to obviously wrong verdicts about whether a certain belief is justified in certain cases involving trade-offs between items of epistemic value. Teleological theories, it is said, have to deem as justified those beliefs that best promote the overall epistemic good, even when such beliefs are formed in a way that, intuitively, is sufficient to make them unjustified.[21]

Teleological theories in ethics are (in)famously known to countenance trade-offs between persons. Trade-offs cases arise when an doing a certain action would incur some disvalue, but best promote overall value – maximizing teleological views will yield the verdict that the agent ought to commit the action, regardless of what it involves. As the classic case goes: suppose that a doctor has a choice between killing one perfectly healthy patient to collect their organs and save five critically injured ones, and letting the healthy patient live and the five injured die. On a simple version of act-utilitarianism, on which what one ought to do is what maximizes expected utility, it appears that what the doctor ought to do is to kill the one to save the five. However, this verdict seems wrong: far from it being the case that the doctor ought to kill the healthy patient, it seems impermissible for the doctor to do so, even if doing so would save the five injured.

Selim Berker, among others, argues that teleological theories of epistemic normativity suffer from a parallel problem. They have to countenance trade-offs between items of epistemic value *with respect to different propositions*, when doing so is what best promotes the overall epistemic good. Consider the following case:

21. The initial version of this argument is found in Firth (1981). More recent and sophisticated versions are found in Berker (2013a,b) and Littlejohn (2012). I will primarily be discussing Berker's formulation.

> **God grant**
>
> You are an atheist seeking a grant from a religious foundation. This grant would allow you to pursue research and come to a great number of true beliefs about important matters of fact. However, you know that the foundation only funds researchers who believes that God exists. You have pondered over this question in the past, weighed evidence for and against believing that God exists, and have come to the conclusion that there is no God. *Should you believe that God exists?*

According to a crude version of epistemic consequentialism on which what one ought to believe simply is what best promotes the overall epistemic good, claims Berker, what you ought to believe is that God exists (and thus, such a belief would be justified). This is because the overall epistemic good is best promoted by forming the belief that God exists. But clearly, your belief would not be justified: it would be formed against what you take to be strong evidence that God doesn't exist. Thus, epistemic teleology ought to be rejected.

The arguments presented above have been directed at crude versions of epistemic teleology, according to which a belief is justified simply in virtue of being conducive to the greatest number of true beliefs.

In what follows, I argue that ceasing to take 'epistemic value' as a substantive, weight-bearing notion allows the teleologist to answer the trade-off objection. The gist of the argument is simple: Berker's objection against epistemic teleology rely on taking epistemic value as a substantive notion; it is based on the idea that epistemic value requires maximization or producing it in a way that satisfies some threshold because that is what thinking that something is valuable commits one to.

Here's Berker, claiming that placing restrictions on value and its promotion in order to avoid trade-offs problem would be a mere *ad hoc* move given that the teleologist is committed to a certain theory of value such that value requires maximization (or at least, requires to be promoted above a certain threshold):

> This restriction would, of course, immediately block any untoward results for the teleologist in the case of the grant-seeking scientist since in that example the scientist's belief in the existence of God is merely a causal means to the acquisition of various true beliefs and the avoidance of various false beliefs sometime in the future. But the question to ask is: why should we restrict what counts as conducing toward or promoting value in this way? What is the motivation, for the veritistic epistemic teleologist, in making this move (other than to avoid a potential problem)? On the epistemic teleologist's basic picture, certain items (such as true beliefs) have nonderivative epistemic value, and certain other items have derivative epistemic value insofar as they conduce toward the items of that first sort. What is the rationale for restricting this conducing relation so that it doesn't cover causal means? Nothing in the epistemic teleologist's basic picture suggests an answer to this question (Berker 2013a, p. 366).

This passage shows that Berker thinks that whatever restrictions the epistemic teleologist places on the conducing relation would be unmotivated *given her teleological commitments*. Why would that be? Because, as Berker sees it, the teleologist is committed to a a view norms are to be explained in terms of promotion of the good. It's easy enough to see how such a move would be arbitrary in the ethical case. Suppose that you are a classical utilitarian. You think that pleasure is the only thing that has final value, and pain final disvalue. In other words, pleasure is the sole thing that is good in and of itself. Pleasure calls for maximization; it's the kind of thing which one ought to bring about as much as one can. The more pleasure contained in an outcome, the better the outcome. It would seem irrational to prefer an option that brings lesser pleasure to an option that brings about more pleasure, because the latter would bring about a better state of affairs. But now, in the face of the objection that one ought to kill the one healthy patient to save the five, you postulate seemingly arbitrary restrictions to avoid the trade-offs problem. This would appear inconsistent with your commitment to pleasure being the sole fundamental good that one ought to promote.

In the epistemic case, things are different. True, it is the case that, in a certain respect, truth plays in epistemic teleology a role that is structurally analogous to the good in ethical consequentialism. But in another respect, this role is very different: in the ethical case, it is because whatever it is that the theory identifies as being of final value is thought to be good that we ought to do what is conducive to it. This is why restrictions on what is conducive to it that are tailor-made to avoid trade-offs objection seem *ad hoc*: if one thinks that only pleasure is good, and why shouldn't every optimal way of bringing it about be commendable? The epistemic teleologist is not (or should not be) claiming that epistemic norms are explained in terms of conduciveness to truth because *of the goodness of truth*. Because she is not claiming that truth is good in the substantive sense of the term, she is free to deny that epistemic value, as it figures in her theory, is the kind of thing that can be aggregated in the same way that value in ethics can, according to the ethical consequentialist. Put in a slightly different way: she is free to think that the conducing relation should be understood in such a way as to respect what Berker calls the 'separateness of propositions.'

One natural way of doing so would be to consider conduciveness to truth with respect only to the belief that is the direct result either of doxastic deliberation or of the application of a belief-forming process, rule, or method. On this picture, the question we should be asking, then, when assessing whether a belief is justified, would be the following: *was that particular belief* formed in a way such that *it* would have been likely to be true?[22] If the answer is yes, then, according to the teleologist, the belief would be justified. On this understanding of justification, justification is still a teleological notion, since a belief is judged to be justified based on the expected truth

22. Either *objectively* or *by the agent's own lights*, depending on whether one thinks that externalism or internalism about justification is correct.

or falsity of the outcome of the doxastic deliberation or of the application of the belief-forming method. But since only *that particular belief* is assessed for whether it was likely to be true, there is nothing between which there could be trade-offs; thus, the trade-offs objection cannot even take off the ground. The reliabilist subscribes to a picture of this kind.

Yet, Berker thinks that any version of the teleological conception of epistemic justification is vulnerable to the trade-offs objection. However, in contrast with the crude teleological view presented above, the reliabilist has to countenance unacceptable trade-offs at the level of *belief-forming processes*, rather than at the level of individual beliefs. The case of trade-offs Berker presents against reliabilism is as follows: consider the belief-forming process <whenever considering whether a certain natural number n is prime, believe that n is not prime>. This process is highly truth-conducive, since most natural numbers aren't prime numbers. Thus, the reliabilist has to say that any belief formed by this belief-forming process is justified. But clearly, if I consider the question of whether 7 is prime, I would not be justified in believing that it is not prime as a result of applying this belief-forming method. Once again, claims Berker, the trade-offs problem resurfaces: 'the false beliefs [that this belief-forming process] occasionally yields are sacrifices for the (much) greater epistemic good' (Berker 2013a, p. 375): one will deem some unjustified beliefs as justified because the method by which one arrived at them is conducive to a greater number of true beliefs on the whole.

In another passage, Berker again suggests that the reliabilist is placing *ad hoc* restrictions on the conducing relation:

> ... the process reliabilist's restriction of the conducing relation is undermotivated, given his or her general teleological commitments. Why, exactly, are nonproximate causal means banned from being a way of promoting epistemic value, whereas proximate causal means are not? Such a fine distinction has an air of arbitrariness to it (Berker 2013a, p. 375).

The first thing that the teleologist should reply to Berker regarding this case is that it is simply not true that the restrictions she places on the conducing relation – according to process reliabilism, only the beliefs that are the direct outputs of a belief-forming process are to be taken into account when assessing whether a belief-forming process is reliable – are arbitrary. The same reasoning as before applies to this aspect of Berker's criticism.

But moreover, the truth is, the prime numbers case does not appear to be a problem involving trade-offs. The main problem raised by this case concerns how to individuate belief-forming process. It is dubious that there is a cognitive process that serves only to form beliefs about primes in the way that Berker's objection presupposes. In this respect, Berker's objection is closer to the generality problem for reliabilism. The generality problem is the problem of specifying in a principled manner what the belief-forming process that led to the formation of a given belief is. The generality problem, however, is not a problem that reliabilism faces in virtue of being a teleological theory.[23] Any view of justification on which whether a belief is justified depends on features of the belief-forming process, rule, or method by which it was produced is *prima facie* susceptible to this problem: it could be that, if we individuate the belief-forming process of which belief B was an output in a certain way, B would count as justified, because that belief-forming process has the justification-making features, whereas if we individuate the belief-forming process in a different way, B would count as unjustified, because the belief-forming process thus individuated lack the features that make its output justified.

Berker could still reply that, once we specify the correct process – one that would serve a broader function than just forming beliefs about primes –, there would still be a trade-offs problem; reliabilism would condone some beliefs that are the outputs of these properly individuated belief-forming processes as justified even though they aren't. These beliefs would still be 'sacrifices for the greater epistemic good.' However, when we look at processes that have the required level of generality, it does not seem like we are dealing with trade-offs anymore. Consider the paradigmatic belief-forming processes that Goldman identifies in 'What is Justified Belief?': forming a belief on the basis of memory, perceptual processes, and reasoning processes such as inferring that p on the basis of antecedent beliefs. In addition, according to Goldman, these processes have to be content-neutral (which would *de facto* exclude Berker's purported process) (Goldman 1979, 2015). Of course, some beliefs that are the outputs of perceptual belief-forming processes, for instance, will be false, but justified. But it would be a stretch to call them 'sacrifices for the greater epistemic good' for that reason: any non-factive theory of justification will entail that sometimes, false beliefs are justified.[24]

I suspect that the main reason there appears to be a trade-off in the prime numbers case is that we think that there is an obviously better belief-forming process available to the agent that the agent was in a position to use when forming a belief about whether 7 is prime – and thus, we have the intuition that the agent's belief in this case is not justified. Since 7 is a small number, one could have easily attempted trial division to test for primality, and trial division would have been a more reliable way to form about whether 7 is prime. But Goldman's formulation of process reliabilism in 'What is Justified Belief?' explicitly specifies that, for a belief that p that is the output of a belief-forming process P to be justified, it must not be the case that there is an alternative reliable belief-forming process Q such that Q would have produced a different belief regarding whether p. This would rule out the belief that 7 is prime formed as a result of applying the rule 'for any prime number n, believe that n is prime' from counting as justified by the reliabilist's own lights.

23. This point is also made in Ahlstrom-Vij and Dunn (2018a).
24. See also Ahlstrom-Vij and Dunn (2014, p. 6).

This rounds out my case against the trade-offs objection to epistemic teleology. I now turn to a different line of criticism, put forward by Jane Friedman in her paper 'Teleological Epistemology'. The answer to Friedman's objection is broadly of the same form as the answer to Berker's objection.

1.4.2 Objection 2: our familiar epistemic norms are not teleological

In her paper 'Teleological Epistemology', Friedman argues that a closer look at our ordinary epistemic norms reveals that these norms are not teleological. If they were, they would be of a very different form from the form they have. The objection is as follows: if epistemic teleology were true, then our ordinary epistemic norms would demand that we maximize or optimize our acquisition of things that have epistemic value. However, our ordinary epistemic norms do not demand that we maximize or optimize our acquisition of things that have epistemic value. Therefore, epistemic teleology is false.

Why think that if epistemic teleology were true, then our ordinary epistemic norms would demand that we maximize or optimize our acquisition of things that have epistemic value? Friedman says that

> ...once we see subjects as in pursuit of some valuable or valued end there is pressure to think of the relevant subjects as having to maximize or optimize with respect to that end. A subject who is trying to acquire knowledge ought to get as much of it as they can going forward; anything less seems irrational (Friedman 2019, p. 681).

In other words, Friedman claims that the proper response to value is to maximize it. This seems plausible when we consider ordinary cases involving valuable things. Suppose that I'm trying to acquire money, and I'm faced with a choice between acquiring one dollar, ten dollars, and a hundred dollars. If I am a rational agent, then, *ceteris paribus*, I ought to choose the hundred dollars, as this is what best promotes my end of acquiring money.

Similarly, if epistemic teleology is true, then it seems that we ought to see believers as being required to maximize with respect to the acquisition of whatever it is that has epistemic value. But then, it seems that our ordinary epistemic norms are not teleological, because they do not require the maximization of value.

To make this argument, we need a way to quantify how much epistemic value a belief yields, or at least a way to rank the various beliefs one may form in any situation. Let's follow Friedman's way of doing so, which is anchored in a proposition's logical strength. Say that all true beliefs that have an atomic proposition as their object are worth the same amount of epistemic value. A true belief in a conjunction of these atomic propositions is worth more than a true belief in just one of the conjuncts. A true belief in a conjunction of a conjunction of atomic propositions is worth even more, and so on. In any given situation, there will be one or more strongest proposition that one could come to know or truly believe, one such that believing it yields the maximum amount of epistemic value one could produce in this situation.

If our norms demand that we maximize with respect to the acquisition of true belief or knowledge, then our norms should demand that we believe this strongest true proposition we can believe in all circumstances in which we form a belief, or that we come to know the strongest knowable proposition in any circumstances. But this does not seem to be the case.

Consider the following norm: for any S, *p*, and t, S's coming to know *p* at t is permissible. There is nothing epistemically wrong or impermissible with coming to know something that one is in a position to know. Suppose that I am wondering where my water bottle is, and, glancing at my desk, I see that it is on it. It would permissible for me to come to know that my water bottle is on my desk. It is plausible that this norm (call it 'Weak Knowledge Norm') is among our ordinary epistemic norms.

Yet, if our norms were teleological, they should only permit coming to know the *strongest knowable proposition* in any circumstances in which one is forming a belief. Let's return to the previous example: when, looking for my water bottle, I glance at my desk, I could come to know that there is yellow highlighter in my pencil jar, that the water bottle is to the left of my computer mouse, that there is a coffee cup right next to it, etc. Nothing short of coming to know the strongest knowable conjunction of all these facts would be permissible. But this is not the case; according to our ordinary epistemic norms, it is always permissible to come to know any proposition one is in a position to know. Thus, our ordinary epistemic norms do not demand that we maximize or optimize with respect to what is of epistemic value – and since teleological views of epistemic normativity entail that our epistemic norms demand just that, we ought to conclude that they are false.

The objection above is directed at a maximizing version of epistemic teleology. However, as discussed previously, some teleological views of epistemic normativity are analogous to satisficing consequentialist views rather than maximizing views. Friedman thinks that her objection can also be made to target these views (Friedman 2019, pp. 687–688). The objection to satisficing epistemic teleology is as follows:

(1) if satisficing epistemic teleology were true, then our ordinary epistemic norms would demand that we satisfice with respect to our acquisition of things that have epistemic value.
(2) However, our ordinary epistemic norms do not demand that we satisfice with respect to our acquisition of things that have epistemic value.
(C) Therefore, satisficing epistemic teleology is false.

The motivation for the first premise is similar to that presented in the discussion of the version of the argument directed at maximizing views. If we see subjects as being in pursuit of some valuable end, and they are faced with the occasion to attain that end, then it would be irrational for them not to do so. On a satisficing view, claims Friedman, although subjects would not be required to maximize the amount of knowledge, they still would be required to form knowledge that has a degree of value that meets a certain threshold. In other words, if this account were correct, there would be a lower limit on how little one is permitted to come to know in any situation.

However, in most circumstances in which it would be possible to come to know things, it is permissible not to form any beliefs. When I look around me, I could come to know a great number of truths about my environment. Yet, it would be perfectly appropriate for me not to come to know any of these truths, and so it would be appropriate not to meet this threshold of knowledge. There would be nothing defective or wrong in doing so. Thus, concludes Friedman, our epistemic norms do not demand that we satisfice with respect to our acquisition of things that have epistemic value.

Just like Berker's objection to epistemic teleology, Friedman's argument rests on illicitly importing assumptions about how value behaves in ethics into the epistemological realm. In Friedman's case, this happens in the course of motivating the first premise of the argument: if epistemic teleology were true, then our ordinary epistemic norms would demand that we maximize, optimize or satisfice with respect to our acquisition of things that have epistemic value. Recall that Friedman motivated this premise by claiming that 'once we see subjects in pursuit of some valuable or valued end there is pressure to think of [them] as having to maximize or optimize [or at least satisfice] with respect to that end.'

But the teleologist can simply reject this claim. She is not attempting to explain our epistemic norms by appealing to the goodness of truth; rather, she is trying to explain the contents of our epistemic norms by appealing to facts about whether conforming to any of these norms on a particular occasion would (tend to) lead one to form a true belief. Again, as discussed in the case of the trade-offs objection, this explanation is still teleological, insofar as it appeals to the consequences of following these norms. Yet it does not make reference to truth being valuable. Relatedly, there is nothing defective or irrational in not maximizing, optimizing, or satisficing with respect to the number of true beliefs one could form on any occasion, because one is simply not thereby failing to respond to value in an appropriate way. Thus, it is not the case that if epistemic teleology were true, our ordinary epistemic norms would demand that we maximize, optimize or satisfice with respect to the acquisition of true beliefs or knowledge. Friedman's objection to teleological views of epistemic normativity fails.

Conclusion

With this, I conclude my case against the objections to the teleological conception of epistemic normativity. The first important takeaway from the previous discussion is that the notion of value as it appears in teleological theories should be understood not as a substantive, weight-bearing notion, but as a placeholder term that is only analogous with the notion of value in consequentialist ethics insofar as it occupies a similar role in the structure of epistemic normativity as that of value in the structure of ethical normativity. In contrast to what is going on for the ethical consequentialist, however, the epistemic teleologist does not and should not appeal to the *value* or *goodness* of true belief in order to explain the contents of our epistemic norms. The second significant takeaway is that the main objections against epistemic teleology fail because they illicitly import into the epistemic realm assumptions about how value behaves in ethics. Although the teleological conception is ultimately unscathed by the challenges that Berker and Friedman raise, the points I have made in this paper should also give pause to the teleologist who wants to frame her view too closely along the lines of ethical consequentialism, lest she find herself vulnerable to these objections. Justification, she should think, is a guide to truth, but not a guide to the good.

References for Chapter 1

Ahlstrom-Vij, Kristoffer (2013). "In Defense of Veritistic Value Monism". In: *Pacific Philosophical Quarterly* 94.1, pp. 19–40.

Ahlstrom-Vij, Kristoffer and Jeffrey Dunn (2014). "A Defence of Epistemic Consequentialism". In: *The Philosophical Quarterly* 64.257, pp. 541–551.

— eds. (2018a). *Epistemic Consequentialism*. Oxford: Oxford University Press. 335 pp.

— (2018b). "Introduction". In: *Epistemic Consequentialism*. Ed. by Kristoffer Ahlstrom-Vij and Jeffrey Dunn. Oxford: Oxford University Press, pp. 1–17.

Alston, William (2006). *Beyond "Justification": Dimensions of Epistemic Evaluation*. Ithaca, New York: Cornell University Press.

Audi, Robert (1988). "Justification, Truth, and Reliability". In: *Philosophy and Phenomenological Research* 49.1, p. 1.

Berker, Selim (2013a). "Epistemic Teleology and the Separateness of Propositions". In: *Philosophical Review* 122.3, pp. 337–393.

— (2013b). "The Rejection of Epistemic Consequentialism". In: *Philosophical Issues* 23.1, pp. 363–387.

BonJour, Laurence (1985). *The Structure of Empirical Knowledge*. Cambridge, Mass.: Harvard University Press. 258 pp.

Craig, Edward (1990). *Knowledge and the State of Nature: An Essay in Conceptual Synthesis*. Oxford: Clarendon Press. 169 pp.

David, Marian (2014). "Truth as the Primary Epistemic Goal". In: *Contemporary Debates in Epistemology*. Ed. by Matthias Steup, John Turri, and Ernest Sosa. 2nd ed. Chichester: Wiley-Blackwell, pp. 363–377.

Firth, Roderick (1981). "Epistemic Merit, Intrinsic and Instrumental". In: *Proceedings and Addresses of the American Philosophical Association* 55.1, p. 5.

Foot, Philippa (1985). "Utilitarianism and the Virtues". In: *Mind* 94.374, pp. 196–209.

Friedman, Jane (2019). "Teleological Epistemology". In: *Philosophical Studies* 176.3, pp. 673–691.

Goldman, Alvin (1979). "What Is Justified Belief?" In: *Justification and Knowledge*. Ed. by George Sotiros Pappas. Dordrecht: Springer Netherlands, pp. 1–23.

— (2002). "The Unity of the Epistemic Virtues". In: *Pathways to Knowledge*. Oxford University Press.

— (2015). "Reliabilism, Veritism, and Epistemic Consequentialism". In: *Episteme* 12.2, pp. 131–143.

Greaves, Hilary and David Wallace (2006). "Justifying Conditionalization: Conditionalization Maximizes Expected Epistemic Utility". In: *Mind* 115.459, pp. 607–632.

Henderson, David K. and John Greco, eds. (2015). *Epistemic Evaluation: Purposeful Epistemology*. First edition. Oxford: Oxford University Press. 293 pp.

Ichikawa, Johnathan Jenkins and Matthias Steup (2018). *The Analysis of Knowledge*. In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta. Summer 2018.

Joyce, James M. (1998). "A Nonpragmatic Vindication of Probabilism". In: *Philosophy of Science* 65.4, pp. 575–603.

Kvanvig, Jonathan L. (2014). "Truth Is Not the Primary Epistemic Goal". In: *Contemporary Debates in Epistemology*. Ed. by Matthias Steup, John Turri, and Ernest Sosa. 2nd ed. Chichester: Wiley-Blackwell, pp. 351–362.

Littlejohn, Clayton (2012). *Justification and the Truth-Connection*. Cambridge: Cambridge University Press. 269 pp.

Percival, Philip (2002). "Epistemic Consequentialism". In: *Aristotelian Society Supplementary Volume* 76.1, pp. 121–151.

Perrine, Timothy (2021). "Conceptions of Epistemic Value". In: *Episteme*, pp. 1–19.

Pettigrew, Richard (2016). *Accuracy and the Laws of Credence*. Oxford: Oxford University Press. 238 pp.

Pritchard, Duncan (2014). "Truth as the Fundamental Epistemic Good". In: *The Ethics of Belief*. Ed. by Jonathan Matheson and Rico Vitz. Oxford: Oxford University Press, pp. 112–129.

Schafer, Karl (2014). "Knowledge and Two Forms of Non-Accidental Truth". In: *Philosophy and Phenomenological Research* 89.2, pp. 373–393.

Scheffler, Samuel (1994). *The Rejection of Consequentialism: A Philosophical Investigation of the Considerations Underlying Rival Moral Conceptions*. Rev. ed. Oxford: Clarendon Press. 196 pp.

Schoenfield, Miriam (2015). "Bridging Rationality and Accuracy". In: *Journal of Philosophy* 112.12, pp. 633–657.

— (2017). "Conditionalization Does Not (in General) Maximize Expected Accuracy". In: *Mind* 126.504, pp. 1155–1187.

— (2022). "XII—Deferring to Doubt". In: *Proceedings of the Aristotelian Society* 122.3, pp. 269–290.

Setiya, Kieran (2018). "Must Consequentialists Kill?" In: *The Journal of Philosophy* 115.2, pp. 92–105.

Sharadin, Nathaniel (2018). "Epistemic Consequentialism: Haters Gonna Hate". In: *Metaepistemology: Realism & Antirealism*. Ed. by Christos Kyriacou and Robin McKenna. New York: Palgrave Macmillan, pp. 121–143.

Singer, Daniel J (2018). “How to Be an Epistemic Consequentialist”. In: *The Philosophical Quarterly* 68.272, pp. 580–602.

— (2019). “Permissible Epistemic Trade-Offs”. In: *Australasian Journal of Philosophy* 97.2, pp. 281–293.

Sosa, Ernest (2007). *A Virtue Epistemology. Apt Belief and Reflective Knowledge*. Oxford: Clarendon Press. 149 pp.

Sylvan, Kurt (2020a). “An Epistemic Nonconsequentialism”. In: *The Philosophical Review* 129.1, pp. 1–51.

— (2020b). “Reliabilism without Epistemic Consequentialism”. In: *Philosophy and Phenomenological Research* 100.3, pp. 525–555.

Wittgenstein, Ludwig (2009). *Philosophical Investigations*. Trans. by G. E. M. Anscombe, P. M. S. Hacker, and Joachim Schulte. Rev. 4th ed. Chichester: Wiley-Blackwell. 321 pp.

Why be antisocial (about epistemic normativity)

Picture this—you are having friends over for dinner, and are preparing a sumptuous beef brisket. When it's time to serve your guests, you realize that you had falsely believed that none of them was vegetarian; however, one is. What an embarrassment! Had you held an accurate belief, this discomfiting situation would have been avoided. This example calls attention to an important feature of true beliefs: as a rule, they help one navigate the world successfully, accomplishing what one intends, whereas false beliefs do just the opposite. It has seemed very plausible to many that this feature of true beliefs is intimately connected with epistemic normativity; that accurate beliefs have such instrumental value explains why we ought to believe what we ought to believe: epistemic norms are meant to help us believe the truth, and we should conform to them because believing the truth is instrumentally valuable.

Yet, this view sits ill with the fact that sometimes, believing the truth is not instrumentally valuable for individuals, as when a friend spoils a movie ending for us; yet, we judge beliefs to be justified or not regardless of whether the beliefs in question are useful to those who have them (Kelly 2003). How should we resolve this tension? Trying to preserve the insight that epistemic normativity is inherently tied to the instrumental value of truth while avoiding the unwanted result that it is unduly dependent on what's useful for individuals, a flurry of epistemologists have argued in recent years that epistemic normativity is a fundamentally social affair, one that is grounded in the value truth has *for communities*.

I will argue that epistemic normativity is not essentially social in this way. If what we ought to believe were to be explained in terms of the usefulness of truth for communities, I show, we should expect our epistemic norms to be wildly different from what they are: they could condone significant trade-offs between true and false beliefs across a community when doing so would be useful for it; in addition, they should sometimes fail to prescribe believing over other doxastic attitudes such as accepting—the attitude of regarding a proposition as true for theoretical or practical purposes—when the epistemic aims of the community would be equally well-served by these alternative attitudes, even in cases where it would intuitively be wrong to not believe. Finally, none of the accounts offers a satisfying way to deal with the problem of what isolated individuals should believe. More broadly, these arguments also cast doubt on the viability of any purely social account of epistemic normativity.

2.1 A puzzle about epistemic normativity

2.1.1 Epistemic normativity is grounded in the instrumental value of truth

The dinner situation described above was meant to call attention to an important feature of beliefs: as a rule, true beliefs help us navigate the world successfully, whereas false beliefs do just the opposite. Even a cursory reflection on the matter should suffice to shore up this idea. Having true beliefs helps with nearly everything one does: it is undeniably useful to know where your phone is when you need to make a call, to know what ingredients you already have at home when shopping for groceries, or to know how to get to your destination without getting lost. And what holds in these mundane situations holds, *a fortiori*, in high-stakes cases: in a medical emergency, knowing precisely what symptoms a patient has is crucial for a doctor to make an accurate diagnosis; in financial planning, knowing your current assets and liabilities allows you to make sound investment decisions. In Frank Ramsey's famous words, "beliefs are a map by which we steer.[1]" We had better hope that this map represents the world accurately, so that we can get where we want to be! We can encapsulate this idea by saying that believing the truth and not believing the false is instrumentally valuable for us: it helps us achieve our aims and goals.[2]

From here, it is only a small step to recognizing the instrumental value of conforming to epistemic norms. How we form and revise our beliefs is subject to norms. These norms specify when forming or revising beliefs is permissible, obligatory, or forbidden. They prescribe responding to evidence in specific ways: for instance, if you encounter decisive evidence that one of your beliefs is false, then you ought to revise it. Similarly, if you encounter someone whom you recognize as an expert on a particular subject, you should defer to their opinion regarding claims in that domain. These norms may also forbid certain combinations of doxastic attitudes, such as believing that p while simultaneously believing that you were likely unreliable when you formed the belief that p. Such norms appear to be directed at getting things right and avoiding error. Call them 'epistemic norms.' If epistemic norms help us get to the truth and avoid the false, and believing the truth and avoiding the false is instrumentally valuable for us, then it seems that conforming to epistemic norms will be instrumentally valuable as well. Furthermore, even if having true beliefs does not seem to matter much now, it may matter greatly in the future. Regularly conforming to epistemic norms now bolsters dispositions that enable one to conform when the stakes are high; conversely, frequently violating these norms could make one less able to conform when it truly matters.

1. The original quote is "[a] belief of the primary sort is a map of neighbouring space by which we steer." (Ramsey 1990, p. 146)
2. In order to avoid burdening the reader, I will use 'believing accurately' and 'believing the truth' as shorthands for 'believing the truth and not believing the false.' If the distinction between believing the truth and not believing the false is relevant in a particular passage, it will be explicitly stated.

It would be surprising if the instrumental value of truth—and, by extension, the value of conforming to norms that promote true beliefs—were merely coincidental and of no deeper significance. This has led many to think that this connection captures a fundamental feature of epistemic normativity: that it is grounded in the instrumental value of truth.[3] That is, the prescriptions and reason-giving nature of epistemic norms derive from the fact that conforming to them helps agents achieve their ends by promoting true beliefs and avoiding false ones.

This view is known as *instrumentalism* about epistemic normativity, and has been endorsed by many philosophers over the past few decades. In addition to explaining why conforming to epistemic norms is generally instrumentally valuable, instrumentalism possesses several theoretical strengths. When trying to understand epistemic norms, there are several questions we may ask. First, we may inquire about *what the correct set of epistemic norms is*. Is a purported norm *N* a genuine epistemic norm? This is what we are asking, for example, when we ask whether we should remain steadfast or update our beliefs in the face of peer disagreement. A related but distinct line of inquiry concerns the *metaphysical status of epistemic norms*: are epistemic norms *sui generis*, or are they reducible to other kinds of norms, such as moral norms or instrumental norms? Finally, we may ask *what explains their normative authority*: why are epistemic norms reason-giving, unlike norms of etiquette or the bylaws of crime syndicates? Instrumentalism offers plausible answers to all three questions.[4]

First, to the question, 'what is the correct set of epistemic norms?', instrumentalism answers: those norms for belief that, when followed, are conducive to the promotion of one's ends. Because believing accurately is what helps one further one's ends, these are the norms are conducive to truth. Second, instrumentalism provides an answer to the question of the metaphysical status of epistemic norms: epistemic norms, on this account, are akin to 'recipes' for promoting our ends; they are reducible to instrumental norms.[5] This answer will be especially appealing to philosophers of a broadly naturalistic bent, who are generally doubtful of the existence of irreducibly normative properties. Instrumental—or *means-ends*—normativity does not entail the existence of such properties: ultimately, what means are conducive to certain ends can be cashed out in purely descriptive terms.

3. See for instance Kornblith (1993).
4. There are about as many ways to frame instrumentalism as there are discussions of it. A particular choice-point is whether to distinguish two versions of instrumentalism, one as a thesis about the content of epistemic norms and the other as a thesis about what explains the normative force of these norms. See Côté-Bouchard (2021). Many discussions, such as Kelly's foundational "Epistemic Rationality as Instrumental Rationality: A Critique", disregard this distinction. I choose to present instrumentalism as a view that answers both the question of what epistemic norms prescribe and the question of why epistemic norms have normative authority. I do not think that the objections I present succeed or fail depending on the particular way it is framed.
5. I take the 'recipe' gloss from Graham (2015).

Last and not least, instrumentalism provides an explanation for the normative force of epistemic norms. When we ask 'what ought I believe?', it seems that the prescriptions of epistemic norms are always relevant—or, some will think, completely settle the question (Nolfi 2021).[6] In other words, epistemic reasons always bear on what we ought to believe. This would not be the case if these norms were not genuinely binding. Moreover, a look at our practices of assessing one another's beliefs suggests that we take one another to be bound by epistemic norms: we criticize beliefs we deem irrational, try to inculcate good belief-formation habits in others, and praise the display of epistemic virtues. We take it that people really believe as they ought not to when they exhibit resistance to evidence, such as is the case with climate change denialism.

Yet, not all sets of norms that prescribe actions or beliefs come with reasons to act or believe in certain ways. Consider, for instance, the bylaws that govern some crime syndicates. These rules may require that members obey their superiors' orders no matter what, even if doing so entails engaging in various harmful activities such as extortion and murder. In such cases, despite the fact that there is a rule that prescribes that members obey their superior's orders, it is not the case that a member has a reason to do as they require.[7] Another way of getting at the claim: even though the rule may be expressed in terms of what members ought to do—members ought to obey their superior's orders—this 'ought' is not genuinely normative; it is not the case that members ought to obey their superior's orders. Contrast with moral norms, such as 'you ought not to steal', 'you ought not kill', etc. Whatever these norms say you ought to do, you genuinely ought to do. Epistemic norms, just like moral norms, are generally taken to express genuine prescriptions.

What makes epistemic norms different from the rules of crime syndicates? What explains their reason-giving force? One explanation appeals to constitutive features of beliefs: belief is an attitude that is constitutively governed by a norm of truth. From this fundamental epistemic norm, we can derive secondary epistemic norms, such as "one ought to believe in accordance with their evidence." Because we have beliefs, we are inescapably subject to these norms.[8] Epistemic norms could also reduce to norms of another domain: for instance, it could be that epistemic norms reduce to moral norms. This is one way of understanding the view held by W. K. Clifford, who, in "The Ethics of Belief", famously wrote that "it is [morally] wrong always, everywhere and for anyone, to believe anything upon insufficient evidence." Clifford might be read as saying

6. I leave open the question of whether they are the only relevant considerations that bear on what one ought to believe.
7. Under the stipulation that no harm would come to the member for disobeying. If that weren't the case, the member would have a prudential reason to do what their superior orders them to.
8. For the idea that it is constitutive of belief that it is correct if and only if true, see Shah (2003), Shah and Velleman (2005), and Velleman (2000). Wedgwood formulates and endorses a version of epistemic constitutivism in Wedgwood (2002). For a critical discussion, see Côté-Bouchard (2016)

that the norm forbidding one to believe that p on the basis of insufficient evidence—an epistemic norm—reduces to a moral norm.

Here is, in contrast, a view that does not by itself explain why epistemic norms issue binding prescriptions: the view that epistemic norms just are effective means of believing the truth.[9] Unless supplemented by an account of why we should believe the truth, the view does not answer the question of why epistemic norms are binding. Compare: there are effective ways to count blades of grass (call them "blades-of-grass counting rules"), but that fact does not entail that I ought to count blades of grass in this way. Unless I have a reason to count blades of grass—and to do so well—, I have no reason to abide by blades-of-grass counting rules. In other words, the mere existence of effective means of bringing about a certain result does not entail that I ought to take these means.

Instrumentalism provides a straightforward explanation of why epistemic norms have normative force. Often, what we have reason to do is dependent upon having certain ends: suppose I want to make myself a coffee; this gives me a reason to turn on the coffee machine, grind coffee beans, etc. Had I not wanted to make a coffee, *ceteris paribus*, I would not have had any reason to do any of these actions. Epistemic instrumentalists hold that the same is true regarding what I epistemically ought to believe. It is because what I epistemically ought to believe contributes to my achieving my ends that the epistemic ought has normative force for me.

2.1.2 Truth is not always instrumentally valuable

So far, I have argued that there is a strong case for the idea that epistemic normativity is grounded in the instrumental value of truth. Here is another plausible claim to consider: having accurate beliefs is not always instrumentally valuable, as the next case illustrates.

> **SPOILER ALERT**
>
> You just discovered a new crime mystery show and are riveted. Because you want to enjoy it to its fullest extent, you want to avoid learning who is the culprit until watching the episode in which the identity of the murderer is revealed. Thinking that you have finished watching the show, your friend Josh, who you know has watched it, asks you: "So, were you surprised to learn that Colonel Mustard was the murderer?"[10]

In this scenario, trusting Josh's testimony and forming the belief that Colonel Mustard is the murderer would actually diminish your enjoyment, since it would spoil the suspense. Here, the true belief works against your goal and is therefore not instrumentally valuable. Other examples of this kind abound: a parent might be deeply distressed to learn that their child has been stealing from family members, or a marriage might never recover if a long-buried affair comes to light. In such cases, true beliefs cause significant harm. Conversely, false beliefs can sometimes be useful.

9. See for instance Quine (1998).
10. The *locus classicus* of this case is Kelly (2003).

A false belief that one's spouse has always been faithful might contribute to a stable and happy marriage, or a false belief that one's child is morally upright might significantly enhance a parent's happiness and flourishing. Hence, accurate beliefs are not always instrumentally valuable.

The idea that truth is not always instrumentally valuable does not, by itself, conflict with the view that epistemic normativity is grounded in the instrumental value of truth. The tension only arises when we introduce a third claim: that epistemic normativity is universal.

2.1.3 Epistemic norms apply universally

Consider SPOILER ALERT again. This example not only demonstrates that truth is not always instrumentally valuable, but it also suggests something important about epistemic normativity: it is *universal*. Any belief is a proper target of epistemic evaluation and ought to be formed and maintained in accordance with epistemic norms. This includes beliefs that, if true, would not be instrumentally valuable to the believer. In SPOILER ALERT, believing that Colonel Mustard is the murderer would spoil the show and go against your desire to avoid learning the identity of the murderer until the big reveal. Yet it seems that you are still epistemically obligated to form this belief, given that you have strong evidence supporting its truth.

Additionally, consider the other examples mentioned earlier. If a parent catches their child red-handed, hiding some of their relatives' belongings in their coat pockets, and has been informed by several family members that this has happened before, they epistemically ought to believe that the child is stealing, even if doing so would be detrimental to their well-being. The same holds true for a spouse who finds undeniable proof that their husband had an affair: the spouse epistemically ought to believe this, whether or not this truth would harm them. In these cases, where an agent's interests lie doesn't seem to affect what the epistemic norms prescribe.

Thus, epistemic norms such as "one ought to believe what their evidence supports" do not appear to be contingent on an agent's goals or desires. They apply universally, regardless of whether following them aligns with the agent's interests. With this final claim, we now are in possession of the three claims that make up our puzzle.

2.1.4 The puzzle

So far, we have considered reasons to think that these three claims are individually plausible.

1. **Epistemic normativity is grounded in the instrumental value of truth**—the prescriptions and reason-giving nature of epistemic norms derive from the fact that conforming to them helps agents achieve their ends by promoting true beliefs and avoiding false ones.
2. **Believing the truth and avoiding falsehoods is not always instrumentally valuable.**
3. **Epistemic norms apply universally**—their prescriptions and reason-giving force do not depend on whether believing in accordance with them furthers an individual's goals or interests in a particular context.

These claims are in tension with each other. If **1** and **2** are true, then we should expect epistemic norms to not prescribe beliefs that go against the believer's ends in cases where believing the truth and avoiding falsehoods would not be instrumentally valuable, making **3** false. If **2** and **3** are true, then it seems we must reject the claim that epistemic normativity is grounded in the instrumental value of truth. Finally, **1** and **3** might be consistent; but the universality of epistemic norms would only be consistent with the claim that epistemic normativity is grounded in the instrumental value of truth if believing the truth were always instrumentally valuable, making **2** false.

But what if there was a way to circumvent the problem altogether, one which did justice to the idea that epistemic normativity is grounded in the value of truth, while accounting for its universality, without being committed to the dubious claim that believing the truth is always instrumentally valuable? This is precisely what *social instrumentalism* promises.

2.2 A social approach to epistemic normativity

Why think of epistemic norms in social terms in the first place? One good place to start is to ask ourselves what we are doing to when we are assessing one another's beliefs as rational or irrational, warranted or not, and so on. Several philosophers have argued that our epistemic evaluations, as well as our attributions of knowledge, fulfill certain social needs: given our cognitive limitations, we need to be able to rely on one another's testimony when forming our beliefs; epistemic evaluations and attributions of knowledge serve to enforce conformity with some set of rules governing the formation and revision of beliefs, as well as to flag potential informers as reliable or unreliable sources of information.[11] Epistemic evaluation plausibly has a social component.

Social instrumentalists, as the name suggests, take this social component and elevate it to a central feature of epistemic normativity. According to them, epistemic normativity is grounded in the fact that conformity to epistemic norms promotes the ends, not of individuals, but of *communities*. By making epistemic normativity a social phenomenon, their view promises to preserve all the benefits of traditional instrumentalism, while avoiding the unwanted result that epistemic normativity is unduly contingent on each and everyone's idiosyncratic goals and desires.

For instance, in her 2021 paper "Could our Epistemic Reasons be Collective Practical Reasons?", Michelle Dyke presents a view on which "epistemic reasons are instrumental in nature..."; however, the ends they serve are not those of those individuals, as traditional instrumentalism posits, but "are attributable to collectives" instead (Dyke 2021, p. 847). These collectives include juries, boards, labs, as well as larger groups such as society as a whole. They have cognitive goals:

11. Craig (1990) is often regarded as the originator of this approach to knowledge and epistemic evaluation. See also Dogramaci (2012) and the collection of essays in Henderson and Greco (2015).

for instance, juries have to determine whether a defendant is, and a lab may aim to uncover a drug's side effects.

In effect in these groups are also sets of epistemic standards—"collections of rules or expectations" that are meant to regulate belief-formation and revision of members—that contribute to achieving its cognitive aims (Dyke 2021, p. 847). They determine, among other things, what counts as evidence, how to revise one's belief on the basis of evidence, and how peer disagreement is to be resolved. Membership in a group generates reasons for one to believe according to these shared standards.

Presenting the view more schematically,

> An agent A has epistemic reason$_c$ to believe *p* just in case:
> (i) A is a member of community C,
> (ii) Community C has some epistemic ends E,
> (iii) C's achievement of E is most effectively promoted by C's adherence to epistemic standards S, and
> (iv) A's believing *p* is called for by standards S.

In their forthcoming "The Construction of Epistemic Normativity," Hannon and Woodard take a similar approach. Drawing on the work of Craig and Dogramaci, among others, they argue that epistemic norms address coordination problems within communities. We need true beliefs as well as to avoid error in order to successfully achieve what we want. However, we are limited agents: it would not only be costly, but beyond our capacity, to rely solely on ourselves to inquire about everything we have an interest in knowing. We thus have to rely on the testimony of others when we form our beliefs. We have to make sure that testimony in the community is reliable and trustworthy. But there are many instances in which oneself and others may pollute, rather than improve, the shared pool of information that circulates in the community: that is the case when we engage in self-deception, wishful thinking, or retain beliefs in the face of very strong evidence that they are false. How, then, can members of the community ensure the general reliability and trustworthiness of testimony and avoid basing their beliefs on unreliable testifiers?

This is where epistemic norms come in, argue Hannon and Woodard: in order to ensure that testimony is indeed reliable and trustworthy, we need to subscribe to an appropriate system of epistemic norms. By assessing one another's beliefs as 'rational' or 'irrational', by praising good doxastic behavior, and by sanctioning non-conformity, we are pressuring members of the community to form their beliefs according to truth-conducive rules.[12] We are thus ensuring that others "play by the same set of epistemic rules," which makes it reasonable for us to trust the testimony of others: if I am aware that others use the same rules as me to form their beliefs, I can

12. Sanctions might include, among other things, reducing trust in norm-violators as well as expressing blame and disapproval.

trust that they arrived at the same beliefs I would have given their evidence, and I can reasonably adopt these beliefs myself. And, as a member of a group, if I violate them, I am criticizable for it; I am thus bound by them (Hannon and Woodard forth. P. 16).

Hannon and Woodard's view, schematically presented:

An agent A epistemically ought[13] to believe *p* just in case:
- (i) A is a member of community C,
- (ii) Members of C enforce a set of epistemic standards S in order to promote the reliability of testimony within their community, and
- (iii) A's believing *p* is called for by standards S.

Social instrumentalism is an appealing position. First, it accounts for the fact that conforming to epistemic norms is useful and, in doing so, retains all of the benefits of instrumentalism. Yet, in contrast to traditional instrumentalism, it locates its usefulness at the level of communities. Because the normative force of epistemic norms is explained by the usefulness of these norms for communities, rather than for individuals, the view does not entail that whether we ought to conform to epistemic norms is unacceptably contingent upon one's individual goals and desires: it might not be useful for an agent to conform to *E*, but because that agent is a member of a community that has *E* as one of their epistemic standards, the agent is bound by the norm.

2.3 Social instrumentalism: a critique

In what follows, I argue that social instrumentalism, and, in fact, any view that explains the normative force of epistemic norms by appeal to the roles that epistemic norms play in communities. The first argument, which draws on Selim Berker's critique of epistemic consequentialism, shows that social versions of instrumentalism condone impermissible trade-offs (§3.1). The second argument is that, because acceptance and belief are often indistinguishable, the social instrumentalist often cannot say that agents who merely accept rather than believe a proposition when faced with sufficient, undefeated evidence of its truth are doing anything wrong (§3.2). The third argument appeals to our intuitions about the reasons to believe that agents would have even if they were not members of any collective. I argue that the social instrumentalist is drawing the wrong conclusion when she claims that epistemic reasons are intersubjective—they are in fact universal (§3.3).

2.3.1 Impermissible Trade-offs

The first issue faced by social instrumentalism is that it condones trade-offs that are clearly impermissible. This is because social instrumentalism views what one ought to believe in terms of

13. Given that Hannon and Woodard do not consider 'ought to' and 'has reason to' to be interchangeable, I frame their view using 'ought', their preferred term.

effective means for bringing about a valuable state of affairs. Sometimes, this may entail trading off something of value for the sake of something that has greater value; in our case, requiring that members abide by standards that are problematic from an epistemic point of view in order to achieve the epistemic ends of the group.

To start, consider the following case:

> **Climate Change Denialism**
>
> A think tank aims to defend the claim that human activity is not a significant factor in climate change. In order to do so, its researchers would have to reject conclusions widely accepted among the scientific community and dogmatically resist any evidence speaking against the claim they wish to defend, regardless of its strength.

Here, the think tank has the end of defending the claim that anthropogenic factors are not a significant cause of climate change, a view that goes against a well-established scientific consensus. Defending that claim would require clearly problematic epistemic standards to be in place. If social instrumentalism were true, then researchers for the think tank epistemically ought to abide by these inadequate epistemic standards; after all, following these standards is the best way to achieve the epistemic end of the collective. But clearly, epistemic norms do not require anyone to dogmatically resist strong evidence speaking against a claim they (or a group of which they are a member) want to defend. In fact, resistance to evidence seems like a paradigmatic epistemic vice. We should thus reject social instrumentalism.

I see two possible responses to this objection. The first is to deny that the end in question is an *epistemic* end. The social instrumentalist could argue that epistemic ends are those that, if achieved, would leave one (or a group) in a better epistemic standing with respect to a question than the one in which one started—or, at least, could be reasonably expected to leave one in a better epistemic standing. There are different ways to spell out this idea. An improved epistemic standing could involve knowing more about the question than at the beginning of inquiry, having more accurate credences, or having more true beliefs. Because the claim the think tank wants to defend is false, and the think tank should know that it is, it cannot reasonably expect to be in a better epistemic standing at the end of inquiry. Thus, the end is not epistemic, and its members are not bound by the epistemic standards that would be most conducive to defending the claim.

If this reply is to succeed, the ball is in the social instrumentalist's court: she owes us an account of what makes an end an epistemic one. Until we have a convincing one, we have no reason to reconsider this argument. As I'll argue below, claiming that the end in Climate Change Denialism is not epistemic only grants the social instrumentalist a temporary reprieve; what generates the problem is something altogether different.

The second response is to deny that the end in question generates reasons to take the means necessary to achieve it. In contrast to the first reply, the social instrumentalist would not be denying that the end in question is epistemic. Instead, she could appeal to a plausible principle about the transmission of instrumental reasons: that ends only generate reasons to take the best means to achieve them when it is permissible to pursue them in the first place. Sometimes, having an end is itself enough for generating reasons. For instance, when I want to make myself a coffee, I thereby have reasons to grind coffee, put the kettle on, etc. However, in other cases, our end may be one we ought not to pursue: e.g., torturing kittens. In this case, it is plausible that, because torturing kittens is impermissible, having that end does not generate reasons to take the means to achieve it. Arguably, the end of the think tank is harmful and should not be pursued. If we deny that the think tank ought to pursue it, then we are not committed to the result that members ought to abide by the problematic epistemic standards that doing so would require. Social instrumentalism is thus safe from this counterexample.

I find both the transmission principle presented above and the claim that the think tank ought not to pursue the end of defending the view that human activity is not a significant factor in climate change plausible, and am happy to grant them for the sake of the argument. However, neither helps resolve the underlying problem: it arises even in cases when the ends are good, as the next case shows.

> **Pharmaceutical Company**
>
> A pharmaceutical company aims to develop a revolutionary drug that would cure a certain autoimmune disorder. A successful drug of this kind would constitute a breakthrough that has long eluded the pharmaceutical industry. The company is much more likely to succeed if researchers are confident that they can and will succeed. However, the researchers at the company do not have any evidence that they will; in fact, they have strong evidence that they won't, given that many companies employing much more renowned scientists and benefiting from much better funding have failed before.

To significantly increase their chances of success, it would be best for the scientists to have high confidence that they will succeed despite having significant evidence that they won't. If social instrumentalism is true, they epistemically ought to exhibit resistance to evidence about their chances of success—if not ignoring this evidence altogether, then at least downplaying it. Once again, this cannot be the case: resistance to evidence is an epistemic defect, not something that epistemic norms prescribe.

The rejoinder to Climate Change Denialism, that bad epistemic ends do not give rise to reasons to take the means to achieve them, is unavailable here, since the end in this case is laudable. Instead, the social instrumentalist could bite the bullet but try to mitigate its problematic appearance by insisting on the relativity of what one ought to believe. It's open to her to claim that the researchers ought to believe that they will succeed relative to the standards of the pharmaceutical company, but not relative to the standards of everyday reasoning. Standards of everyday reasoning serve a different end, namely, the end of promoting the reliability of testimony within a community. Given that they serve different ends, the standards of the pharmaceutical company and the standards of everyday reasoning may yield conflicting prescriptions, as in the case above. When we judge that social instrumentalism yields the wrong verdict in Pharmaceutical Company, we are doing so using the standards of everyday reasoning. However we resolve the tension (denying that there's something that the researchers ought, all things considered, to believe; resolving in favor of the standards of everyday reasoning; or resolving in favor of the standards of the pharmaceutical company), there is a sense in which it is true that the scientists ought to believe that they will succeed.

Regardless of whether this reply succeeds, social instrumentalism would still be in the hot seat, because the problem arises again in the case of everyday reasoning. Social instrumentalists understand the standards in effect in communities at large (the standards of everyday reasoning) as aiming at ensuring the reliability of testimony within a community. This is briefly mentioned as the epistemic end of societies as a whole by Dyke and is central to Hannon and Woodard's view. But even at this level of generality, social instrumentalism countenances impermissible trade-offs, as the following example illustrates.

> **Cosmological Beliefs**
>
> A community holds a set of cosmological beliefs that were not formed and are not maintained according to reliably truth-conducive norms. These cosmological beliefs are the glue that binds the community together. Were they to be revised according to truth-conducive norms, they would be abandoned. The sense of meaninglessness generated by that loss would lead to widespread disarray and, potentially, to a partial disbandment of the community. This, in turn, would lead to a breakdown in the overall reliability of testimony in what would remain of the community, because many of its most reliable testifiers would have left upon abandoning the cosmological beliefs in question.

If social instrumentalism were true, then members of the community ought to hold the cosmological beliefs that play a central role in promoting the reliability of testimony within the community, despite not having been formed and not being maintained in accordance with truth-conducive norms themselves. Again, this is the wrong verdict: these beliefs are not epistemically justified. Thus, social instrumentalism is false.

The objection above recalls Selim Berker's objections against epistemic consequentialism. It involves cases in which an effective—or the most effective—way to bring about the achievement of a valuable or valued end involves trade-offs (Berker 2013a,b). Trade-offs arise when doing a certain action would incur some disvalue but best promote overall value. As the classic case goes: suppose that a doctor has a choice between killing one perfectly healthy patient to collect their organs and save five critically injured ones, and letting the healthy patient live and the five injured die. On a simple version of act-utilitarianism, on which what one ought to do is what maximizes expected utility, it appears that what the doctor ought to do is to kill the one to save the five.

If we ought to take the means to a given end, and these means involve trade-offs, then we ought to make these trade-offs. However, as Berker has argued, our views about what we ought to believe do not countenance trade-offs. To see this, consider the following case:

> **God Grant**
>
> You are an atheist seeking a grant from a religious foundation. This grant would allow you to pursue research and come to a great number of true beliefs about important matters of fact. However, you know that the foundation only funds researchers who believe that God exists. You have pondered over this question in the past, weighed evidence for and against believing that God exists, and have come to the conclusion that there is no God. Should you believe that God exists?

According to a crude version of epistemic consequentialism on which what one ought to believe simply is what best promotes the overall epistemic good, claims Berker, what you ought to believe is that God exists (and thus, such a belief would be justified).[14] This is because the overall epistemic good is best promoted by forming the belief that God exists. But clearly, your belief would not be justified: it would be formed against what you take to be strong evidence that God doesn't exist. Thus, epistemic teleology ought to be rejected.

What cases like *God Grant* show is that when we assess whether a belief is justified, we do not take into account its downstream consequences, such as whether it leads to forming many other true beliefs. Judging that a belief is justified on account of its consequences, no matter how epistemically valuable, violates what Berker labels 'the separateness of propositions,' the epistemic analog to the Rawlsian 'separateness of persons': just like we ought not to make trade-offs across persons, we ought not make trade-offs across beliefs. Our epistemic evaluations are *backwards-looking*, not *forward-looking*. Any plausible account of epistemic norms should respect this feature. Social instrumentalism does not. Furthermore, the problem does not arise because of a merely superficial feature of instrumentalism; it is generated by its central claim, the claim that what one ought to believe is determined by what brings about certain ends. In order to bring about these ends, it may be useful to violate what we would regard as epistemic norms and form unjustified beliefs. Yet, the social instrumentalist has to say that, in these cases, we epistemically ought to have these beliefs. If this diagnosis is right, then there is little hope for social instrumentalism: it is bound to run afoul of the separateness of propositions.

The social instrumentalist could claim that, if the authority of epistemic norms were made void on every occasion where not conforming to them was more useful to the community, the

14. For present purposes, we can think of each true belief as one 'unit' of epistemic value.

benefits to be gained from the practice of enforcing conformity to epistemic norms would be lost: the practice is beneficial only if we know we can rely on one another without having to verify, for each and every belief, whether the agent respected the same standards that I would or whether it was beneficial to them not to believe according to these standards, something that would often be very costly. But the case above is not a case in which individual members of the community are deliberately violating epistemic norms when it is more useful for them to do so; it is one in which the *actual epistemic standards* in force in the community prescribe beliefs that are unjustified. The members of the community are abiding by the standards in effect in the community. And these standards are effective at fulfilling their function: they *do* promote the reliability of testimony within the community. Therefore, the social instrumentalist is committed to the verdict that members of the community epistemically should believe in the way they prescribe.

2.3.2 Accepting and believing

The previous considerations showed that social instrumentalism, just like its individualist analogue, countenances unacceptable trade-offs in which some degree of reliability is sacrificed in order to bring about a greater degree of reliability in the community. This will sometimes entail that individuals ought to believe against their evidence. But there's another way in which social instrumentalism condones failing to believe what the evidence supports: when achieving the ends of the community will be served *equally well* by holding a belief as by holding other doxastic attitudes that, in context, are indistinguishable from belief in terms of how an agent is disposed to behave. Social instrumentalism entails that as long as one or the other of the doxastic attitudes is adopted—whether it is belief or not—the agent is not doing anything impermissible.

Here is the general argument:

1. If one has decisive evidence that p and considers whether p, one ought to believe p;
2. If social instrumentalism is true, it is not the case that if one has decisive evidence that p and considers whether p, one ought to believe p;
3. Therefore, social instrumentalism is false.

I will devote most of this section to defending the second premise. But first, a brief note on premise one. Although the existence of positive epistemic duties to believe has been a matter of some controversy, many deem it plausible that we sometimes have an epistemic duty to form the belief that p when there is sufficient and undefeated evidence supporting p that we can easily take up or of which we are aware.[15] If one rejects the existence of such duties, then one has further reason to reject social instrumentalism, as the social instrumentalist is committed to it: sometimes, as in the cases presented in previous sections, the most effective way to achieve the

15. For arguments in favor of the existence of positive epistemic duties, see Ichikawa (2022) and Simion (2024). For arguments against, see Nelson (2010).

ends of an epistemic community requires that agents have certain beliefs. These beliefs are not merely permissible; they are obligatory.

As I'll argue below, sometimes *accepting* that *p* can be just as effective for bringing about the epistemic ends of communities as believing *p*. This creates a problem for the social instrumentalist: in cases in which accepting that *p* is as effective a means to bringing about the realization of the epistemic ends of a community as believing that *p*, if it is permissible to believe that *p*, then it should also be permissible to accept *p* and vice-versa. But accepting *p* does not entail believing that *p*. Therefore, there will be cases in which one fails to believe that *p* but, because one accepts *p* and accepting *p* is just as effective a means as believing *p*, one is epistemically blameless according to the social instrumentalist. This is the wrong verdict in cases in which the agent possesses sufficient and undefeated evidence that *p*: the agent should not simply accept *p*. By not believing *p*, they are failing from the epistemic point of view: they are violating an epistemic duty to believe.

Acceptance is the state one is in in virtue of having accepted that *p*. One accepts that *p* when one decides to regard *p* as true for the purposes of their theoretical and practical deliberation in a given context. As L. Jonathan Cohen puts it:

> ... to accept that *p* is to treat it as a given that *p*. More precisely, to accept that *p* is to have or adopt a policy of deeming, positing, or postulating that *p*—i.e., of including that proposition or rule among one's premises for deciding what to do or think in a particular context, whether or not one feels it to be true that *p* [i.e., believes that *p*]. (Cohen 1992, p. 4)

In contrast to belief, acceptance is voluntary. Furthermore, although beliefs aim at truth, acceptance need not do so; while in scientific contexts one may accept a hypothesis for the sake of testing it and confirming or disconfirming it, in other contexts one may accept that *p* for considerations that are not truth-related. That is the case, for instance, when someone accepts a proposition because others in their group do so out of a need to coordinate (e.g., a group of lawyers who have to defend a client in court may accept that their client is innocent for that very purpose without anyone *believing* that their client is in fact innocent). Here it is the need for coordination that drives acceptance, not evidential considerations. Lastly, as the case of the lawyers illustrates, acceptance is circumscribed in a way that believing is not: to accept that *p* is to treat *p* as true *for certain purposes, in a certain context*, whereas to believe *p* is to regard *p* as true *tout court*.

The following examples present situations in which agents accept rather than believe a certain proposition despite having decisive and undefeated evidence that the proposition is true. These agents are not believing as they should. Yet, because their acceptance of the proposition translates to the same dispositions to act as a belief would, it is equally well-suited to achieving the epistemic ends of the community.

Creationist Teacher

> Mr. Martin is a biology professor at his local high school. Privately, he is a Young Earth creationist; he believes in a literal interpretation of the creation story from Genesis and thus believes evolutionary theory to be false despite being well aware of the evidence of its truth. However, Mr. Martin keeps these beliefs to himself and teaches evolutionary theory as required by the curriculum. Through his teaching, students come to know that evolutionary theory is true.[16]

Science Skepticism

> Neda is an anxious cognizer; in particular, she is very careful when it comes to accepting science communication: whenever well-recognized, reliable experts assert that anthropogenic climate change is occurring or that vaccines are safe, Neda suspends belief, thinking "Well, scientists sometimes get it wrong! I'll do my own research." But Neda also knows that she would be shunned by her friends if they knew that she does not believe that anthropogenic climate change is occurring or that vaccines are safe. For the sake of fitting in, she accepts these propositions and behaves as if they were true even though they do not reflect her privately held views.[17]

The cases above show that acceptance can suffice for someone to contribute to the reliability of the shared pool of information in a community. Consider Creationist Teacher. In this case, Mr. Martin's beliefs about evolutionary theory, despite being held in the face of strong evidence to the contrary, do not prevent him from teaching the truth when it comes to this very same subject matter. Let's stipulate that there is no other context in which Mr. Martin would share his beliefs. He is an intensely private man, after all, and solitary on top of all that. Yet Mr. Martin is not believing as he should. Yet again, the social instrumentalist cannot give this verdict: Mr. Martin's doxastic attitude is as good as belief when it comes to promoting the reliability of testimony in the community.

Something analogous is going on in Science Skepticism. Neda accepts the scientific claims at issue for the purpose of fitting in, not for epistemic reasons. Her failure to believe in the face of sufficient and undefeated evidence makes her epistemically blameworthy: she is not believing as she should. However, since her acceptance translates to the same dispositions to act as would her belief, it is innocuous; it has no harmful impact on the reliability of the community. Thus, the social instrumentalist is committed to saying that her mere acceptance without belief is epistemically permissible, although it clearly is not.

One may claim that by accepting rather than believing, individuals are forming cognitive dispositions: they are impeding their future ability to respond properly to sufficient and undefeated evidence. This is, in general, bad for the reliability of the community. Therefore, it is impermissible to merely accept without believing what one's evidence supports when one has

16. This case originates in the literature on testimony. See Lackey (1999).
17. This case is adapted from Simion (2024).

sufficient and undefeated evidence that *p* and considers that *p*. But we can imagine cases in which accepting without believing despite strong evidence has no impact on someone's cognitive dispositions. As the next example shows, agents in these cases would still be epistemically blameworthy, a verdict that the social instrumentalist would not be able to give.

> **Mad Scientist**
>
> A mad scientist injects his patients with a drug that will revert them back to their original cognitive dispositions after a certain time. In the meantime, his patients accept rather than believe propositions when faced with sufficient and undefeated evidence of their truth. After having achieved the purpose for the sake of which they were accepting the proposition at issue, they revert back to their ordinarily truth-conducive cognitive dispositions.

In Mad Scientist, the scientist's patients are not believing as they should. However, because their resistance to the evidence during the time they are under the drug's effect will not impair their ability to respond properly to evidence in the future, the social instrumentalist cannot say that the patients are doing anything epistemically blameworthy. It is probable that such a fanciful case would not sway the social instrumentalist. But even if it does not, I think that it brings into focus the same problem that the trade-offs cases brought up: the social instrumentalist is evaluating whether beliefs are epistemically permissible in virtue of their downstream consequences. But whether an agent's belief is epistemically permissible or not is not a matter of the impact it will have on the agent's (or the community's) future beliefs. It is a matter of how the agent's belief fits their evidence and how it was formed.

Finally, the social instrumentalist could appeal to the claim that, in fact, our epistemic norms usually prescribe *beliefs* as a response to strong, undefeated evidence that *p*. If they do, it would be impermissible to merely accept *p*, even when doing so promotes the epistemic ends of communities just as well as believing *p*. Yet, that they do so is merely a contingent matter of fact. We can imagine a community in which acceptance is treated as permissible when belief also is. Given this community's norms, the social instrumentalist would be in no position to claim that they ought to believe that *p* when possessing strong, undefeated evidence that *p*. Once again, this is the wrong verdict.

2.3.3 The 'Desert Island' Objection

Lastly, it is worth considering what the social instrumentalist would say about putative cases in which an agent is not a member of any epistemic community. If epistemic normativity is essentially social, then such an individual does not possess any epistemic reasons for belief; put otherwise, there is nothing that this individual epistemically ought to believe. But this seems absurd: this individual would clearly possess epistemic reasons in favor of and against various doxastic attitudes. If this objection succeeds, then it would preclude a complete reduction of epistemic normativity to social normativity. It would thus be a problem not simply for social instrumentalism, but for any view that makes epistemic normativity a form of social normativity.

To drive home how unpalatable rejecting the idea that an isolated agent has no epistemic reasons for belief is, consider the following scenario:

> **Desert Island**
>
> Jan lives alone on a desert island. The island is a safe haven, free of any substantial dangers, and its fertile soil provides all that is needed for Jan to survive without much work. To pass the time, he likes to toss a coin and compile the results. So far, the coin has landed heads ninety-nine times out of a hundred tosses. Jan stops to wonder whether the coin is fair, concludes that it is, and continues to toss it.

It is obvious that Jan should not believe that the coin is fair—his belief would not be justified. After all, he has excellent evidence that it is not and either disregarded the evidence or drew the wrong conclusion from it. However, because the social instrumentalist is committed to the view that what we epistemically ought to believe is set by the ends of the communities of which we are members, she has to deny that Jan is not believing as he epistemically should. Truth-related considerations, on her view, are normatively inert from the epistemic point of view unless one is a member of a group. Reflection on Desert Island suggests that this way of thinking is erroneous; truth-related considerations are relevant to what one ought to believe even when one is not a member of a group.

To mitigate the problem, the best strategy available to the social instrumentalist is to try to explain away our intuitions. She can claim that there are still reasons for this agent to believe things; however, instead of being epistemic, these reasons are purely practical. As Michelle Dyke puts it, this agent:

> would still have purely instrumental reasons to promote her own actual interests. So, for instance, if the last human survivor on Earth aims at surviving for much longer, then she had better learn to form reliable beliefs about where she has just left her provisions of food and water. What she would lack are the kinds of distinctively epistemic reasons for beliefs that could not be accounted for by any purely individualistic instrumental conception of epistemic reasons. (Dyke 2021, pp. 853–854)

I think that we should be dissatisfied with this explanation. First of all, we are free to stipulate that the belief is of no practical import and that this occurrence of believing against the evidence would not bring about changes in Jan's otherwise reliable belief-forming dispositions. In this case, the 'ought' in 'Jan is not believing as he ought to' is not the practical ought. If it is true that Jan is not believing as he ought to, it has to be that he is not believing as he *epistemically* ought to, which the social instrumentalist cannot accept.

Thus, the social instrumentalist has to claim that it is false that Jan is not believing as he ought to. But biting the bullet here is ill-motivated: social instrumentalism derived its appeal from its ability to account for the apparent universality—or intersubjectivity—of epistemic reasons that traditional instrumentalism could not countenance. Now, faced with an example that appeals to the *very same* feature of epistemic reasons, social instrumentalists are biting the bullet. Their view still does not respect this apparent universality and thus does not constitute (at least in this respect) a superior alternative to traditional instrumentalism. Epistemic reasons are certainly less fickle on social instrumentalist views. But the problem, it turns out, is not that they merely ought to be less fickle than the traditional instrumentalist could account for; it is that they need to be universal. What we should conclude from the 'Desert Island' objection is that epistemic reasons are not merely intersubjective: they truly are universal. The intuition at work in DESERT ISLAND is the same that is at work in SPOILER ALERT, and only a conception of epistemic normativity that makes epistemic reasons universal will respect it.

Another way of looking at the problem is as follows: on the social instrumentalist's view whether a certain fact that supports the truth of a proposition is an epistemic reason to believe is a contingent fact. It depends upon further contingent facts about the ends of communities or those of individual agents. But pieces of evidence are not contingently related to what we ought to believe. Any piece of evidence that bears on whether *p* is relevant to the doxastic attitude we should have towards *p*. Necessarily, any piece of evidence is a reason to believe.

It seems that the only option left to the social instrumentalist is to deny that the fact that the coin landed on head 99 times out of a 100 tosses is evidence in favor of the proposition that the coin is not fair: if what counts as evidence is determined by the epistemic standards in effect within a society, then, according to the social instrumentalist, that the coin landed heads nine-nine times out of a hundred may simply *not* be evidence for the proposition in question. But this fact should uncontroversially count as evidence against the fairness of the coin. It thus seems that the social instrumentalist cannot provide a satisfying answer to the 'desert island' objection.

Conclusion: epistemic normativity is not essentially social

The arguments presented above cast doubt, not just on social instrumentalism, but on any view of epistemic normativity on which it is essentially social, that is, any view that entails that epistemic norms get their contents and reason-giving force solely from contingent facts about our social settings. This is immediately clear from the 'desert island' objection, which trivially generalizes to any such account. But I think that the two other arguments I present, the argument from trade-offs and the argument from acceptance, also generalize. Here is why, succinctly put: it is plausible that what explains why we have genuinely normative reasons to act in accordance with a set of social norms is at least in part the fact that these norms serve some legitimate social purpose.[18] In other words, it is that they serve to bring about some good for the community members, in contrast to other social norms that are harmful, inefficient or costly. But whether a set of norms brings about some good for community members, and how efficiently it does so, is bound to be contingent—to depend on the specific circumstances in which a community finds itself. This will, in turn, make epistemic normativity dependent on these contingent features, and thus invite the objections from trade-offs and acceptance, which turn on this very aspect of social instrumentalism.

Social instrumentalism brings to the forefront an important aspect of epistemic normativity: that it fulfills certain important social functions. However, the fact that our practices of epistemic evaluation answer some collective needs does not entail that epistemic normativity is, at its core, a social phenomenon. Social instrumentalism, and other views that would make epistemic normativity fundamentally social, go wrong in claiming that epistemic norms get their contents and reason-giving force from some contingent facts about our social settings, such as what best promotes a community's goals. There is much one can accept about the social features of epistemic norms without accepting this thesis.

To conclude, let's return to the puzzle presented in §1. The arguments presented above point toward a certain picture of epistemic norms: that they, at their core, are binding in a way that does not depend on contingent features of our individual or social make-up. It seems that no instrumentalist view, however sophisticated, can do justice to that. If we are to find a good account of epistemic normativity, we ought to look elsewhere,s away from instrumentalism.

18. This point is also suggested by Hannon and Woodard (forth. P. 13)

References for Chapter 2

Berker, Selim (2013a). "Epistemic Teleology and the Separateness of Propositions". In: *Philosophical Review* 122.3, pp. 337–393.

— (2013b). "The Rejection of Epistemic Consequentialism". In: *Philosophical Issues* 23.1, pp. 363–387.

Buckley, Daniel (2020). "Varieties of Epistemic Instrumentalism". In: *Synthese* 198.10, pp. 9293–9313.

Cohen, L. Jonathan (1992). *An Essay on Belief and Acceptance*. Oxford: Clarendon Press.

Côté-Bouchard, Charles (2016). "Can the Aim of Belief Ground Epistemic Normativity?" In: *Philosophical Studies* 173.12, pp. 3181–3198.

— (2021). "Two Types of Epistemic Instrumentalism". In: *Synthese* 198.6, pp. 5455–5475.

Craig, Edward (1990). *Knowledge and the State of Nature: An Essay in Conceptual Synthesis*. Oxford: Clarendon Press. 169 pp.

Dogramaci, Sinan (2012). "Reverse Engineering Epistemic Evaluations". In: *Philosophy and Phenomenological Research* 84.3, pp. 513–530.

Dyke, Michelle M. (2021). "Could Our Epistemic Reasons Be Collective Practical Reasons?" In: *Noûs* 55.4, pp. 842–862.

Graham, Peter J. (2015). "Epistemic Normativity and Social Norms". In: *Epistemic Evaluation*. Ed. by David K. Henderson and John Greco. Oxford University Press, pp. 246–273.

Hannon, Michael and Elise Woodard (forth.). "The Construction of Epistemic Normativity". In: *Philosophical Issues*.

Henderson, David K. and John Greco, eds. (2015). *Epistemic Evaluation: Purposeful Epistemology*. First edition. Oxford: Oxford University Press. 293 pp.

Ichikawa, Jonathan Jenkins (2022). "You Ought to Have Known: Positive Epistemic Norms in a Knowledge-First Framework". In: *Synthese* 200.5, p. 400.

Kelly, Thomas (2003). "Epistemic Rationality as Instrumental Rationality: A Critique". In: *Philosophy and Phenomenological Research* 66.3, pp. 612–640.

Kornblith, Hilary (1993). "Epistemic Normativity". In: *Synthese* 94, pp. 357–376.

Lackey, Jennifer (1999). "Testimonial Knowledge and Transmission". In: *The Philosophical Quarterly* 49.197, pp. 471–490.

Lockard, Matthew (2013). "Epistemic Instrumentalism". In: *Synthese* 190.9, pp. 1701–1718.

Nelson, M. T. (2010). "We Have No Positive Epistemic Duties". In: *Mind* 119.473, pp. 83–102.

Nolfi, Kate (2021). "Epistemic Norms, All Things Considered". In: *Synthese* 198.7, pp. 6717–6737.

Quine, Willard Van Orman (1998). “Reply to Morton White”. In: *The Philosophy of W. V. Quine*. Ed. by Lewis Edwin Hahn and Paul Arthur Schilpp. 2nd. expanded ed. Chicago, IL: Open Court, pp. 663–665.

Ramsey, Frank P. (1990). “General Propositions and Causality”. In: *Philosophical Papers*. Ed. by David H. Mellor. Cambridge: Cambridge University Press.

Shah, Nishi (2003). “How Truth Governs Belief”. In: *The Philosophical Review* 112.4, pp. 447–482.

Shah, Nishi and David Velleman (2005). “Doxastic Deliberation”. In: *The Philosophical Review* 114.4, pp. 497–534.

Simion, Mona (2024). “Resistance to Evidence and the Duty to Believe”. In: *Philosophy and Phenomenological Research* 108.1, pp. 203–216.

Velleman, David (2000). “On the Aim of Belief”. In: *The Possibility of Practical Reason*. New York, NY: Oxford University Press, pp. 244–281.

Wedgwood, Ralph (2002). “The Aim of Belief”. In: *Philosophical Perspectives* 16, pp. 267–297.

What's truth got to do with it?

Epistemic normativity, the value of truth, and the limits of veritism

"Granted, we will truth: *why not untruth instead?* And uncertainty? Even ignorance? The problem of the value of truth came before us,—or was it we who came before the problem?" Thus writes Nietzsche in the opening aphorism of *Beyond Good and Evil* (2002 [1886]). It is telling that, in the first work in which he made explicit his project of critiquing our present values in order to replace them with what he thought were "life-affirming" ones—the "revaluation of all values"—, among a vast set of values from which to pick, the first target he singles out for critical scrutiny is *truth*. No matter what one thinks of Nietzsche's program, it is undeniable that he had recognized in truth something that was, and continues to be, an object of unrivaled reverence in our axiological landscape.

We care about truth. We think that it matters. We orient our inquiries towards it, and devote considerable resources, both individually and collectively, to attain it. And it seems very likely that this fact, that truth matters, bears some important relation to what we ought to believe—to how we should go about living as agents who are in the predicament of having to form and revise doxastic attitudes. Why should I, a physician, carefully assess the evidence to determine what illness a patient has? Because believing the truth about the patient's illness matters for choosing the proper course of treatment, and this calls for a careful assessment of the evidence. Why should I not let my fear of embarrassment distort my judgment about whether my friend truly enjoyed the book I recommended? Because believing the truth about their reaction matters: genuine friendships depend on having an accurate picture of one another, and that includes having true beliefs about what they like. Explanations that appeal to the value of truth have also been endorsed by many of those less inclined to understand our obligations to believe in practical or ethical terms: truth, it is said, has epistemic value, or has value from the epistemic point of view; insofar as the way to attain this value is to have justified beliefs, or to believe in accordance with the evidence, that's what I ought to do.

This general view of the relationship between the value of truth and what I ought to believe is understood in many ways, all of which share the following commitment: that, at the most fundamental level, any explanation of what I ought to believe will appeal to the value of truth. Following others, I'll call this view *veritistic value monism*, or *veritism*.[1] It is a form of *value*

1. As far as I can tell, the term "veritism" entered the philosophical vocabulary through Goldman (1999). For notable discussions of veritism, see Ahlstrom-Vij (2013), Dandelet (2024), Pritchard (2014), and Sylvan (2018). Throughout the chapter, my main foil will be Dandelet (2024); it offers the clearest presentation of veritism

monism, because it recognizes only one fundamental source of value when it comes to the explanation of norms of belief-formation and belief-revision. It is *veritistic*, because this source of value is truth.[2] Veritism has been widely endorsed across the spectrum of epistemologists, be they reliabilists, accuracy-firsters, or proponents of virtue epistemology.[3]

A central part of my project in this dissertation has been to draw attention to the fact that *how* we understand the claim that truth is valuable is crucial to how we understand norms of belief-formation and belief-revision, and vice versa; being clear on why or in what sense truth matters is a task not to be neglected if we are to provide a satisfying account of epistemic normativity. Epistemologists have created a sizable body of literature on how the value of truth explains epistemic normativity. But they often have dealt quickly with the question of the value of truth itself by appealing to the idea that truth is epistemically valuable—it is of value *from the epistemic point of view*. In the first part of this chapter, I will argue that once we look at accounts of the value of truth, we find none that simultaneously meets these two desiderata: 1) it provides a satisfying explanation of why truth matters and 2) it can explain *all* that is subsumed under the label "epistemic normativity." And, furthermore, I will argue that we shouldn't expect to find an account that satisfies both desiderata.

Does that mean that we should abandon the idea that there is an important relation between the value of truth and epistemic normativity? Others may want to draw this conclusion and seek another fundamental explanation of our epistemic norms. I think that veritists are on to something important when they recognize that the value of truth has a significant connection with norms of belief-formation and belief-revision. Many—and probably our most important—epistemic norms are grounded in the value of truth. The issue with veritism is the attempt to explain *all* of them by appeal to the value of truth. We can accommodate the veritist's insight about the value of truth if we embrace pluralism about epistemic normativity: the idea that there is no single, fundamental principle in which every explanation of what we should believe bottoms out. I distinguish between what I call "substantive epistemic normativity" and "basal epistemic normativity", and argue that each constitutes an irreducible component of epistemic normativity. *Substantive epistemic normativity* covers norms of belief-formation and belief-revision that are grounded in the value of truth. *Basal epistemic normativity* covers norms that are explained by appeal to what constitutes ordinary well-functioning mechanisms of belief-formation and belief-revision for beings like us.

available, and at the same time exemplifies precisely what I think goes wrong in many discussions of how epistemic norms relate to the value of truth.

2. In contrast with others cited above, I do not use the term "epistemic" to describe the view of the value at play; the reason for this will become clear later in the paper.
3. Reliabilists: Goldman (1999). Accuracy-firsters: Pettigrew (2016). Virtue epistemologists: Pritchard (2014).

Two clarifications before proceeding further. Importantly, the distinction between substantive and basal normativity does not amount to, nor does it track the distinction between *obligations* or *duties* to believe and *permissions* to believe. Some attitudes may be permissible according to substantive norms, and some attitudes may be obligatory or forbidden, according to basal epistemic norms, and vice versa. Additionally, the distinction is not and does not track the distinction between *positive* epistemic norms (norms requiring that one adopt a certain doxastic attitude) and *negative* epistemic norms (norms forbidding holding certain doxastic attitudes). There are both positive and negative substantive epistemic norms, and this holds equally well of basal epistemic norms.

Finally, it is worth noting that this chapter does not go the full way toward defending a theory of epistemic normativity. Such a project would require defending the superiority of the view developed here over rival approaches that do not take the value of truth to play a central explanatory role when it comes to epistemic norms—most notably, constitutivist views, on which epistemic norms are explained, at the most fundamental level, by the very nature of belief, as well as accounts that explain epistemic normativity solely in terms of biological function.[4] My aim here is more modest: to begin from the idea—a very appealing one—that the value of truth plays an important role in determining what we should believe, and to work out the best way of making sense of that idea.

Here is the plan for the rest of the chapter. In §1, I explain veritism and its appeal. In §2, I discuss extant accounts of the value of truth and consider how they fare when it comes to explaining epistemic normativity. In §3, I present my own view of how the value of truth plays a role in the explanation of epistemic norms. The type of epistemic normativity I discuss in this section is what I label "substantive epistemic normativity." In §4, I turn to cases in which truth does not seem to matter in any substantive way, yet we are still inclined to think that certain doxastic responses are expected or appropriate. I argue that such cases are best explained by what I call "basal epistemic normativity." In §5, I recapitulate the bipartite view I have developed over the chapter and respond to some objections to the account.

3.1 Veritism

What role does the value of truth play in explaining our epistemic norms? By far the most discussed idea in the literature on this question is *veritism*. The central idea behind veritism is that epistemic norms can be explained in terms of fitting responses to a single source of value: truth.[5]

4. More on constitutivism in section §1.

5. Because there is a wide diversity of ways of conceive of the demands that epistemic norms place on us, I will follow Dandelet (2024)'s terminology and talk of doxastic attitudes being "fitting" or "appropriate" to hold, given the evidence, in virtue of being "fitting responses" to the value of truth, or alternatively, of "proper" or

Proponents of veritism disagree on what the proper way of responding to the value of truth is—epistemic consequentialists tend to construe it in terms of accuracy maximization (e.g., Pettigrew 2016), virtue epistemologists understand it as virtuous belief (e.g., Pritchard 2014), and deontologists as respect for truth (e.g., Sylvan 2020). But all these share the core commitment that the doxastic attitudes sanctioned by our epistemic norms are so because they constitute proper responses to the value of truth.

As Dandelet (2024) presents it, veritism contains three claims: "(1) epistemic rationality [what one should believe/what doxastic attitude is appropriate to hold given one's evidence] is defined by its relationship to some *value*, (2) there is only *one* such value, and (3) this value is *truth*." I want to insist on (1) (and on (3), by extension), as my discussion of veritism largely serves as a springboard for an examination of the relation between the value of truth and our epistemic norms, and the view I end up proposing is compatible with different claims about whether it is the only value that makes some doxastic attitudes appropriate.[6] The veritist attempts to explain what attitudes are sanctioned by norms of belief-formation and belief-revision by appeal to value; said another way: norms of belief-formation and belief-revision are *analyzed in terms of value*. According to the veritist, what doxastic attitude is appropriate/justified/permissible/rational is given one's evidence is appropriate *because* of the relation in which it stands to some value: it is a fitting response to the value of truth.[7]

It is important to note, however, that not all explanations of normative phenomena need involve an appeal to value. Some are *norm-first*. Consider punctuality: one might say that arriving late is inappropriate simply because it violates the norm to arrive on time—that is, the explanation of why being late is inappropriate bottoms out in by an appeal to a norm that require us to be punctual. But one might also offer a value-based explanation: being late is inappropriate because it wastes other people's time. On this second view, what fundamentally explains why it is inappropriate to be is how it relates to something independently (dis)valuable.

In the case of epistemic normativity, epistemic constitutivism paradigmatically exemplifies a norms-based approach.[8] In its traditional form, epistemic constitutivism takes as its starting

"appropriate responses", to talk about the attitudes that are sanctioned by epistemic norms; I do not mean much to hinge on this choice of vocabulary.

6. Other values that have been thought to place demands on how we ought to believe that are in the same cluster as truth include *understanding* and *knowledge*, but the debate goes on about whether their value derives from that of truth or is equally fundamental.

7. It is worth noting that how much interest one takes in this view, and in the debates surrounding it, may hinge on one's preferred account of value. Much of the literature seems to presuppose, or at least fit neatly with, a robust conception of value on which it is a primitive, unanalyzable notion. By contrast, if one endorses a fitting-attitude analysis of value—on which to be valuable *just is* for certain responses to be fitting (e.g., to be admirable is for it to be fitting to admire)—then the idea that appealing to value can *explain why* certain responses are fitting may seem misguided. In what follows, I adopt the basic presupposition of the literature.

8. See, for instance, Nolfi (2015) and Wedgwood (2002).

point the claim that belief aims at truth, understood as the idea that belief, *qua* belief, is subject to a norm of correctness: one ought to believe *p* if and only if *p*.[9] It is constitutive of what it is to be a belief to be subject to this norm; a doxastic attitude that is not subject to this norm is not a belief. Since we are believers, we are subject to this norm. Other epistemic norms are explained by the fact that they stand in some appropriate relation to the norm of belief. For instance, the norm requiring one to believe in accordance with one's evidence could be explained by the fact that it provides guidance for satisfying the norm of belief: believing in accordance with one's evidence is typically the best one can do, from one's own perspective, to satisfy this norm. In the constitutivist's explanation of epistemic norms, value plays no role at all. At the most fundamental level, the constitutivist's explanation of the content of epistemic norms appeals to the norm of belief; it thus gives a *norms-first* explanation of epistemic normativity.

Additionally, it is worth noting that an explanation of epistemic norms and of epistemic notions such as rationality or justification, can appeal, at the most fundamental level, to truth *without* appealing to the *value* of truth. Consider the norm requiring one to believe in accordance with one's evidence. An answer to the question of what explains why one ought to believe in accordance with the evidence could go as follows: one ought to believe in accordance to their evidence *because* believing in accordance with one's evidence is truth-conducive. In this case, the explanation has a teleological shape, and is explained by the fact that believing in this way is likely to lead to forming a true belief. But it doesn't appeal to truth's being valuable in any way.[10]

Veritism is an appealing view. It both easily accounts for the central role that truth plays in epistemic evaluations—*via* the claim that it is valuable—, while being extremely parsimonious, as it takes truth to be the sole fundamental value for the explanation of epistemic norms. Other normative or evaluative epistemic notions, such as those related to the epistemic status of a belief (whether it is justified, rational, or warranted), are explained, at the most fundamental level, by appeal to the value of truth. Any other property that is epistemically valuable derives its value from its relation to truth's value. As Sylvan writes,

> [Veritism's] explanatory power is intuitive. It is intuitive that we care about things like rationality and justification because we care about accuracy. (Sylvan 2018, p. 407)

If, for instance, some intellectual virtues are epistemically valuable, it is because they lead to something that is of more fundamental value: truth; similarly, we can show that coherence matters if it stands in the right relation to truth, such as being conducive to accuracy maximization or manifesting respect for truth. Thus, veritism is axiologically parsimonious and accounts for the central

9. Compare: a map, *qua* map (and not, for instance, *qua* aesthetic object), is correct only if it represents the spatial relations of the territory it depicts accurately enough for its user. This standard is constitutive of what it is to be a map.
10. This distinction is important for the first chapter of my dissertation.

role that truth plays in epistemic evaluations: it only posits one source of value when it comes to epistemic normativity, and attempts to explain all epistemic norms by their relation to this value.

Finally, and importantly, although it is, at its core, a view about the *contents* of epistemic norms—it attempts to explain why epistemic norms recommend (permit/require/forbid) what they recommend (permit/require/forbid), what makes a belief justified or rational, etc.—, it also has the added benefit of explaining why we care about epistemic normativity. Truth matters; the fact that truth matters further explains why we care about other epistemic goods, such as justification and rationality. Here's Dandelet expressing this claim:

> [A] fundamental question about epistemic rationality that veritism might help us answer is the question of significance. This question asks: What is the significance of epistemic rationality? An answer to this question should "help those who are concerned with [epistemic rationality] understand why the reasons it provides are so important, or what, exactly, their importance is." In other words, it should help us see why failures of epistemic rationality are not just failures to conform to some standard or other, but failures to live up to a standard that really matters in some way.

Veritism thus provides an answer to the question of why we care, or ought to care, about the epistemic status of beliefs—about having rational or justified beliefs, about having beliefs that accord with the evidence, etc.: it's because beliefs that have these properties stand in the right relation to something that matters: truth.

3.2 The value of truth

In what follows, I survey different accounts of the value of truth, and I argue that none of them is simultaneously plausible as an account of the value of truth *while* doing the work the veritist wants it to do, that is, provide a fundamental principle that explains our epistemic norms.

Pure intellectualism: truth is valuable from the epistemic point of view

Many veritists tend to deal quickly with questions related to the value of truth by claiming that truth has value "from the epistemic point of view", and then proceeding to explain epistemic normativity in terms of fitting responses to this value. Because of the popularity of this approach, I'll deal with it first. I'll call this view of the value of truth "pure intellectualism".

> PURE INTELLECTUALISM Truth is theoretically/cognitively/epistemically valuable.[11]

According to PURE INTELLECTUALISM, truth is a domain-relative good. Domain-relative goods are things that are valuable *from the standpoint of a certain normative domain*. Take com-

11. See for instance Kvanvig (2008) for an explicit defense of the view; the view is found expressed and endorsed throughout the literature on epistemic normativity.

petitive chess playing as an instance. From the point of view of competitive chess, there is one goal toward which one should aim: winning. Winning, then, is valuable within the normative domain of chess. But this does not entail that winning a chess game is valuable or good *outside* that domain. A domain-relative good that lacks value beyond its normative domain is a *merely* domain-relative good. Contrast this with beauty or friendship: these are goods that matter always and everywhere—they are good *simpliciter*.

The appeal of PURE INTELLECTUALISM lies in this analogy: just as winning is *always* good *within* the domain of chess, truth is *always* good within the theoretical or epistemic domain. On this view, we have an account that purports to explain the value not just of truths that are useful, but of all truths. However, the appeal to domain-relativity threatens to make the proposal worryingly uninformative. The theoretical domain is, by definition, the one concerned with the pursuit of truth and knowledge. To say that truth is theoretically valuable amounts to saying that, from the standpoint concerned with truth, truth matters. That seems trivial. If we already had an independent grip on the notion of the "epistemic domain," the proposal might carry more weight. But "epistemic" is widely acknowledged to be a philosophical term of art, and whatever informativeness the claim seems to have may disappear once we stipulate (as we must) how the term is to be used.[12] Although I take this worry to be worth flagging, I don't want too much to rest on it. There is, I think, a much deeper problem for PURE INTELLECTUALISM: although it is presented as an account of the value of truth, it fails to capture a crucial feature of our ordinary discourse about truth: that it often *matters*—and that it does so *simpliciter*. By construing the value of truth entirely in domain-relative terms, the view leaves open the question of why we should care about it at all. Compare: from the standpoint of the Guild of Right Angle Appreciators, it is good that the two pens on my desk form a right angle. We can still ask: *but does it matter*? And it is perfectly compatible with the view to answer "no."

Some veritists may not be worried by the fact that they cannot answer this question. These veritists adopt a two-tiered view of epistemic normativity: there are epistemic norms explained by truth's being epistemically valuable, and then there is a further, distinct question of whether we ought to abide by those norms, one whose answer may not involve the value of truth at all. I think they should be worried.

First, it's worth nothing that an important part of the appeal of veritism comes from the fact that it answered the question of why we should care about believing well. The veritist who pairs her view with PURE INTELLECTUALISM thus foregoes one of veritism's chief attractions. Second, I think it's harder to cleave the question of what we should believe apart from the ques-

12. For a discussion of different ways of specifying the meaning of "epistemic" and their implications for epistemological theorizing, see Cohen (2016). Alston, for instance, defines the epistemic domain as the one concerned with the maximization of true beliefs and the minimization of false beliefs. See Alston (2006, pp. 29-30).

tion of the value of truth in the way that the two-tiered veritist thinks she can. The veritist thinks that the value of truth plays a real explanatory role when it comes to accounting for what epistemic norms prescribe, permit or forbid, or what constitutes a "fitting" response to this value. To begin, it's hard to even wrap one's head around the idea that there could be "proper" or "fitting" responses to something that is not truly valuable. In fact, it seems that it is the very fact that something is valuable *simpliciter* that makes it the case that there are norms regarding how we should relate to it and that determines what these norms are.

Consider pleasure. Pleasure is valuable in a way that makes it the case that, *ceteris paribus*, there is normative pressure not to settle for less pleasure than one could have. Now contrast that with friendship. Friendship is valuable, but the value it possesses does not require us to maximize the number of friends we have. Likewise for politeness: it matters, but it does not require us to stand at the door all day holding it open for every passerby. Beauty is valuable, but it does not demand that we create as many beautiful objects as possible; it calls, instead, for preservation or appreciation. In each case, it is the distinctive way in which something is valuable that determines the kind of responses to it that it calls for.

Now contrast these with right angles. Imagine that the Guild of Right Angle Appreciators engages in a debate about what constitutes a fitting response to right angles, which are of fundamental value "from the standpoint of the Guild of Right Angle Appreciators". Some argue that we should maximize the amount of right angles in the world; others advocate for attitudes of respect towards them; still others propose that we should act in a way that conforms to a form of right-angle-directed virtue. How should this question be settled? The claim that "right angles are valuable, from the Guild's perspective" does no normative work here. It cannot adjudicate between competing responses and cannot explain why any response is fitting in the first place.

The point of these examples is that not everything valuable generates the same kinds of demands; value generates normative pressures that vary in kind depending on the nature of the value. The fittingness of a response depends on what sort of value is at stake. But to be able to make intelligible the idea that a response to something is fitting, or that it places normative requirements on us, this thing must be valuable.

Go back to the Guild of Right Angles Appreciators once again. The only way in which we can make sense of the claim that "right angles are valuable, from the standpoint of the Guild of Right Angle Appreciators", is to observe their norms, and see that they prescribe some ways of relating to right angles. But then, the explanatory direction does not flow *from the value of right angles* to *the fitting responses to this value*, quite the opposite! We say that "right angles are valuable, from the perspective of the Guild of Right Angles Appreciators" *because* the norms adopted by the Guild prescribe these ways of favorably relating to right angles. What we are dealing with is really a norms-first view in disguise! Having a view in which "right angles" occupies a role that is structurally similar to the role that values *simpliciter* have in theories about how we should act is not equivalent to having a view that explains norms in terms of values. In the case of the Guild, the norms are given as brute facts; they fix what counts as "value" within the standpoint.

To drive the point home further, suppose that the norms of the Guild recommended maximizing the number of right angles in the world. But then we discover that one of their norms prescribes not maximizing in certain circumstances as a kind of side constraint. Should we think that they have misunderstood the value of right angles from their own standpoint, or that this norm doesn't prescribe a fitting response to the value of right angle from their own standpoint? No! We should instead revise our conception of what the value of right angles are, from the standpoint of the Guild. This further shows that the merely domain-relative value of right angles plays no role in explaining these norms.

It should thus come as no surprise that Dandelet (2024) puzzles over how to derive epistemic norms from the value of truth in the way the veritist attempts to, and concludes that the veritist must retreat to viewing certain responses as fitting or unfitting as a matter of brute fact. As I've argued here, this outcome follows directly from the veritist's (and Dandelet's) reliance on a merely domain-relative notion of value, one that cannot bear the explanatory weight it is meant to carry. By appealing only to truth as "epistemic value," the veritist deprives herself of the very resources needed to sustain the explanatory project she sets out to pursue. Dandelet takes this failure to show that veritism is untenable. But that conclusion is too quick. The right response is not to abandon veritism, but to ask whether truth, when understood as valuable *simpliciter*, can fulfill the explanatory role the veritist thinks it does.

Because it has no trouble treating *all* truths as valuable (in the domain-relative sense of being epistemically valuable) pure intellectualism initially appeared well-suited to ground an explanation of our epistemic norms. But we have seen that it fails to meet this explanatory demand. I now turn to other extant accounts of the value of truth: instrumentalism, Andrew Reisner's welfarist pluralism, the natural curiosity view, and Zagzebski's view that the value of truth is grounded in the fact that we are rationally required to care about truth. Each of these, I argue, offers a compelling explanation of why *some* truths matter, but none succeeds in explaining why *all* truths do. What follows is that the claim that epistemic norms are explained by the value of truth ought to be restricted: *some* epistemic norms are explained by the value of truth. This motivates a turn towards pluralism: there is not just one way in which truth matters, but many, and all are relevant to explaining epistemic normativity. In the next section ("Substantive epistemic normativity and the value of truth"), I show how this radically pluralist account can explain a broad range of norms of belief-formation and belief-revision while preserving some of the spirit, although not the letter, of veritism.

While this may seem to strike a decisive blow against veritism, that need not be the case. The veritist's insight is that the value of truth is central to what we ought to believe, and this very appealing idea be preserved if we adopt a pluralist account of epistemic normativity. What emerges then, I believe, is a richer and more realistic picture of our lives as believers, one on which these lives are deeply and inextricably intertwined with our various other normative commitments.

Instrumentalism: truth is valuable as a means of achieving our aims

An intuitively plausible explanation of the value of truth appeals to the idea that, as a rule, true beliefs tend to be useful. As Ramsey (1990) famously put it, "beliefs are a map by which we steer." If I want a snack, it's helpful to know where my snacks are stored; if I want to get to work efficiently, it's useful to know the quickest route. By contrast, a false belief about whether I have flour at home might frustrate my cooking plans, and thinking my train leaves at 8:30 when it actually leaves at 8:00 could cause serious logistical problems. The usefulness of having true beliefs and the hindrance caused by false beliefs are already apparent in these low-stakes cases; but the stakes of having accurate beliefs increase with the stakes of the actions they would motivate. In general, then, true beliefs help us navigate the world, while false beliefs hinder us. Call this view *instrumentalism*.

> Instrumentalism True beliefs are valuable insofar as they are useful for navigating the world successfully.[13]

Instrumentalism cannot provide an adequate account of why all truths matter: first, it yields the verdict that certain truths are disvaluable; second, it fails to account for the value of some truths that are not useful for navigating the world successfully. If usefulness is what makes truth valuable, then false beliefs might often turn out to be just as valuable, or more valuable, than their true counterparts. Usefulness is not the prerogative of true beliefs. Whether a belief is useful is contingent; for instance, you may struggle to be timely, in which case a false belief that your work commute takes 45 minutes (rather than the actual 30 minutes it takes) might help you get to work on time; this belief would not be useful if you were a timely person. McKay and Dennett (2009) have identified a set of false beliefs that were predictably adaptive in the context of natural selection and continue to be beneficial to those who have them: positive illusions—*viz.*, unrealistically positive self-appraisals and appraisals of sexual partners and children as well as

13. The *locus classicus* of this view is Kornblith (1993). For a more recent defense that aims to avoid the objections to accounts such as Kornblith's, see Willoughby (2022). For a social version of instrumentalism, see Hannon and Woodard (forth.)

unwarranted optimism. Such beliefs are thought to propel adaptive action, directly sustain and enhance health, and strengthen social bonds. For example, believing that one will succeed can lead to greater effort and achievement, and overestimating a partner's positive attributes can bolster romantic relationships. Thus, instrumentalism yields the verdict that believing a falsehood about p may be more valuable than believing the truth about that proposition.

Second, it is not necessary that a truth to be useful for it to be valuable. To see this, consider "basic" or "fundamental" research—research that is driven simply by fascination with a question, or aimed at increasing our knowledge and understanding of the world without consideration of the practical upshots of the findings, such as research into the foundations of physics and mathematics. Intuitively, it seems that the knowledge thus acquired can be of deep value, providing us insight and understanding into important aspects of our world; it is doubtful, however, that we reap any practical benefits from them. The pragmatist could try to explain the value of these truths by appealing to their future potential usefulness: we cannot know whether basic truths will be valuable in the future, so these truths are worth caring about. Plausibly, however, we can and do often know that basic truths will not turn out to be useful. Thus, while it is plausible that some truths matter in virtue of the fact that believing them helps us achieve our aims and fulfill our desires, instrumentalism does not provide an account of why all truths matter.

Welfarism: truth is valuable in virtue of its contribution to well-being

Reisner (2023) develops an account of epistemic normativity grounded in the value of truth. He argues that the value of true belief lies in its contribution to well-being. Part of this contribution, he thinks, is as the instrumentalist conceives of it: it comes from the fact that, as a rule, accurate representations, including beliefs, help us achieve our goals and fulfill our desires. In contrast with the instrumentalist, however, he holds that their contribution does not end there. A belief's being true or justified—what Reisner calls "being in a positive epistemic state"—*directly* contributes to well-being simply in virtue of its being true or justified. On the welfarist pluralism account, the etiology of our mental states—crucially, whether they are properly connected to the world—is a partial determinant of our well-being.

Being connected in the right way to the world appears to be an important component of well-being, regardless of the practical benefits to be reaped from it. Consider a case in the style of Nozick's experience machine: two perfect internal duplicates who seem to be living the same life. One, however, inhabits a simulated world, experiencing only simulacra, while the other's experiences are anchored in reality. Even if their experiences are phenomenally identical, our intuitions tend to support the idea that the life of the person whose experiences are tethered to reality is preferable; that person seems to be better off. Thus, it seems that being connected in the right way to the world is a partial determinant of well-being. According to Reisner, there are two distinct ways in which a belief can instantiate the property of being in a positive epistemic state: by being true, or by being formed on the basis of epistemic reasons. Each contributes to well-being even if the other is absent.

Reisner's proposal does much better than regular instrumentalism at accounting for the value of truth *in general*; because true beliefs are thought to contribute to well-being not merely in virtue of their consequences, but also directly, in an unmediated manner, pluralist welfarism can account for the value of truths that are of no practical benefit: believing them contributes directly to our well-being. Additionally, it does explain what is distinctively valuable about true beliefs: it's that they instantiate the property of *being properly connected to the world*, which matters. As I see it, a notable set of truths whose value the welfarist account seems especially well-suited to explain are truths that can be said to have *existential* import; truths that are meaningful regardless of one cares about them or of their usefulness, the paradigmatic example of which is the truth about whether God exists. There is certainly a practical import to believing the truth about this question; however, the value of that truth does not seem reducible to it. Other examples include, for instance, truths about our origins, about the existence of intelligent life elsewhere in the universe, about whether free will is compatible with determinism, and so on. Given the kind of beings we are, someone who is right about such matters does seem better off, all else being equal, than someone who is not.

However, Reisner's argument falls short of establishing that all truths are valuable. What it shows is that being properly connected to the world matters and makes us better off, on the whole, than someone who substantially or thoroughly lacks this connection. But this does not entail that every increase in this connection makes us better off. In fact, there are some ways in which we should desire *not* to be connected to the world. Put another way: sometimes, ignorance is valuable. Consider believing intrusive truths about friends, such as those revealed in jealously guarded, intensely private diary entries. Suppose Jim and Jones are counterparts inhabiting different possible worlds. They are alike in nearly every respect, including in their deep friendship to Jean. Where they differ is that Jim has true beliefs about all the intimate thoughts Jean has ever recorded in her diary—thoughts never meant for anyone else to read. Jim's connection to the world is greater than Jones's. But Jim is not better off than Jones for having these additional true beliefs. We should prefer to be like Jones than like Jim, if given the choice. This shows that, although true beliefs may generally contribute to well-being in the way Reisner thinks they do, not *all* true beliefs do.

Natural curiosity: truth is valuable because it satisfies our natural curiosity

A venerable and popular account of the value of truth in its capacity to satisfy our *natural curiosity*: as Aristotle famously wrote, we humans naturally desire to know. And this natural curiosity is fulfilled by true beliefs.

> Natural curiosity Truth is valuable insofar as it answers our intellectual inquiries and, by doing so, satisfies our natural curiosity.[14]

The problem with the natural curiosity account is that satisfying our intellectual inquiries is not necessary for truth to matter. Some truths—for instance, the truth about whether God exists—would matter even if no one were curious about it. Faced with this problem, the proponent of the natural curiosity account might borrow a move from the playbook of the proponents of the desire-satisfaction theory of well-being and *idealize*: an idealized agent *should* be curious about these truths (as discussed in Brady (2009)). But this risks reversing the order of explanation: we think that an agent should be curious about these truths because they are valuable, not that they are valuable because they would satisfy the curiosity of an ideal agent. Therefore, even if some truths may be valuable in virtue of the fact that we are naturally curious about them, it is not necessary that they be for them to be valuable.

Furthermore, it seems unlikely that we are naturally curious about everything, or that all truths we are naturally curious about are thereby valuable. There are, in fact, propositions we should want to remain ignorant of. Return to the case of intrusive truths about friends discussed above. Suppose one could, at will, discover all the most private thoughts a friend has ever recorded in their diary. Even if this yielded a large number of true beliefs, and even if these could be acquired at no cost, one should still prefer not to believe them. Hence, the natural curiosity account cannot explain the value of all truths. Once again, we are dealing with an account of the value of truth that explains the value of some, but not all truths.

Truth is valuable in the sense that we are rationally required to care about it

According to Zagzebski, if one cares about a certain subject matter, then one rationally ought to care about the truth with respect to it, and, by extension, to be epistemically conscientious when it comes to forming beliefs about it. This is the sense in which truth matters. Suppose you care about your friend. Then, there will be certain things you should want to know about your friend in order to act in ways that align with your caring about them, and you should also want to know things about her regardless of how they might guide your actions. By extension, you should also try to have, seek, and form beliefs about them in an epistemically virtuous manner, says Zagzebski.

14. See for instance Goldman (1999, p. 3).

This argument generalizes: for any domain or subject matter *D*, if *S* cares about *D*, then *S* should want to act successfully when it comes to *D* or to know things about *D*. If *S* wants such things, then *S* ought to care about having true beliefs and knowledge with respect to *D*, and, by extension, to try to have, seek and form beliefs in an epistemically virtuous manner. Hence, for any domain *D* and agent *S*, if *S* cares about *D*, then *S* ought to care about the epistemic goods (truth, knowledge, justified beliefs, epistemic virtues) with respect to D.

On that proposal, the value of truth is conditional upon caring, and the requirement to care about truths regarding the subject matters we care about is a requirement of rationality. As Zagzbeksi writes,

> What unifies the demands to be epistemically conscientious is what we care about. This explains why we think that persons with epistemically unjustified beliefs are irrational. They are irrational because they have beliefs that oppose demands of what *they* care about. (Zagzebski 2004, p. 376)

It does seem that the fact that I care about something or someone can make it the case that I ought to care about certain truths. If I care about astronomy, then there are facts about stars, planets, and so on, that I should want to know—these truths matter to me in virtue of my interest in astronomy. If I were to lose that interest, they might no longer matter to me. However, if a truth is valuable in virtue of the fact that I rationally ought to care about it *given what I care about*, then it is simply not the case that all truths are valuable, because there are many domains about which I do not care and about which it's not the case that I ought to care.[15]

However, the account does point out to one important feature of the value of truth that hasn't been discussed yet. Often, its value is *relative to agent*, and sensitive to what moral obligations and obligations anchored in their social roles the agent is under. As a doctor, the truth about what the most effective ways to cure common ailments are is valuable in a way that it isn't for someone who is not a doctor. As a friend, the truth about whether my friend is trustworthy matters in a way that it doesn't to someone who isn't their friend. Belonging to a group that has been historically privileged may make it the case that believing the truth about a my privilege, its nature, extent, and effects has a special kind of value.

Therefore, although Zagzebski's account does not explain the value of all truths, it underscores an important aspect of the value of truth: some truths matter for particular agents in virtue of the social roles they occupy, the relationships they stand in, and the moral obligations under which they are.

15. Even more, there are domains about which I ought *not* to care: consider, once again the contents of a friend's diary.

3.3 Substantive epistemic normativity and the value of truth

Before I go on to explain my proposal about how we should think the value of truth relates to epistemic normativity, I want to highlight what I take to be an important feature of our lives as believers: that our identity as agents who form and revise beliefs and are subject to norms in doing so is inextricably tied to the broader web of normative commitments—moral, prudential, social, etc.—that shape our lives. When we ask ourselves the question *"What should I believe?"* we do so not in a vacuum, not in circumstances that include just our evidence, but in circumstances in which practical interests, moral requirements, social roles, and other normative considerations are already in play. While it may often be helpful to abstract away from this normative web in order to isolate one specific phenomenon, I think that, when it comes to explaining how epistemic normativity relates to the value of truth, such abstraction is misleading, because whether and how a particular truth is valuable cannot be determined *apart* from this broader normative web.

The previous survey illustrated many ways in which truth matters, many reasons to care about it, and it seems unlikely that one is more fundamental than the others for the purpose of explaining what we ought to believe. Sometimes a truth matter because believing truly would help us achieve our goals; sometimes we happen to care about it for idiosyncratic reasons; being properly connected to the world does seem to be a component of well-being, as Reisner has identified, especially when it comes to important or meaningful truths; we have moral obligations to others that often involve caring about truth; to be good occupants of certain social roles, we need to be attentive to certain truths. My proposal is that there is an important and distinct type of normativity subsumed under the label "epistemic normativity" that is explained by appeal to the value of truth, understood in this pluralistic way. I call this type of normativity "substantive epistemic normativity."

Let's break down the view I am putting forward in a few core claims. My aim is to articulate a form of epistemic normativity that is grounded in the value of truth, understood in terms of how particular truths matter in context.

1. Epistemic normativity includes a distinct type I call *substantive epistemic normativity*.
2. Substantive epistemic normativity is explained by appeal to the value of truth—that is, to facts about whether and how certain truths matter.
3. Facts about whether and how a truth matters can only be understood in relation to the broader normative web—including moral, prudential, and social considerations, among others—in which the agent is situated.

4. Facts about how and why a truth matters determine both (a) whether we are subject to a doxastic obligation concerning it, and (b) what doxastic attitudes are fitting or unfitting in the current circumstances.
5. This type of epistemic normativity has a strong agential flavor—that is, the agent is treated as responsible for her doxastic attitudes in a distinctive way, such that failures to adopt the appropriate attitudes makes her liable to blame.

To put it succinctly, the idea is that the way in which a particular truth matters determines *both* whether a doxastic response is required in some circumstances *and* what is a fitting response to these circumstances. So far, the discussion has been fairly abstract; I now want to consider how substantive epistemic normativity works by looking at cases. Consider this case:

> **Negligent landlord**
>
> Paola owns an apartment building. She glosses over the annual fire safety report and fails to form any belief about the safety of the property. In fact, the building is not up to fire safety standard.
>
> **Performance report (schoolboard member)**
>
> Leah is a school board member reviewing a report on district-wide student performance. The report shows that while reading scores have improved significantly overall, there has been a slight drop in performance among students being taught with method X, due to factors Y and Z. Leah reads the report but fails to register or form the belief that performance among students using method X has declined.

Both Paola and Leah's doxastic conduct is problematic: they fail to attend to and form beliefs about true propositions they are obligated to form beliefs about, given the evidence that they have. On the view I propose, their failure stems from a deeper normative failure; namely, they fail to respond appropriately to the significance of these truths (in this case, this importance stems from their roles and responsibilities, but this needn't be the case). The fire safety status of the building matters in a way that makes it incumbent on Paola, as a landlord, to be attentive to evidence that concerns it and to form a belief about it. Likewise, the drop in performance under method X matters in a way that makes it incumbent on Leah, as a school board member, to register the evidence and respond to it. They fail to track the value of the truth in question. In both cases, the truth's significance generates an obligation to believe. Contrast with the following case:

> **Performance report (passerby)**
>
> Rachel is a traveler passing through the district. Curious, she picks up a copy of the same report at a public library. She reads it casually and comes away with a general sense that "things are improving." She does not notice the finding about method X. When someone later asks her about it, she shrugs: "Oh, I didn't notice that; I couldn't tell you."

Leah and Rachel are both presented with the same evidence for the proposition that students taught with method X have shown a slight decline in performance. They both read the same report, which makes this fact available to them. Leah's failure to form the belief about the drop in performance among students is blameworthy; she *should have* formed that belief. In contrast, Rachel's isn't—it is in fact absolutely unproblematic. What explain this discrepancies is that, given the context in which Rachel is, the truth about the drop in performance *does not matter*. Therefore, there is no duty for Rachel to form a belief about it.

Truth need not matter in a specifically practical way for it to generate obligations. All it takes is that it matters. Here is another example to bolster this claim.

> **Pauline fragment**
>
> Albert is a scholar trained in New Testament textual criticism, though his current research is in Neoplatonic ethics. He comes across a papyrus fragment catalogued under a miscellaneous heading. The Greek phrasing, the epistolary structure, the historical hints, and the theological motifs all point to the unmistakable conclusion that the fragment belongs to a previously unknown letter by the apostle Paul. It is such that anyone with Albert's expertise could preliminarily authenticate it. But Albert doesn't register the evidence. He turns the page, makes no note of the fragment, and never forms any belief about its provenance or implications.
>
> **Parliamentary horticulture**
>
> Albert is a scholar trained in New Testament textual criticism, though his current research is in Neoplatonic ethics. While flipping through a government landscaping report for unrelated reasons, he reads a sentence indicating that there was an even number of blades of grass on the lawn of the Parliament in Ottawa on August 5, 1991. He is in a position to recognize the statement as true. But Albert doesn't register the evidence. He moves on, makes no note of the report, and never forms any belief about the number of blades of grass on the lawn of the Parliament in Ottawa on August 5, 1991.

In Pauline fragment, Albert is in a position to make an incredibly meaningful discovery. He just needs to form a belief that is well-supported by the evidence he encounters. The significance of the discovery is immense: it would deepen our understanding of an enormously influential historical and religious figure, shed light on early Christian thought, and answer questions that have occupied scholars for centuries. Because of the significance of the fact that the fragment was authored by Paul, and given the evidential circumstances in which he finds himself, Albert should have registered its significance and formed the belief supported by his evidence. To underscore this, imagine that the fragment is identified five hundred years later by another scholar. By a curious turn of events, this person uncovers clear documentation showing that Albert had seen the fragment centuries earlier. The sense of exasperation we would feel—"how could he not have seen what he was looking at?"—points to a normative failure: Albert missed something that, given the significance of the truth at issue, and given Albert's evidential circumstances, he ought to have grasped. Some form of blame would seem appropriate.

The contrast with Parliamentary Horticulture makes this point vivid. Here, too, Albert is presented with evidence that supports a true proposition. But the truth in question is trivial, devoid of any real significance. It is clear that Albert is entirely blameless in the latter case. Nothing makes the truth matter in a way that would demand uptake. That is the difference. And certainly, no one should blame Albert five centuries later for failing to form the belief supported by his evidence.

The cases discussed above strongly support the idea that there are obligations to form doxastic attitudes that are explained by the value of truth. Many would be willing to countenance the existence of such obligations and may even find it trivial. But I think that the role that the value of truth plays in explaining our norms of belief-formation and belief-revision goes much further. I think that appeal to this very same notion of the value of truth can determine *what doxastic responses are fitting*; *pace* Dandelet (2024), appealing to the value of truth can settle the question of whether a doxastic attitude is appropriate in a given case.[16]

> **Friendship suspicion**
>
> Nora hears a rumor that her close friend Sam has been secretly leaking confidential information at work. The rumor is based on circumstancial, good, but not conclusive evidence. Despite their long-standing friendship, Nora quickly forms the belief that Sam is guilty.
>
> **Colleague suspicion**
>
> Leo, who works in the same office but has no personal relationship with Sam, encounters the same rumor at the same time and also forms the belief that Sam has been leaking information.

The difference about what would count as a fitting response to the evidence is striking. Nora and Leo are presented with the same evidence, but they stand in different relations to the truth about whether Sam has been leaking information. For Nora, the truth bears directly on her relationship with a close friend. In light of this, she should be much more cautious in forming a belief on the basis of the evidence she currently has; she should take the question to be settled on that basis. In contrast, in the case Leo, who lacks that personal connection to Sam, it is less clear that the response is unfitting. What is clear is that he is not subject to the same requirements concerning how he treats the evidence as Nora is, and this is explainable in terms of their different relations to the truth in question: Nora is Sam's friend, and this makes the truth about whether Sam has been leaking information matter differently for her than it does for Leo. It gives rise to stricter requirements on how she should treat the evidence than the ones to which Leo is subject.

16. As I've argued above, if we attempt to strip any substantive sense of truth's being valuable and reduce the view to a claim about truth being "epistemically valuable", we cannot do that. I believe that, in this case, Dandelet's conclusion would be warranted.

The same phenomenon is at play in the pair of Performance report cases. Leah (the school board member) and Rachel (the passerby) are both presented with the same report, which contains evidence for the true proposition that students taught with method X have shown a decline in performance. Setting aside their obligation—or lack thereof—to form a doxastic attitude, we can say that if they do form a belief, what would count as a fitting doxastic response differs between them, because the truth's significance differs between them. Leah ought to be cautious and deliberate in assessing the data, to consider how strongly it supports the conclusion, and to attend to factors that might cloud her judgment (suppose, for instance, that she was previously a vocal proponent of method X) in ways that are not required of Rachel. A belief formed hastily, without attention to the reliability of the data, would not be fitting for Leah, even if it happens to be true; it would not be proportioned to the importance that the truth holds for her. By contrast, it seems entirely appropriate for Rachel to believe the report outright at a quick glance.

I admit that it is difficult to specify *exactly* what would constitute a fitting response and what would constitute an unfitting, brash or overly skeptical belief. I nevertheless think that the cases I've presented strongly support the claim that the significance of the truth in question *does make* a difference to what would count as an appropriate belief in light of the evidence.

It's worth noting that, when the truth under consideration matters, there is a strong sense in which agents are responsible for their doxastic attitudes. We hold the agents to be blameworthy for their failures. This is easily explained by the view I am proposing: being attuned to the value of truth, to how and why a particular truth matters, means that agents are disposed to engage in both mental and nonmental actions regarding it: to attend to relevant facts, to suspend judgment, to consider relevant alternatives, to examine potential sources of bias in their assessment of the evidence, to determine whether the evidence as it currently stands warrants belief, and even to seek further evidence when appropriate. In such cases, agents may be required to engage in conscious doxastic deliberation rather than rely on automatic mechanisms of belief-formation and belief-revision.[17]

This view does not rest on an *ad hoc* notion of "fitting responses to value," but on the familiar idea that certain things call for certain responses. The fact that a friend is in a deep depression makes it the case both that I should positively *do* certain things (given our circumstances), and these responses should be proportionate to the seriousness of the situation—relieving them of certain chores, checking in with them often, helping them get professional help, letting them know they can call at any hour of the day or night, and so on. Furthermore, certain ways of responding to this fact would be *unfitting*—for instance, brushing it off or treating it as a mere in-

17. I do not, however, think that the presence of an agential component means that what I call "substantive epistemic normativity" is, in fact, ethics as it concerns belief-formation, belief-revision and inquiry. I discuss why in §5.

convenience. Less serious matters may require less of me and make different responses unfitting: the fact that my friend is currently stuck in traffic trying to cross the BU bridge might warrant mild sympathy or a quick text of commiseration, but it would be unfitting to respond with deep alarm. This kind of phenomenon extends to, and explains, part of the norms of belief-formation and belief-revision to which we are subject.

I want to flag one issue that arises at this point, which concerns what happens, on this account, when a truth is harmful or disvaluable in some other way. Once the full picture of epistemic normativity I develop is laid out, how my account deals with this problem will be clear, but we have to set the problem aside until then.

Thus, my view appeals to a familiar normative phenomenon to explain what we ought to believe: genuine value gives rise to normative requirements. This stands in contrast to accounts that invoke a more mysterious principle—e.g., that "domain-relative value, considered as such, without appeal to whether it is valuable *simpliciter*, places normative demands on us." Moreover, because my account explains both *when* we are subject to obligations to adopt a doxastic attitude and *what* attitudes are fitting to hold, all by appealing to the value of truth, it is explanatorily simpler than two-tiered accounts of epistemic normativity. It is also more powerful: it offers a unified explanation of a wide range of norms that apply to belief—positive and negative norms alike—as well as, potentially, norms of evidence-gathering and norms of inquiry.

3.3.1 *But it's not epistemic!*

One may object that my proposal explains, not epistemic normativity, but moral, social, or prudential normativity as it applies to belief. The idea, I take it, is that I have abandoned the project of explaining what we *epistemically* ought to believe, and instead offered an account of what we ought to believe *all things considered*, or from the standpoint of some other normative domain. On this view, I'm no longer giving an account of epistemic normativity; I'm changing the subject. The worry is that epistemic normativity is supposed to be autonomous: there is, it is thought, a special class of norms that apply to beliefs that don't have anything to do with moral, social or prudential considerations. If we explain norms governing belief by appealing to these other domains, the thought goes, then whatever we are explaining, it isn't *epistemic* normativity.

I must admit that I do not know exactly what is at stake in this objection. What I am trying to account for is the relation between the value of truth and how we should go about forming and revising our beliefs. If that leads us to think that why truth matters has to do with moral, prudential, and social considerations, among others, and if the value of truth, understood in light of these considerations, shapes what we should believe, I see no reason to dismiss the resulting norms as somehow "not epistemic." They are norms that govern what we ought to believe, they take into account our evidence, they give truth pride of place, and they are explained by the fact that truth matters. I am not sure what is gained by carving the normative joints differently, especially since my aim is precisely to capture the idea that our doxastic lives are, by and large, subject to norms that are explained by the value of *truth*.

3.3.2 *But that's just standard pragmatic encroachment! or: moral encroachment!*

How does my view differ from standard pragmatic encroachment about justification?

> Pragmatic encroachment about justification Practical considerations are relevant in determining whether or not the strength of a subject's epistemic state is strong enough to meet the epistemic standards required for justification.[18]

First, the scope of my thesis is more ambitious than that of the pragmatic encroachment thesis: it covers not only what makes a doxastic response appropriate, but also when we are subject to duties to form doxastic attitudes in the first place, and it does so in a unified way. Second, my view is that the value of a truth explains what response to it would be fitting. But the value of a truth differs from stakes as the pragmatic encroacher conceives of them. It is neither reducible to the costs of being wrong, to the benefits of being right, nor to a combination of both. Although the importance of getting it right sometimes has to do with what actions would follow from the belief, sometimes that's simply not the case. A truth may matter in virtue of my relationship to someone, or in virtue of its being meaningful. Return to Friendship suspicion, in which Nora believes that her close friend and colleague Sam has leaked confidential information at work. Even if she is not about to act on that suspicion—she's not confronting Sam, not ending the friendship, etc.—the truth about whether Sam has leaked information still matters in a distinctive way that determines what doxastic attitude is appropriate. It matters simply *because of her relationship* with Sam. The view also doesn't amount to the combination of practical and moral encroachment. Moral encroachment is the view that *moral stakes* (e.g., whether a belief would be racist) are relevant in determining whether or not the strength of a subject's epistemic state is strong enough to meet the epistemic standards required for justification.[19]. Again, as with practical stakes, although the importance of getting it right may depend on moral considerations at play in the agent's circumstances, that need not be the case. Additionally, as I argue below, substantive epistemic normativity is only one type of epistemic normativity, alongside what I call basal epistemic normativity.

18. This definition is adapted from Kim (2017). Often, the view is defended about knowledge rather than justification. See for instance Fantl and McGrath (2009).
19. See for instance Basu (2019).

3.3.3 *It doesn't explain all that falls under the label "epistemic normativity"!*

On even a starkly pluralist account of the value of truth like the one I endorsed, some truths are not valuable. Plausibly, the truth about whether the number of blades of grass on the lawn of the Parliament in Ottawa on the 5th of August of 1991 was odd or even does not matter in any way to any ordinary person. *But*, the objection goes, if these truths don't matter, and if what we ought to believe is explained in terms of fitting responses to the value of truth, then we should not be under any epistemic norms concerning them at all. Yet, there seems to be cases in which truths don't matter, but we are still subject to norms regarding them: if I find myself having any doxastic attitude about the proposition in question, it should be one that accords with my evidence.

There are three main ways, it seems to me, to deal with this objection. The first is to bite the bullet and insist that, despite the appearances, we are not subject to epistemic norms when forming beliefs about valueless truths. This is a hard bullet to bite; we would rejecting the idea the plausible idea that epistemic norms have a universal scope. The second is to argue that, despite the appearances, these truths are valuable in some way. This move looks suspiciously *ad hoc*; it threatens to render the relevant notion of value so expansive that it loses explanatory power—we would lose our grip on how, exactly, the value of truth is supposed to figure in an explanation of epistemic normativity; it would involve something like an appeal to "truth as epistemically valuable", of which, as I have argued, we should be suspicious. The third response, which I think is the most promising, is to concede that the objection gets something right (some truths lack value), but to argue that it simply shows that the account of epistemic normativity given above is incomplete. Although appealing to the value of truth is the best way to explain many epistemic norms (and, arguably, to explain the *most important* of these norms), *not all normative phenomena that are grouped under the label "epistemic normativity" are explained by the value of truth*. Of course, given the ease with which one may simply "patch" an inadequate view by adding further components to it, there is always an air of *ad hoc*ness to a move towards pluralism. My aim is now to present considerations that show this pluralism to be not only plausible, but well-motivated.

3.4 Basal epistemic normativity and cognitive well-functioning

Consider the two following sets of cases, some of which we have already encountered.

> **Group A**
>
> Negligent landlord Paola owns an apartment building. She glosses over the annual fire safety report and fails to form any belief about the safety of the property. In fact, the building is not up to fire safety standard.

Friendship suspicion Nora hears a rumor that her close friend Sam has been secretly leaking confidential information at work. The rumor is based on circumstancial, good, but not conclusive evidence. Despite their long-standing friendship, Nora quickly forms the belief that Sam is guilty.

Friendly detective Detective Dave is investigating a crime scene. Dave is extremely thorough, but, at the same time, a close friend of the butler. Dave finds very strong evidence that the butler did it—the butler's gloves covered in blood, his fingerprints on the murder weapon, a letter written by the butler confessing to the crime—but fails to form the corresponding belief: Dave just can't get himself to believe that his friend would do such a thing. (from Simion 2024)

Group B

Perceptual non-responsiveness Alice is looking straight at the table in front of her and fails to form the belief that there is a table in front of her. (from Simion 2024)

Busy street Despite the loud ambient noise suggesting that the street she is about to cross is a high-traffic road, Elise fails to form the corresponding belief.

Delusion Yanis is convinced that a stranger is impersonating his brother. Despite his brother giving him information that he alone could have, Yanis stay firm in his belief that the person is not his brother, but a stranger.

Although they both involve cases of defective belief-formation, two sets should strike us importantly dissimilar. One key difference lies in the responses they elicit. The agents in **Group A** are apt targets of blame. It seems fitting to sanction them for their doxastic attitudes or lack thereof. As discussed earlier, their failures are not just puzzling or unfortunate mishaps: they are something they are responsible for. This, however, doesn't seem true of the agents in **Group B**, who do not seem to be subject to the same kind of criticism. Their failures strike us as strange or concerning, but not blameworthy in the same way. We might respond with surprise, puzzlement, or some correction that doesn't function as a sanction, but not with blame. *There is* something defective in how they respond to their environments, but the defect is of a different kind. Relatedly, there is a sense in which agents in group A are responsible for their doxastic attitudes in a way that agents in group B aren't. That's because, in **Group B**, what's going on is cognitive malfunctioning; not a failure to engage in conscious doxastic deliberation, to consider relevant alternatives, all things over which an agent has some degree of control; rather, the failure is located at a subagential level.

The two differences between **Group A** and **Group B**—first, with regards to the type of reaction a failure would elicit, and second, whether the failure is attributed to the agent or placed at the subagential level—support the idea that epistemic normativity is at least a bipartite phenomenon. We have already seen that we can explain what goes awry in **Group A** by appeal to fitting responses to the value of truth. How is what goes wrong in **Group B** to be explained?

I think that we can diagnose what is defective in **Group B** by appeal to the idea that there are norms of belief-formation and belief-revision that track the dispositions of a typical, well-functioning cognitive system. As agents, we have some degree of control over our doxastic states—we can engage in conscious deliberation about whether p, consider relevant alternatives and assess whether our evidence rules them out, explore different ways of weighing the support our evidence provides, and so on. We ought also, as we have seen above, be attuned to the importance of certain truths and adopt doxastic attitudes towards them. But a large part of our doxastic life unfolds automatically, without much conscious direction: we register perceptual inputs, update beliefs in light of striking counterevidence, and draw simple inferences without deciding to do so. Failures in **Group B** belong to this largely subagential type of epistemic normativity. Because it is grounded in norms that are basic, minimal, and essential to cognitive good functioning, I call it "basal epistemic normativity". These norms possess the following features:

1. They are not explained by appealing to the value of truth;
2. They are fairly permissive;
3. They are neither mere descriptive regularities nor full-blown, "authoritative" norms, and this is shown in the reactions that are warranted by their violation;
4. They are disparate enough to make it difficult to capture them in a unified way using standard epistemological notions such as reliability or evidential fit.

To get a better grasp on the type of norms to which I am referring, it is worth considering a few additional examples. They are general norms about how a typical, well-functioning cognitive system should operate in ordinary conditions.

(a) Forming true beliefs about specific contents of one's immediate perceptual experience when directly attending to these contents (e.g., believing that there is a screen in front of me, that the lights are on);
(b) Forming true beliefs about suddenly salient features of one's environment (e.g., believing that someone has entered the office after hearing the door open and a voice, having previously been alone);
(c) Believing basic facts about one's immediate environment (e.g., that one is at home, that it is not raining);
(d) Forming true beliefs about what one is presently doing (e.g., that one is walking, writing, or speaking);
(e) Revising an occurrent belief in response to clear perceptual evidence that contradicts it (e.g., realizing that the red object you thought was an apple is a tomato once you pick it up);
(f) Not arbitrarily forming beliefs about things regarding which one has no evidence whatsoever;
(g) Retaining some details, but not all, of what one has just experienced (i.e., a novel one has just finished, a movie one has just watched, what happened in the past few minutes);

(h) Not simultaneously holding beliefs whose contradiction is immediately salient without inference (e.g., not believing both that the key is in your pocket and that you left it on the table).

The norms cited above cover, for instance, awareness of one's environment and actions (a-d), norms prohibiting holding belief without or against one's evidence (e-f), norms on what should be retained from one's recent experiences (g), and norms prohibiting basic and flagrant inconsistencies across beliefs (h). In a sense, they spell out baseline conditions for a minimally well-functioning cognitive system. Although seemingly trivial, norms like these account for a significant part of what we expect when it comes to belief-formation and belief-revision.

While they appear to have to do producing with generally accurate representations of the world, they are not explained, at the most fundamental level, by the value of truth. An explanation of epistemic norms that appeal to the value of truth must rely on a notion of value, *simpliciter*. But basal norms, as exemplified above, are indifferent to the significance of the truths that would be believed if one's doxastic states were to accord with them: they apply even when the resulting beliefs are utterly pointless—such as believing that the lights are on when you look at them. They simply express the minimal standards of belief-formation that are characteristic of ordinary, minimally well-functioning cognitive agents.

Until here, I have not said much about the notion of cognitive well-functioning. One useful way to think about it borrows from Millikan's notion of "proper function." Roughly, for an item *A* to have a function *F* as a proper function, items like *A* must have performed *F* in the past in a way that explains why such items exist now (Millikan 1989). Thus the heart's proper function is to pump blood (and not to produce a rhythmic thump) because pumping blood explains why hearts exist. Similarly, mechanisms of belief-formation and of belief-revision are generally taken to have the proper function of producing true beliefs.[20] True beliefs are typically adaptive, false beliefs maladaptive, and this helps explain why such mechanisms have persisted. Notably, their function is not to produce as many true beliefs as possible—a mechanism with that aim would generate countless trivial beliefs and overburdens our limited cognitive resources; their function is to produce enough true beliefs for a cognitively limited agent to navigate and act effectively in ordinary circumstances.

I draw on this notion of proper function in a limited way. Basal normativity tracks whether, in ordinary, everyday contexts, our mechanisms of belief-formation and belief-revision are meeting *minimal* standards of good functioning, or, in other words, whether they are performing as

20. See for instance Millikan (1984).

they are supposed to (where "supposed to" appeals to the notion of proper function) in ordinary circumstances.[21]

That basal norms track whether our belief-forming mechanisms are functioning as they should explains several features of these norms. First, these norms are fairly permissive. The governing norm is something like "don't be weird"—i.e., don't diverge too much from the ways in which people with ordinarily functioning cognitive systems tend to believe. But the ways in which people with ordinarily functioning cognitive systems tend to believe often countenance considerable variation. For instance, two people with equally well-functioning cognitive system should form some true beliefs about *some* features of their environment, but could permissibly form beliefs about *different* features of that environment. What would be puzzling, and an indicator of cognitive malfunction, would be forming *no beliefs at all* about one's environment. Similarly, two people with a well-functioning cognitive system should, assuming they are attentive to the movie, form some true beliefs about the contents of a movie they just watched; however, the beliefs need not be about the same propositions.

Additional, note that basal norms appear to occupy a middle ground between mere descriptive regularities and "full-blooded", authoritative norms. Violations are noticeable, often puzzling, and may occasionally invite mild correction; however, they are not usually met with a stinging form of blame. There are other norms that occupy this middle space. Compare: we operate with the presumption that people will not sit uncomfortably close to us when ample space to sit elsewhere is available. We rarely notice this norm being in effect while it is being met; it becomes salient only when it is breached. And even then, unless aggravating factors are present—such as warnings or a clear disregard for social cues—the breach is registered less as a blameworthy behavior than as something odd, intrusive, or unsettling. Other examples include norms about speaking turns and interrupting. In much the same way, we would judge it very strange if someone who is otherwise functioning normally were to fail to form a belief about something plainly in front of them—for instance, failing to believe that a person has just walked into the room, despite clearly perceiving it. Blame does not seem to be the proper response to this; nevertheless, it would strike us as off. Just like in the cases above, failures of this kind prompt puzzlement or concern, rather than blame: the agent believes in ways that are at odds with what's typical of cognitively well-functioning agents.

The fact that these norms occupy this "middle ground" between mere descriptive regularities and authoritative norms is not incidental. It is of a kind with and is what we should expect from other norms that are such that deviating from them is an indication of some physiological malfunction. We expect, for instance, ordinary adults to breathe steadily and fairly slowly when

21. Notably, the appeal to Millikan's account of proper functioning does not amount to taking a stance on what the correct account of biological function is.

at rest; it is not merely that it is statistically true that most ordinary adults breathe in this way: it is that they should, under ordinary circumstances. When someone's breathing does not follow this pattern (for instance, when it is unusually fast or appears difficult without any obvious cause), we tend to think that something is awry with their respiratory system. *Blame* is certainly not the reaction warranted by this, however. Concern seems much more apt. Failures to believe in accordance with basal epistemic norms, such as failing to form a belief that a person has entered the room despite clearly perceiving it, prompt the same kinds of responses. They are taken as indications of a an issue in the normal operations of a well-functioning cognitive system, rather than as breaches of obligations in a strong sense.

Finally, appeal to basal normativity as tracking minimal cognitive well-functioning explains neatly the heterogeneity of these norms. By opposition, the norms that served as example above are disparate enough to resist unified explanations that appeal to standard epistemological notions like reliability or evidential fit. Although defending this claim thoroughly is too great a task for the present chapter, I want to point out that it has proven remarkably difficult to give a unified, principled account of what one ought to believe. On the view I'm proposing, this difficulty is partly explained by the fact that what one ought to believe is, to some degree, what an ordinary agent with a well-functioning cognitive system would believe, and that cognitive well-functioning, in turn, involves a patchwork of norms that do not appear to flow from a single fundamental principle such as "believe in accordance with your evidence." Basal norms are heterogeneous: some are positive norms, prescribing that we form doxastic attitudes (e.g., noticing salient features of one's environment and forming the corresponding belief), while others are negative norms, forbidding attitudes that violate them (e.g., simultaneously holding salient, obviously contradictory beliefs).

Consider, for instance, views such as reliabilism, that focus on the etiology of doxastic states to assess whether they are permissible. Such views, by their very nature, cannot account for norms that require one to form a doxastic state in response to their evidence. They can only evaluate doxastic states once they exist. But it seems that many of our expectations regarding how we should believe require people to form doxastic attitudes in certain circumstances; some aspects of one's environment, for instance, should be salient in a way that prompts one to *form* the corresponding doxastic attitude. The "duty", if one wishes to use the term, is a positive one: in response to a certain piece of evidence, one should form a fitting doxastic attitude. Some forms of evidentialism face a similar change: those that countenance only negative duties to believe (e.g., "do not believe against the evidence"), and those on which one must be in a doxastic state to be the proper subject of epistemic evaluation (e.g., "if one is in doxastic state D regarding p, D must fit the evidence that one has regarding p") are unable to account for positive norms.

The solution, it may appear, is simply to modify standard epistemological theories so that they can accommodate the existence of positive norms. The modification cannot be as simple as requiring that one believe what their evidence supports: it would place on us demands that we can't possibly meet, since our evidence supports an infinity of propositions, a phenomenon often called "evidential fecundity" (see Harman (1984) as well as Friedman (2018) for a discussion). This is a choice point for the epistemologist who wants her account to countenance positive duties to believe: *either* she admits that we have all these epistemic duties, but that practical considerations such as the cost of forming these beliefs or the fact that we have to attend to more important things make it permissible for us to violate these duties, *or* she attempts to find a principled way to restrict what we ought to believe on the basis of our evidence.

The first option, although popular, is problematic: it means that ordinary agents are *constantly*, *at all times*, violating epistemic obligations to believe, but they are blameless for doing so because these obligations are generally overridden on the grounds that we are finite agents, that is, agents with limited time, reasoning capacity, and so on. First, it places full compliance with epistemic norms out of reach of cognitively limited agents like us *by principle*. This is implausible: it's one thing for norms to be *very* demanding, it's another for them to be so demanding that there is no way even the most capable of human agent can even come close to satisfying them.

So one has to take the second option and find a principled way to restrict what we ought to believe on the basis of our evidence. This seems, however, an insurmountable task. One would have to explain, using only "epistemic" notions like those of evidential fit, reliability, or being in a position to know, why one should form selective beliefs about *some* of what one's evidence support, but not everything, despite having evidence for much more than these selective beliefs. For instance, we take it that, as they watch a movie, an ordinary believer *should* form *some* selective beliefs about what happens in the movie, but not a near-infinite number of beliefs about every minor detail of the movie. Accounting for the idea that there are positive duties to believe while respecting the idea that such selectivity is perfectly appropriate is an uphill battle.

By contrast, if basal norms are understood as reflecting minimal standards of cognitive well-functioning, their heterogeneity, including the facts that some are positive norms and others negative, is no surprise. Basic cognitive functioning involves a range of distinct requirements: noticing salient features of one's environment, avoiding obvious contradictions, retaining relevant information, and so on. These demands do not reduce to a single principle like "believe in accordance with your evidence." On this view, the disparate character of these norms simply mirrors the variety of conditions that must be met for ordinary belief-formation and belief-revision.

In the previous section, I argued that there is a distinct type of epistemic normativity that reflects minimal standards of cognitive well-functioning. This type of normativity forms one part of the overall picture. In what follows, I turn to that full picture—one that combines these basal norms with the substantive norms grounded in the value of truth—to show how they together form a coherent account of epistemic normativity.

3.5 The final picture: a bipartite account of epistemic normativity

We opened our inquiry with the following question: *what does the value of truth have to do with what we ought to believe?* I have argued that, to answer this question in a satisfying manner, we need a account of the value of truth, but not one on which truth is merely valuable "from the epistemic point of view." A look at several accounts of the value of truth followed. None of the accounts could explain *all* the norms that are subsumed under the term "epistemic normativity." But appealing to the value of truth, I have argued, has powerful explanatory value nonetheless, and we shouldn't take the fact that we cannot appeal to the value of truth to explain each and every epistemic norm to mean that we should reject wholesale the attempt to explain epistemic norms by appeal to the value of truth. Instead, we should embrace a pluralism about epistemic normativity: the view that epistemic norms are explained, at the most fundamental level, by more than one principle.

The pluralist picture I have sketched above is bipartite: it distinguishes between norms for belief-formation and belief-revision that are explained, at the most fundamental level, by appeal to the value of truth ("substantive epistemic normativity", thereafter "substantive normativity"), and those norms that are grounded in what constitutes cognitive well-functioning ("basal epistemic normativity", thereafter "basal normativity"). Besides this, they are distinct in the following respects:

(a) **AGENCY** Substantive normativity has a significant agential aspect: mental and non-mental actions are often required of agents (e.g., directing one's attention towards certain things, considering relevant alternatives, suspending judgment, engaging in conscious doxastic deliberation); in contrast, basal normativity is sub-agential: it concerns processes over which an agent's control is largely indirect;

(b) **RESPONSIBILITY AND BLAME** Relatedly, agents are held responsible for responding fittingly to the value of truth and are liable to blame if they violate substantive norms, whereas violations of basal norms are not blameworthy, but rather puzzling or warranting some corrective.

(c) **NORMATIVE STRENGTH** Substantive norms are full-blooded norms that have the same authoritative character as moral norms (e.g., norms on promise-keeping, etc.); basal norms occupy a middle space between such authoritative norms and mere descriptive regularities (e.g., the norm that the heart *should* pump blood).

The question remains of how these two types of normativity interact, which I now address. On the account I propose, these two types of epistemic normativity *jointly govern* our doxastic life: on the one hand, basal norms are akin to universal side-constraints that set the outer limits on what counts as acceptable belief-formation and belief-revision; any fitting response to the value of truth—with which substantive epistemic normativity is concerned—must take place *within* the space carved by basal norms. Within that space, what's required varies in relation to the value of the truth at issue.

I think that this general conception of the interaction between the two is well-supported by our responses to cases. Consider a public health official who sincerely doubts the safety of routine vaccinations despite having access to a body of evidence that overwhelmingly supports their safety. As long as the doubt is sustained in the face of very strong, yet still defeasible evidence, our natural diagnosis is that the official is violating a substantive norm and is responsible for this violation. In this situation, blame feels entirely apt.

However, as the strength of the evidence increases, and reaches a point where it is, for all intents and purposes, conclusive—when the official has in hand clear, unassailable, decisive data confirming vaccine safety that one would think they should be able to assess—our evaluative stance may shift. Apparence of continued disbelief seems to be either due to insincerity or vicious stubbornness (in which case it involves an agential component and is culpable, qua violation of a substantive norm) *or* starts to look more and more like a breakdown of the basic responsiveness to evidence that any well-functioning cognitive system should display. The more extreme the unresponsiveness to the evidence, the more our reaction moves away from blame and toward puzzlement or concern. At this point, we start to think that a basal norm has been violated and that there is something wrong with the agent's cognitive machinery. The official might need to be reallocated to some other position (or so we would hope), but it doesn't seem, at this point, that it's due to a fault of his.

This suggests that basal norms circumscribe a kind of territory or constitute a backdrop against which substantive normativity applies: we presuppose that the cognitive machinery of an agent is working, in which case the failure is a failure to adopt a fitting response to the value of truth. At some point, however (such as in cases of complete unresponsiveness to conclusive evidence), the cognitive machinery itself appears to be malfunctioning, and the intuition that the agent is blameworthy starts to diminish, even when the truth of the belief matters deeply. The general principle of interaction, as I see it, is along the following lines:

> **GENERAL INTERACTION PRINCIPLE** Fitting responses to the value of truth are limited to those doxastic attitudes that do not violate basal norms; substantive normativity never recommends what basal norms deem defective.

The question of why this principle holds remains, and answering it would require going beyond the scope of this chapter. However, I think its holding supports the view that these two types of normativity—substantive and basal—are both species of one broader genus of normativity, rather than one (substantive) being merely practical or ethical and the other (basal) alone being epistemic. Normative domains typically exhibit internal consistency: we expect morality not to issue conflicting verdicts about what one ought to do, and the same holds for prudence or etiquette. For instance, it should not be the case that φ-ing is both morally required and morally forbidden. By contrast, prescriptions across domains, such as morality and prudence (assuming they are distinct), may very well conflict: it may be that you morally ought not ψ (e.g., lie to your parents), but that you prudentially ought to ψ (say, because you would avoid significant punishment). If substantive and basal epistemic normativity belonged to two different genera of normativity, nothing would prevent substantive norms from requiring or permitting what basal norms forbid. That we do not in fact find such conflicts is evidence that they belong to a single normative domain.

At this point, I want to consider how the account fare when it comes to dealing with some of the issues faced by accounts that have been discussed throughout this thesis; I take the capacity of the account to deal with each of these issues to weight considerably in its favor.

Balancing the universal and the relative aspects of epistemic norms

Remember that one persistent problem for instrumentalist accounts of epistemic normativity, including the social instrumentalist account discussed in the second chapter, is their inability to account for the universality of epistemic norms—the fact that there are norms of belief-formation and belief-revision that apply to agents regardless of their (contingent) personal interests or to what would further the (contingent) ends of the community to which they belong. On social the social instrumentalist account, for instance, what one ought to believe is fully determined by what would further the ends one's community happens to have. Consequently, in cases where an agent is not part of any community, such as in the Desert Island case, the view predicts that no norms of belief-formation or belief-revision apply to the agent at all.[22]

Yet social instrumentalism does capture something important: it easily explains why, in many ordinary contexts, the norms of belief-formation and belief-revision that apply to us seem to depend in part on contingent features such as the social roles we occupy or facts about the communities to which we belong. There is *some* degree of relativity to epistemic norms. The difficulty for social instrumentalism was that it could not balance this relativity with the universality

22. For the Desert Island case, see pp. 53-54.

of epistemic normativity (even though it fares much better than traditional instrumentalism in this regard).

My bipartite account resolves this tension. On the one hand, the value of a truth is often agent-relative: it depends on the broader normative web in which the agent is situated—a web that includes the agent's interests, the social roles they occupy, the obligations they have towards others, and so on. This explains why some substantive epistemic norms vary across contexts—for example, why a juror has especially stringent obligations regarding how she assesses the available evidence in a trial and forms a doxastic attitude about the guilt or innocence of the accused, in contrast with an uninvolved bystander who happens to possess the same evidence.

On the other hand, basal norms are binding universally. They apply to any agent, regardless of contingent practical or social circumstances. Moreover, they cover exactly the kind of failures illustrated in Desert Island, in which an agent forms a belief contrary to conclusive evidence. Even if nothing hinges on the belief, basal norms still apply. We thus have a principled explanation of what is amiss in such cases, an explanation that preserves the genuine universality of epistemic normativity while accommodating some very real degree of relativity.

Unacceptable trade-offs

One objection often lobbed at value-based accounts of epistemic normativity (a version of which I also used against the social instrumentalist account) is that these accounts countenance unacceptable trade-offs between propositions. The general worry is that any view that grounds epistemic norms in the value of truth must accept that these norms sometimes permit, or even require, belief against one's evidence, or belief formed by defective procedures, if doing so leads to more value (e.g., more true beliefs, or true beliefs of great value).

My account does not allow for these problematic trade-offs in which clearly impermissible doxastic conduct is supposedly sanctioned by epistemic norms: the doxastic conduct required in response to the value of a truth cannot conflict with basal norms, which act as side-constraints on fitting responses to this value. Believing against conclusive evidence, on the basis of wishful thinking, or without any evidence, is not permissible, no matter how significant the truth or truths that would come to be believed as a downstream consequence of holding these beliefs. However, *within the parameters set by basal normativity*, some degree of trade-off-like doxastic conduct may turn out to be permissible.

Aronowitz (2021), for instance, has argued that it is sometimes permissible to explore by believing. This is because of the connection between our current beliefs and "our dispositions to conduct experiments and explore the space of possibilities" (for instance, our beliefs constrain our imagination, an important driver for exploration), and because, in turn, of the connection between exploration and gains in accuracy. The possibility of such gains—and hence, the value of exploration—is especially notable at the beginning of inquiry, and diminishes as inquiry progresses. This may warrant adopting the belief that is not optimal in light of one's current evidence for the purpose of exploring and potentially ending up with a more accurate doxastic attitude about the same question in the longer run. I do not want to take a strong stance on the permissibility of exploring by believing. But I think that, as long as it stays within the confines of what is permitted by basal norms, my account could accommodate the permissibility of trade-offs of this kind.

For example, suppose that one is at the beginning of an inquiry into whether a particular defendant is guilty of a serious crime—one for which, if found guilty, the penalty would be life in prison—and one's evidence at this early stage only slightly favors guilt over innocence. The proper way of valuing this truth may be to explore by temporarily adopting the belief that the defendant is innocent, in order to more fully investigate what could serve as exculpatory evidence, at least until the evidential situation becomes more decisive. Because exploring by believing in this way would not involve believing against conclusive evidence or violating other basal constraints, it could very well be a fitting response to the value of the truth in question.

Thus, we see that the account can deal satisfyingly with problems that typically afflict views of epistemic normativity that give value a prominent explanatory role.

Conclusion

A full vindication of this bipartite account would require showing its superiority over other accounts of epistemic normativity, most notably of the constitutivist kind.[23] Such a project is beyond the scope of this chapter. I nevertheless want to consider for a moment some of the account's appealing features.

First, it preserves what was most attractive about veritism: the central role that the value of truth plays in explaining our epistemic norms, while avoiding its shortcomings (such as the retreat to a "merely epistemic" notion of value and the difficulty this creates for explaining what makes a response to the value of truth fitting). On my view, that central role is secured by substantive epistemic normativity. *Second*, it preserves the intuition that epistemic norms are universal—basal norms apply to any agent, regardless of their interests, social roles, or community membership—while accommodating the fact that what one should believe is also sensitive to the agent-relative value of a given truth. Attempts to explain epistemic normativity solely by appeal to the value of truth understood in a substantive sense—such as instrumentalism—struggle to account for this universality. Moreover, value-based accounts are prone to sanctioning impermissible trade-offs,

23. See for instance Nolfi (2015) and Wedgwood (2002).

a difficulty that the present account avoids. *Third*, the existence of each type of normativity is independently motivated. Value gives rise to norms, and proper functioning gives rise to norms; in the present case, both types of normativity govern belief-formation and belief-revision and hence, both properly fall under the label "epistemic normativity."

We often see pluralism as ungainly and cumbersome, an ill-fitting patchwork stuffed with more than it needs and is able to gracefully carry. But pluralism needn't be "We Are the World"; it can be Lennon–McCartney too. And if I had to choose, I'd take the duo over either solo any day. This is one such case. Truly.

References for Chapter 3

Ahlstrom-Vij, Kristoffer (2013). "In Defense of Veritistic Value Monism". In: *Pacific Philosophical Quarterly* 94.1, pp. 19–40.

Alston, William (2006). *Beyond "Justification": Dimensions of Epistemic Evaluation*. Ithaca, New York: Cornell University Press.

Aronowitz, Sara (2021). "Exploring by Believing". In: *The Philosophical Review* 130.3, pp. 339–383.

Basu, Rima (2019). "Radical Moral Encroachment: The Moral Stakes of Racist Beliefs". In: *Philosophical Issues* 29.1, pp. 9–23.

Brady, Michael (2009). "Curiosity and the Value of Truth". In: *Epistemic Value*. Ed. by Adrian Haddock, Alan Millar, and Duncan Pritchard. Oxford University Press, pp. 265–284.

Cohen, Stewart (2016). "Theorizing about the Epistemic". In: *Inquiry* 59.7–8, pp. 839–857.

Dandelet, Sophia (2024). "Epistemic Rationality and the Value of Truth". In: *Philosophical Review* 133.4, pp. 329–365.

Fantl, Jeremy and Matthew McGrath (2009). *Knowledge in an Uncertain World*. Oxford ; New York: Oxford University Press. 251 pp.

Friedman, Jane (2018). "Junk Beliefs and Interest-Driven Epistemology". In: *Philosophy and Phenomenological Research* 97.3, pp. 568–583.

Goldman, Alvin I. (1999). *Knowledge in a Social World*. Oxford: Oxford University Press.

Hannon, Michael and Elise Woodard (forth.). "The Construction of Epistemic Normativity". In: *Philosophical Issues*.

Harman, Gilbert (1984). "Logic and Reasoning". In: *Synthese* 60.1, pp. 107–127.

Kim, Brian (2017). "Pragmatic Encroachment in Epistemology". In: *Philosophy Compass* 12.5.

Kornblith, Hilary (1993). "Epistemic Normativity". In: *Synthese* 94, pp. 357–376.

Kvanvig, Jonathan (2008). "Pointless Truth". In: *Midwest Studies In Philosophy* 32.1, pp. 199–212.

McKay, Ryan T. and Daniel C. Dennett (2009). "The Evolution of Misbelief". In: *Behavioral and Brain Sciences* 32.6, pp. 493–510.

Millikan, Ruth Garrett (1984). *Language, Thought, and Other Biological Categories*. Cambridge, Mass.: MIT Press. 355 pp.

— (1989). "In Defense of Proper Functions". In: *Philosophy of Science* 56.2, pp. 288–302.

Nolfi, Kate (2015). "How to Be a Normativist about the Nature of Belief". In: *Pacific Philosophical Quarterly* 96.2, pp. 181–204.

Pettigrew, Richard (2016). *Accuracy and the Laws of Credence*. Oxford: Oxford University Press. 238 pp.

Pritchard, Duncan (2014). "Truth as the Fundamental Epistemic Good". In: *The Ethics of Belief*. Ed. by Jonathan Matheson and Rico Vitz. Oxford: Oxford University Press, pp. 112–129.

Ramsey, Frank P. (1990). "General Propositions and Causality". In: *Philosophical Papers*. Ed. by David H. Mellor. Cambridge: Cambridge University Press.

Reisner, Andrew (2023). "Welfarist Pluralism: Pluralistic Reasons for Belief and the Value of Truth". In: *Philosophical Topics* 51.2, pp. 53–72.

Simion, Mona (2024). "Resistance to Evidence and the Duty to Believe". In: *Philosophy and Phenomenological Research* 108.1, pp. 203–216.

Sylvan, Kurt (2018). "Veritism Unswamped". In: *Mind* 127.506, pp. 381–435.

— (2020). "An Epistemic Nonconsequentialism". In: *The Philosophical Review* 129.1, pp. 1–51.

Wedgwood, Ralph (2002). "The Aim of Belief". In: *Philosophical Perspectives* 16, pp. 267–297.

Willoughby, James Bernard (2022). "Embedded Epistemic Instrumentalism: An Account of Epistemic Normativity". In: *Philosophers' Imprint* 22.

Zagzebski, Linda (2004). "Epistemic Value and the Primacy of What We Care About". In: *Philosophical Papers* 33.3, pp. 353–377.