

Epistemic Norms and the Normativity of Belief

by

Anna K. Edmonds

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Philosophy)
in the University of Michigan
2018

Doctoral Committee:

Professor Maria Lasonen-Aarnio, Chair
Professor Peter Railton
Associate Professor Chandra Sripada
Professor Brian Weatherson

Anna K. Edmonds

annaedmo@umich.edu

ORCID iD: [0000-0003-1702-1417](https://orcid.org/0000-0003-1702-1417)

© Anna K. Edmonds 2018

Acknowledgements

I came here from two rounds of mountains, Williamstown, Massachusetts and Denver, Colorado, and was dismayed at Michigan's flatness. But Ann Arbor and its people won me over and made itself my home. I've felt nothing but lucky to be immersed in Michigan's fantastic philosophy department (for so long).

My dissertation and I are grateful for the support of the Herbert Boynton Fellowship, a Weinberg Pre-Doctoral Fellowship, a Rackham One-Term Fellowship, and Peter Railton's John H. D'Arms Distinguished Faculty Mentor Fellowship.

Endless thanks to my committee, Maria Lasonen-Aarnio, Peter Railton, Brian Weatherson, Chandra Sripada and my de facto fifth committee member, Daniel Drucker, for the tireless and sharp commentary, the ongoing intellectual inspiration, and definitely not least, for getting me through it. I couldn't imagine a better chair, both academically and personally, than Maria. I am very lucky. The car ride back to Ann Arbor from Purdue was a turning point for my work, and even though I know my driving concerned her, I got only one ticket. Enormous thanks also to my end-of-the-road editorial assistant, Chloe Armstrong, who (of course) did far more than that.

These people shaped my mind and my days: Chloe Armstrong, Sara Aronowitz, David Baker, Gordon Belot, Mara Bollard, Sarah Buss, Mercy Corredor, Daniel Drucker, Betsy Edmonds, Christine Edmonds, Ross Edmonds, Stella Frances, Alex Fraser, Dmitri Gallow, Zoe Johnson King, Jim Joyce, Sydney Keough, Maria Lasonen-Aarnio, Ishani Maitra, David Manley, Peter Railton, Laura Ruetsche, Patrick Shirreff, Cat St. Croix, Rohan Sud, Brian Weatherson, and André Weimerskirch.

This is for Benjamin, a sheer joy who every second is missed, for Henry, who I get to meet very (very) soon, but mainly for André, my everything and my always, who has kept me here, stood me up, made it worth it, and on top of that somehow continued to make it wonderful – and who told me he would still love me even if I never finished my dissertation.

Preface

Early on in grad school I had a few experiences that were formative for this thesis. Their content didn't appear to be *about philosophy* at the time, but they did collectively start to give me an impression of the characters that made up my lot – the lot in which I, of course, was a fairly stereotypical member. On an early phone conversation with my mother, I made some remark about my cousin. It was, as far as I recall, true, but unflattering. My mother responded, “You shouldn't think that about her – she's your cousin!” I remember being struck at the time with the seeming absurdity of that response. Because (somewhat unfortunately for my mother) I am who I am, I asked her to clarify: “Wait, do you mean that I shouldn't say it? You do agree it's true, right?” She demurred about whether it was true – a telling sign, I thought! – and doubled down: “No, I'm saying that you shouldn't have even believed it in the first place!” My mom, I concluded to myself, was evidently deeply confused. Sure, there was no reason to go around saying unflattering things that didn't need to be said, but surely the only thing that determined whether I should *believe* those things was whether or not they were true!

Pub hangouts were populated by groups of grad students from different departments – economists, art historians, political scientists... I noticed early on that the philosophers had a rather distinctive habit that no one else seemed to share. We were prone to correcting people – each other, friends from other departments, ourselves, sometimes all of the above in a single utterance. The striking part was how little the corrections usually mattered. More often than not they were simply distractions from the overall trajectory of some tale or other. Someone would set out to tell a story and a philosopher would interrupt to provide

some correction of a trivial detail. Or sometimes it was a sustained interruption in order to debate the importance of a nearby distinction or the truth of some factual matter. Whatever the variety, the corrections apparently seemed sufficiently fitting or necessary for us to judge them worth it. I started talking about this habit with friends and we gave it a name – fact-a-donning! Fact-a-don strikes again!

Maybe a good explanation of what unites these occurrences is just rudeness or some kind of curious social ineptitude. That is almost certainly true. But it's at least a very particular *breed* of curious ineptitude and one that seems to plague philosophers at rates higher than the general population. And besides, there were other pieces of the puzzle. It didn't take long to realize that a certain way of carrying oneself in seminars or a certain attitude towards one's own work seemed straightforwardly linked to success. The surer one seemed about the quality of their work, the likelier they were to talk about it and write about it and solicit feedback. And, it's not hard to see, the more one produced and the more feedback one got, the better things seemed to go. So even though at the beginning, the confidence and the beliefs in the goodness of one's work in many cases appeared unwarranted, the strategy (if it could even be called that) seemed like a winner. The trick was just how to cultivate those beliefs.

These things together started to shake loose a kind of commitment that I think I had before held pretty deeply. It was a commitment to the value of getting at the truth – wherever and whenever. How important *was* believing the truth? Was it important to set one's beliefs right about insignificant features about one's family members? Or to make sure that everyone had the correct background beliefs when listening to a friend's story, no matter their relation to the narrative arc? I take it the answers to these questions are likely obvious to a lot of people (probably people whose friends like them more). But I think there's something about philosophers that can make it seem to us as if the value of true belief reigns supreme – that if one is in a position in which one's evidence licenses the adoption of a belief, don't suspend or think about something else - go ahead! Get to the bottom of it! Believe it! We will have thereby added a true belief to the stock and

whether or not we choose to act on that belief is a separate matter entirely! That if believing that one's work is good or one is talented in philosophy correlates with eventually becoming a better philosopher, well, *ideally* one would get themselves to do the thing that creates the better results and figure out how to do it without deceiving oneself!

In the *Republic*, Socrates talks about the nature of philosophers. They always love learning, he says, about that which is unchanging and real. "They must be without falsehood – they must refuse to accept what is false, hate it, and have a love of the truth." Is it possible, he asks, "for the same nature to be a philosopher – a lover of wisdom – and a lover of falsehood?" Of course not, agrees Glaucon. Socrates goes on, "Then someone who loves learning must above all strive for every kind of truth from childhood on" (485a-d).

If, by striving for truth and refusing to accept falsehood, what we uncover is that we were wrong – that it's not true that every truth is worth having and that we must hate all falsehoods – then what? Maybe Socrates is right that the *true* philosophers love truth indiscriminately. I, however, difficult as it often feels, am striving for what I see as a better – a lesser – relationship with the truth, and even more popularity in pubs.

Table of Contents

Acknowledgements	ii
Preface	iii
Abstract	vii
Chapter 1 - The Doxastic Account of the Epistemic	1
1 Introduction to the Thesis	1
2 What are Epistemic Norms?	7
3 The Epistemic Constitutivist Response	13
4 The Value of True Belief	27
5 An axial account of the epistemic and the norms of belief	35
References	40
Chapter 2 - Evidential Exclusivity, Correctness, and the Nature of Belief	45
1 A Pair of Claims about Doxastic Deliberation	46
2 Two observations about Evidential Exclusivity and Transparency	49
3 The Standard of Correctness Proposal	57
4 The Standard of Correctness: The Weak Version	63
5 An Alternate Explanation	68
References	74
Chapter 3 - Tradeoffs and Epistemic Value	78
1 The Ethics and Epistemology Analogy	80
2 Tradeoffs: Interpersonal and Inter-Propositional	87
3 Teleology and its Relation to Value	91
4 Epistemic Value and its Relation to Normativity	97
5 Rethinking Epistemic Normativity	104
References	112

Abstract

Epistemologists frequently claim that the question “What should I believe?” demarcates the field of epistemology. This question is then compared to the question asked in ethics: “What should I do?” The question and the ensuing comparison, it is thought, specify both the content and the normativity at stake in epistemology. I argue that both of the assumptions embedded in this demarcation are problematic. By thinking of epistemology’s focal question in this light, first, we risk importing our assumptions about the epistemic domain into our understanding of the nature and normativity of the belief state, and second, we come to have a false picture of the normativity that supposedly underlies the domain.

In Chapter 1, “The Doxastic Account of the Epistemic”, I explore a range of views that assume there to be an essential connection between belief and truth. I look at views that treat all beliefs as attempts to believe the truth, views that consider belief’s biological function to be accurate representation, and views that hold that the very concept of belief is a normative concept. I go on to explore instrumentalist conceptions of belief’s truth connection and conduct an inquiry into the value of true belief. I conclude that neither the value of true belief nor an essential connection between belief and truth can explain epistemic normativity.

In Chapter 2, “Evidential Exclusivity, Correctness, and the Nature of Belief” I note that epistemologists have recently argued that the best explanation for the apparent truth of a pair of claims - “Transparency” and “Exclusivity” - is that belief is subject to a standard of correctness such that a belief that p is correct if and only if p is true. I argue that the proposed explanation unduly privileges one part of

belief's full functional profile – its role in deliberation – and that a more complete consideration of belief's role in cognition suggests an alternative explanation for Exclusivity and Transparency but denies belief's standard of correctness.

In Chapter 3, "Tradeoffs and Epistemic Value", I look at a debate about whether epistemic norms are teleological. Though it's standard to assume in keeping with teleology that certain goals or values explain the content of our norms, a collection of recent papers have aimed to show that this assumption can't be correct because teleological norms countenance tradeoffs but epistemic norms don't countenance tradeoffs. I note that the kind of non-teleological view that countenances no tradeoffs whatsoever is actually quite extreme and virtually unheard of in ethics. I go on to make the case that norms that license no tradeoffs can't reasonably be taken to be grounded in value at all, and thus can't be understood to give rise to necessary normativity. I conclude by suggesting that we broaden our conception of the epistemic domain to recognize teleological norms that provide recommendations for methods of inquiry or pursuit of significant truth or knowledge.

Chapter 1 - The Doxastic Account of the Epistemic

1 Introduction to the Thesis

Epistemologists often identify “What should I believe?” as the question that demarcates the field of epistemology. The question, I think, embeds two assumptions. The first assumption is that *epistemology* can supply the answer to the question of what we should believe. The second assumption is that epistemology is *normative*. Both of these assumptions warrant further scrutiny. Take the first assumption: though some epistemologists go on to clarify the question they’re asking – “What (epistemically) should I believe?”, in many cases the question remains unqualified. Regardless of whether qualifications are made explicit, epistemologists proceed by adopting a certain perspective in answering the question. What characterizes the epistemic perspective? Consider the following cases: Bonjour asks us to

[s]uppose that I have a dear friend who has stood by me and supported me through many trials and crises, often at considerable cost to himself. Now this friend stands accused of a horrible crime, everyone else believes him to be guilty, and there is substantial evidence for this conclusion. Suppose too that I have no independent evidence concerning the matter and also that my friend knows me well enough that an insincere claim to believe in his innocence will surely be detected. If in these difficult circumstances I can bring

myself to believe in his innocence, it is surely plausible to say that there is a sense in which I am justified in so believing; indeed such a belief might well be regarded as obligatory. But the justification in question is plainly not epistemic justification, but rather a kind of moral justification: even if my friend is in fact innocent, I obviously do not know on this basis that he is innocent, no matter how compelling a reason of this sort I may have for my belief. (1985, 6)

A second example, reminiscent of one by William James: Suppose that I am a mountaineer. The snow bridge I traveled on the way to the summit has melted and to safely return home I must leap a crevasse. It's fairly wide, and so my evidence that I will make it is not particularly strong. In order to succeed, I need to feel confident. If I hesitate, all will be lost. Despite the fact that my evidence does not justify my belief that I will clear the gap, I should nevertheless believe that I will.

In the first case, reasons of friendship, perhaps moral reasons, seem to license my belief in the innocence of my friend. But these reasons, Bonjour claims, are not epistemic reasons. He tells us two things about the epistemic perspective: first, "epistemic justification is that species of justification which is appropriate to beliefs or judgments, rather than to actions, or decisions" and second, "[t]he distinguishing characteristic of epistemic justification is...its essential or internal relation to the cognitive goal of truth" (1985, 6). The epistemic reasons in the first case, then, are the reasons exclusively related to the truth of whether or not the friend committed the crime. In the second case, my wellbeing depends on having enough confidence to succeed. So for prudential reasons, we think I should believe that I can leap the crevasse. My epistemic reasons, in contrast, concern only my evidence that I will, in fact, succeed in making a leap of that length.

These cases are typical in their treatment of the epistemic perspective as marked by its connection to truth. But notice Bonjour's other claim: that what makes a species of justification *epistemic* justification is that it's the kind of

justification which is *appropriate to beliefs*. When we combine the thought that epistemic evaluation is that evaluation which is appropriate to beliefs, and the thought that the epistemic perspective is marked by exclusively truth-related concerns, it's easy to see how we begin to develop a picture in which beliefs are essentially evaluated in relation to considerations having to do with their truth.

The second assumption is that the epistemic domain is *normative*. This claim is often made explicitly, as in Alston and Feldman and Talbot, but it's also implicit in the language that epistemologists employ.¹ Epistemologists speak of *justification* and *warrant* and *virtue* and what we *ought* or *should* believe. Bonjour talks of what justification is *appropriate* to beliefs and refers to truth as the *cognitive goal*. This evaluation treats truth as something of *value*, to be pursued or admired or promoted. The normativity of the epistemic domain is often brought into further focus through a comparison with ethics. Some epistemologists refer to epistemology as the "ethics of belief". Though Clifford's 1877 eponymous piece which christened the phrase did in fact hold that it was *morally* wrong to believe on insufficient evidence, the contemporary "ethics of belief" literature captures a wide array of views that attempt to characterize epistemic normativity.

It's also common to analogize the epistemologist's focal question "What should I believe?" to the ethicist's guiding question "What should I do?". Consider the book blurb for Kelly and Pryor's suggested textbook *What Should We Believe?: The Fundamentals of Epistemology*. In introducing the text, they write:

Consider the following question: 'What should I believe?' This question is a normative question. It is, of course, notoriously obscure exactly what one is claiming when one claims that a given question is a normative question. (Indeed, this issue is among those which we will explore.) But intuitively, the question 'What should I believe?' differs

¹ Alston 1988; Feldman 1988; Talbot 2014.

from purely descriptive questions such as 'What do I believe?' or 'What will I believe?' in a way in which it resembles other paradigm normative questions such as 'What should I do?' It is this question which serves as both the starting point and guiding focus of our book, an exploration of select issues in contemporary epistemology...²

In the same vein, Kim claims that “[e]pistemology is a normative discipline as much as, and in the same sense as, normative ethics” (1988, 383). Morality is a paradigm normative domain. We take moral reasons to bestow a kind of necessary force on us. Though we need not think they give us overriding reasons – sometimes we think we have better reason to go on vacation with our families or friends than to volunteer at the local shelter, we typically think they always give us at least some reason to act in accordance with them. The comparison between epistemology’s and ethics’ guiding questions suggests that we view epistemic normativity in the same light: epistemic norms bestow a kind of necessary force that always give us some reason to comply with them.

This thesis is dedicated to examining both of these assumptions: first, that epistemology can tell us what to believe – in other words, that the epistemic perspective, with its commitment to evaluation exclusively in accordance with truth-related considerations, can determine the norms for the state of belief, and second, that epistemic norms are normative in much the same way as ethical norms. In what follows, I will argue that the first assumption serves to collapse the distinction between epistemic norms and norms of belief. It’s to give a *doxastic account of the epistemic*. By conceiving of the norms of the epistemic and the norms of the doxastic as co-extensional, we run the risk of importing our assumptions about the epistemic domain into our understanding of what the belief state is like and what norms govern (or do not govern) it, thereby pre-supposing

² Thomas Kelly and James Pryor: <https://www.amazon.co.uk/What-Should-Believe-Fundamentals-Epistemology/dp/140512363X>. Though it appears the book never quite came to be, its description is sufficient for current purposes.

an answer to what beliefs we should adopt. It's no mere terminological dispute. Though some philosophers appear to unreflectively treat the epistemic domain and the doxastic domain as interchangeable, others do so to signal a deep-seated commitment to understanding the normativity of belief as *essentially epistemic* normativity. Whereas the norms of the epistemic domain are stipulatively truth-directed, an examination of the functional role belief plays in our overall cognition suggests that belief norms shouldn't be thought to be similarly truth-directed, but are perhaps better understood as being oriented toward informing our practical cognition in ways that stably contribute to wellbeing. In respect to the second assumption, I will argue that whether or not epistemic norms can be thought to be normative in a relevantly similar sense as ethical norms – in other words, whether they always give us at least some reason to follow them, depends on how we understand epistemic value. I look at the existing options that attempt to explain epistemic normativity and argue that none of them result in a picture in which epistemic norms bestow necessary normative force the way moral norms do. Here's a more detailed roadmap:

In the remainder of Chapter 1, I will look at how epistemologists characterize epistemic normativity. In particular, I'll take account of the fact that a main data point about epistemic evaluation is that it applies necessarily – in other words, it does not depend on one's individual goals or desires. I'll then ask *in virtue of what* do epistemic norms derive their necessary force? I take there to be two prominent ways of attempting to ground the necessary normative force of epistemic norms. The first is via *epistemic constitutivism* about belief. I will look at the three main varieties of epistemic constitutivism and will argue that none of them vindicate epistemic normativity's necessary force. The second way to ground necessary epistemic normativity is through the value of true belief. I will argue that we don't have good reason to view true belief as instrumentally or finally valuable in a way that guarantees necessary normative force. Where does this leave us in regard to the question about how the normativity of belief relates to epistemic normativity? By eliminating the epistemic constitutivist route, we will

see that there's nothing essential about *belief* or about *believers* that dictates that doxastic norms are necessarily truth-directed. Further, when we examine belief formation, we see that there are many evidentially isolated doxastic processes that systematically contribute to our wellbeing. We should then conclude that the norms of belief are not co-extensive with epistemic norms. This distinction will lead the way to an *axial* account of the epistemic, in which epistemic norms are defined, not as the norms of belief, but in terms of their orientation towards truth or knowledge or other cognitive goods.

Chapters 2 and 3 should be seen as applications of the work of Chapter 1 to current problems in epistemology. In Chapter 2, "Evidential Exclusivity, Correctness, and the Nature of Belief", I set out to look at a pair of related claims – "Exclusivity" and "Transparency" – that have been used to draw conclusions about features of the belief state. Both claims target the way in which only considerations having to do with the truth of a candidate belief are relevant during deliberation about what to believe. I look at the suggestion that the best explanation for the existence of these constraints on deliberation is that beliefs are subject to a standard of correctness such that a belief that *p* is correct if and only if *p* is true. There are a number of interpretations of how we ought to see the normativity at play in belief's standard of correctness. I argue that none of them succeed and that the standard of correctness explanation for Exclusivity and Transparency unduly privileges belief's role in deliberation, but fails to account for other parts of belief's functional profile. Ultimately I reject the standard of correctness proposal and suggest the beginning of an alternate explanation.

Chapter 3, "Tradeoffs and Epistemic Value", takes a look at an ongoing debate about whether epistemic normativity is teleological. Though it's standard to assume that certain goals or values, such as believing the truth and avoiding falsehoods, explain the content of epistemic norms, a few recent papers have drawn attention to the fact that teleological theories of normativity countenance tradeoffs. Since epistemic norms, or at least the norm that tells us to believe the truth during deliberation, appear not to countenance tradeoffs, these views claim

that epistemic normativity can't be teleological. I look in more detail at what it would be for a theory to license no tradeoffs and point out that it would be a quite extreme sort of view – extreme enough that it appears to receive little to no representation in the ethics literature. I argue that norms that never license tradeoffs, even when great value in the domain is at stake, can't reasonably be taken to be grounded in value at all. I go on to claim that this should lead us to understand the normativity that underlies the epistemic domain as in fact very different from moral normativity. Whereas moral norms confer necessary normative force, epistemic norms of the kind that don't license tradeoffs do not. This conclusion paves the way for the suggestion that we expand our conception of epistemic normativity to include norms aimed at inquiry or the pursuit of significant truths.

2 What are Epistemic Norms?

So to start: what are epistemic norms? Philosophers commonly talk about epistemic normativity as the normativity that governs belief. And it's commonly held that epistemic norms are the norms that describe correct belief acquisition and retention. But in beginning an investigation into the source and content of epistemic normativity, it's impossible to ignore the fact that uses of "epistemic" are mostly confined to academic language. And the addition of talk about the epistemic to the academic lexicon appears to be relatively recent. If we consider the epistemic to be a domain of value, we find that "epistemic" is unlike "moral", "prudential", "aesthetic", or other domains of value whose lexical entries are more familiar outside of usage in academic works. We must recognize, then, that "epistemic" (and its related concepts like "epistemic justification" or "epistemic warrant") is a term of art. We look to epistemologists to learn how the concept is deployed, but also must acknowledge that the fact that there exists little for us to grasp onto outside of its academic usage puts more pressure on us to be clear about what work we want it to do in epistemology.

I want to make two central observations about “epistemic” and its use in related concepts. First, as I noted earlier with the example of leaping the crevasse, we apparently have (at least) two competing kinds of reasons. On the one hand, given the fact that I have a strong desire to make the jump and the additional fact that a belief that I’ll succeed raises my likelihood of success, I have a reason to believe I’ll make it. Statistical evidence in the form of, say, past failures of jumps over lesser distances, seems to give me a reason to believe I won’t make it. We can distinguish between two broad kinds of reasons. The first kind we might call a “state-given reason”. It tracks the benefit of being in a state such that I believe that I will make the jump. The second kind is an “object-given reason”. We can call the proposition “I will make the jump” the object of the belief that I will make the jump. Typically we take object-given reasons that p to simply be evidence that p . Where state-given reasons are usually thought to map onto practical reasons, object-given reasons are thought to map onto epistemic reasons.³

State-given reasons have come to be known in the literature as “wrong kinds of reasons” for some kinds of attitudes and object-given reasons have come to be known as “right kinds of reasons”. What makes state-given reasons reasons of the right kind for belief and object-given reasons reasons of the wrong kind? One thing to note is a motivational asymmetry. Whereas it usually feels easy to believe *on the basis of* object-given reasons, it feels much less easy to believe on the basis of state-given reasons. At best, it often feels like we can only make use of indirect strategies to adopt beliefs in accordance with our state-given reasons. Schroeder notes two more reasons for the right/wrong reasons distinction. He claims that, even if we become aware of “pragmatic” reasons, they do not seem to make belief on their basis “more rational qua belief” (2012, 459). He goes on: “So there seems to be a distinctive dimension of rational assessment of beliefs – sometimes called *epistemic rationality* – that is affected by the epistemic reasons of which the subject is aware but not affected by the pragmatic reasons of which the

³ See Parfit 2001 for the distinction, and Schroeder 2012 for discussion of the distinction and how it maps onto the right/wrong kinds of reason divide.

subject is aware” (Ibid.). In addition, a third important difference that makes epistemic reasons right reasons and pragmatic reasons wrong reasons is that “epistemic, but not pragmatic, reasons appear to bear on the correctness of belief. A belief is correct just in case it is true, and epistemic reasons for belief bear on whether that belief is true, but pragmatic reasons are irrelevant to its truth” (Ibid., 460).

In “Evidential Exclusivity, Correctness, and the Nature of Belief,” I try to argue that the motivational asymmetry shouldn’t lead us to treat state-given reasons as wrong kinds of reasons, incorrect reasons, or in any sense lesser reasons for belief. For now I want to note that Schroeder’s second and third reasons for the distinction amount to little other than restatements of the assumption that epistemic reasons are reasons that have to do with the truth of a belief and that only truth-based reasons can be considered reasons of the right kind. When Schroeder says that state-given reasons don’t make belief “more rational *qua* belief” (and that object-given reasons *do* make belief more rational *qua* belief), we can see a direct sense in which truth-directed reasons are *privileged* in determining what we ought to believe.⁴ Why does truth-directedness get privileged in the evaluation of goodness *qua* belief?

We began by noting that epistemic norms are commonly taken to be norms that govern *correct* belief acquisition. This is part of a long and by now familiar practice in philosophy that assumes an essential connection between belief and truth. In the *Nichomachean Ethics*, Aristotle says regarding the intellect that the “good and bad state are truth and falsity respectively (for this is the work of everything intellectual)” (1139a 27-29). We find expression of this essential connection in contemporary epistemologists’ claims that truth is the “aim of belief”, the cognitive or biological function of belief, or the constitutive or

⁴ Though sometimes epistemologists use “rational” to target the idea of “structural rationality”, where being rational consists in ensuring that our attitudes satisfy certain structural or coherence requirements, (see for example John Broome’s 2013 “Rationality through Reasoning”), there’s no *prima facie* reason for us not to interpret “rational” here in keeping with the more familiar notion of “response rationality”, where being rational consists in correctly respond to *whatever* reasons that one has.

conceptual standard of correctness for belief. Later, I will explore each of these options for understanding belief's essential truth connection. But for now, take for examples the following claims:

Williams, on truth as the aim of belief: "It is not a contingent fact that I cannot bring it about, just like that, that I believe something...Why is this? One reason is connected with the characteristic of beliefs that they aim at the truth." (1973, 148)

Gibbard, on belief's standard of correctness: "For belief, correctness is truth. Correct belief is true belief. My belief that snow is white is correct just in case the belief is true, just in case snow is white." (2005, 338)

Grimm, on the primacy of distinctively epistemic evaluation: "It seems clear enough...that even though we can evaluate beliefs relative to countless different 'fundamental values'...the end of realizing the truth enjoys a special sort of status when it comes to the evaluation of belief. (2009, 256)

Velleman, on how norms flow from the very nature or concept of belief: "I take it to be a conceptual truth that beliefs are correct when true and incorrect when false: false beliefs are necessarily faulty or mistaken. What's more, I don't think that the fault in false beliefs can consist in their tendency to misdirect our behavior, and even some false beliefs can direct us well enough. False beliefs are faulty in themselves, antecedently to and independently of any untoward practical consequences." (2000, 277-8)

The thing that these various interpretations of the belief-truth connection have in common is that they take the evaluation of a belief to *necessarily* or *essentially* depend on how the belief satisfies their preferred articulation of the truth connection. Thus, what it is to be an *epistemic* evaluation of a belief as opposed to, say, an evaluation of its moral or prudential value, is to be an evaluation of a belief *qua belief*. We can begin to make sense of such essentialist evaluation by thinking about some artifact. In considering, say, a hole puncher, it's appropriate to make certain evaluations, such as the sharpness of the hole puncher and how well it punches holes, but not others, such as how well it works as a doorstop. The appropriate evaluations seem to be ones that are set by its function or purpose, which flow from the essential nature of the hole puncher. This is the doxastic assumption about the epistemic – that consideration of the normativity of *beliefs in general* is essentially *epistemic* normativity – and all of these authors that make this assumption hold a doxastic account of the epistemic.

The first observation shows the centrality of the truth connection to the epistemic perspective. The second observation I want to make is about the apparent *universal applicability* of the epistemic perspective. Whereas the views above seem to attribute our reason to believe in accord with the evidence to the nature of belief itself, there's another seemingly natural answer to the question about the source of normativity of epistemic reasons. This sort of view explains why facts about evidence ground facts about what we have reason to believe in terms of the relationship between believing in accord with the evidence and achieving the ends or goals that we have. This kind of instrumentalist view sees the epistemic requirement to believe in accord with the evidence as a specific instance of the more general requirement to take the means to our ends. Though in some sense this can seem like a more natural and satisfying explanation of the source of epistemic normativity, we have to take note of an immediate problem the instrumentalist about epistemic reasons seems to face.⁵

⁵ This is not to claim that there can't be an acceptable instrumentalist account of universal reasons. One might claim, as Kornblith attempts in his 1993 "Epistemic Normativity" that everyone has instrumental

Make a distinction between a universal reason and a particular reason. Universal reasons are reasons for all agents, regardless of their own goals and desires. The fact that someone is hurt seems to be a universal reason, because it's a reason for any person to help. Particular reasons, on the other hand, apply only to agents whose interests would be served by compliance. The fact that Christine is in Boston is a particular reason for me to purchase plane tickets to Boston, since it depends on my desire to visit her and doesn't provide such a reason for other agents who don't share my desire. If there's anything that seems to be held fixed in accounts of epistemic normativity, though, it's that epistemic reasons don't depend on anyone's particular interests. Kelly puts this point nicely:

...[F]rom the fact that some subjects are matters of complete indifference to me, it does not follow that I will inevitably lack epistemic reasons for holding beliefs about those subjects. If, despite my utter lack of interest in the question of whether Bertrand Russell was left-handed, I stumble upon strong evidence that he was, then I have strong epistemic reasons to believe that Bertrand Russell was left-handed. Indeed, my epistemic reasons will be no different than they would be if I had acquired the same evidence deliberately, because I did have the goal of finding out whether Russell was left handed. Once I come into possession of evidence which strongly supports that claim that p, then I have epistemic reasons to believe that p, regardless of whether I presently have or previously had the goal of believing the truth about p, or any wider goal which would be better achieved in virtue of my believing the truth about p. (2003, 625)

reason to favor a cognitive system which generates true belief. See also Sharadin's 2016 "Epistemic Instrumentalism and the Reason to Believe in Accord with the Evidence." The challenge, of course, is in making such an account plausible, given the radically varied ends and goals of individual agents. I will raise a criticism about such views later in response to constitutivist accounts that depend on the nature of believers.

Feldman confirms this data point. He takes the following norm to characterize epistemic normativity's fundamental commitment: "[I]f a person is going to adopt any attitude toward a proposition, then that person ought to believe it if his current evidence supports it, disbelieve it if his current evidence is against it, and suspend judgment about it if his evidence is neutral (or close to neutral)" (2000, 679). He claims that this norm applies regardless of a person's particular goal to believe the truth or acquire knowledge. He goes on to say that "all people epistemically ought to follow their evidence, not just those who have adopted some specifically epistemic goals" (Ibid., 682). In the same vein, Sosa writes that the "aptness" of a belief is not proportional to how well the belief furthers our goals of believing truths and not believing falsehoods (2003, 159). Such assertions of the universality of epistemic reasons are very common.⁶

3 The Epistemic Constitutivist Response

What could explain the fact that epistemic norms seem to apply to agents necessarily, regardless of their particular interests or goals? A feature essential or constitutive of belief or believing is a good candidate for such necessity. The idea that beliefs "aim at the truth" has been assumed to be just such a feature. Williams wrote in 1973 that "belief aims at truth" and a large literature has spawned in which philosophers have tried to capture the exact sense in which this slogan seems right.⁷ In this section I'm going to explain the epistemic constitutivist's general strategy and then look at the three main breeds of constitutivist response that appear in the literature.

Constitutivist arguments in general claim that there are certain features of some, for example, attitude or action, that necessarily make it the case that, insofar as we act in that way or have that attitude at all, we are committed to some further

⁶ For other instances, see Zagzebski 2003, David 2001, and Grimm 2008.

⁷ Williams "Deciding to Believe" in *Problems of the Self*.

action, aim, following of principles, etc. In the epistemic constitutivist's case, to claim that "aiming at truth" is constitutive of belief is to claim that part of what it is for some mental attitude to be a belief is that it is directed at or regulated by truth. Some candidate mental attitude, then, will not count as a belief if it isn't aimed at truth in whatever relevant sense the particular constitutivist view requires. If the epistemic constitutivist goes on to claim that believing is a non-optional activity, then it will follow that the aim that believing requires is also non-optional. If in addition, conformance to the non-optional aim requires conformance to some norm, then it will follow that that norm necessarily applies to us.

This is indeed the strategy of epistemic constitutivists.⁸ They claim that belief constitutively aims at the truth: that part of what it is for a mental state to count as a belief is that it is directed towards being true. And, as we saw in the last section, since epistemic norms are norms that describe the truth-directed conditions under which we are permitted, required, forbidden, etc. from believing propositions, it follows that by engaging in belief formation at all, we plausibly are required to conform to epistemic norms. In replying to Gary, an anti-epistemology student who inquires about the possibility of "opting out of the belief business", Railton describes the inescapability of belief: All agents act and form plans. These actions and plans constitutively involve beliefs and are deliberately formed in part on the basis of beliefs. If we were to try to replace belief with some other mental attitude, say, wishing, we wouldn't form any intentions at all. And since our conception of ourselves as agents existing through time involves memories and expectations which themselves involve beliefs, deleting beliefs from our mental stock would leave us with no notion of personal identity:

Being 'in the belief business' therefore isn't as optional as you imagine. It is a precondition of agency. So the argument is non-hypothetical in a familiar sense: as an agent you must possess beliefs;

⁸ See Côté-Bouchard 2016 for a description of the constitutivist strategy and directions to Wiland and Steglich-Petersen for quote selection.

as a belie[f] you must represent certain of your propositional attitudes as accountable to truth and as disciplined by truth-oriented norms...; therefore, as an agent you must so represent at least some of your attitudes, irrespective of what other goals this might or might not serve. A self-representation of certain of one's attitudes as "aiming at" truth is *partially constitutive* of belief, which in turn is *partially constitutive* of agency. Let us, then, call this sort of argument a *constitutive argument*. (1997, 98-99)

So we can't help but be believers. And once we are, we necessarily have mental states that aim toward truth. Wiland writes:

We can extract a norm about how we *should* believe based upon what it *is* to believe. More specifically, the nature of belief tells us something about reasons for belief...The justification for believing the truth doesn't come from something external to the nature of belief itself...Rather, if you are even in the business of believing things, you thereby have reason to believe what's true. Truth is the constitutive aim of belief, and so reasons to believe are necessarily related to considerations concerning the truth of what's believed. (2012, 117-118)

And finally, we move from the claim that some attitude (belief) is inescapable for us, and the claim that some aim (truth) is necessary for that attitude, to the claim that the non-optional activity's non-optional aim requires conformance to a certain norm: Steglich-Petersen tells us that aim theorists explain the norms of "epistemic justification governing belief" – beliefs should be formed in ways that make it likely that they are true – and that a "natural explanation of such epistemic norms is that following them promotes the aims of believing truly...The aim theory promises, in other words, a simple, unified, and prima facie unproblematic explanation of

epistemic normativity” (2009, 396). And thus the epistemic constitutivist has laid the path for an explanation of epistemic normativity that secures the universal applicability that we were looking for. What we’ll look at next are the available suggestions for how to interpret the sense in which belief “aims at the truth”.

We can call the first epistemic constitutivist construal of the way in which beliefs aim at the truth the *Attempt Model*. On the Attempt Model, because of the nature of the state of belief, every time someone believes a particular proposition, they attempt to believe the truth about that proposition. The necessity of epistemic norms is explained by the universality of the claim – every instance of belief formation is an attempt to believe the truth, and the normativity of epistemic norms is explained by the fact that attempts are successes when they accomplish their aim and failures when they fail to accomplish their aim. Sosa has an apparently straightforward interpretation of the way in which beliefs are attempts. He compares beliefs to shots made by an archer. “The good shot is the central value that organizes the sport of archery and the criticism proper to it...Truth is similarly a fundamental value of epistemology. Evaluation is distinctively epistemic when it is concerned with truth” (2007, 77-78).

In a later work he goes on:

The archer’s shot is a good example. The shot aims to hit the target, and its success can be judged by whether it does so or not, by its accuracy...A performance is better than otherwise for not having *failed*, i.e., for not having fallen short of its objective. In line with that, it is *good* if it succeeds, if it reaches its objective...Take any performance with a first-order aim, such as the archery shot and the tennis serve. That performance then has the induced aim of *attaining* its first-order aim... The case of knowledge is just the special case where the performance is cognitive or doxastic. Belief aims at the truth, and is accurate or correct if true. Belief has accordingly the induced aim of *attaining* that objective. (2009, 9)

So the shot hits the target and is successful just in case the belief is true and is a failure just in case it misses the target by being false. Epistemic virtues like true belief and knowledge are *achievements* on Sosa's view. Sosa thinks that shots are evaluated not just on their accuracy, but on the skill that is manifested in bringing about that accuracy. He compares the adroitness necessary for a shot that hits the target because of an archer's skill to knowledge, which he takes to be a more skillful and thus more valuable performance than true belief. But knowledge can manifest success only if the belief reaches its objective and is thus successful. He goes on to refer to his account of epistemic normativity as a sort of "performance normativity" (2009, 9-10).

Two more examples of the Attempt Model of belief's truth aim:⁹ Humberstone writes that unless we take there to be a "criterion of success" in the case of an attitude towards a proposition and "unless that criterion is truth", then the attitude is not a belief. He continues: "So unless the attitude-holder has what we might call a controlling background intention that his or her attitudinizing is successful only if its propositional content is true, then the attitude taken is not that of belief" (1992, 73). And Steglich-Petersen construes the necessary attempts of believers as intentions. He tells us to "Suppose...you judge that John believes that the London train leaves at 5 p.m.. [Y]ou have [necessarily] judged that John has a certain aim or intention in so doing, namely, the aim of regarding that proposition as true only if it is in fact true" (2006, 499).

So on Humberstone's and Steglich-Petersen's view, what makes a certain mental attitude a belief is essentially that we *want* to believe that proposition only if it's true. This makes epistemic norms out to be a variety of instrumental norms – they help us achieve our goal of believing the truth, yet are built into the very nature of belief. This explanatory route, though, is problematic for the goals of the epistemic constitutivist who seeks to account for the necessary normative

⁹ References sourced from Hazlett's "Humean Approaches to the Epistemic Value of True Belief" in his 2013 *A Luxury of the Understanding*.

authority of epistemic norms. It's problematic, I think, for both targets of explanation. It's problematic for explaining the normative authority because we generally don't take satisfaction of an agent's means-ends goals as resulting in necessary normative force. The epistemic constitutivist thinks that our aim of believing the truth supplies us with good reasons to believe the truth. But we don't take aims of, say, torturing, to supply us with necessary *pro tanto* reasons to torture. This is a general problem for reasons internalism, and I take the more basic issue with the Attempt Model to be with the second target of explanation, namely the necessity, or the universality, of epistemic norms.

Proponents of the Attempt Model often construe the belief aim in terms of intention. Attempts, after all, are usually thought to be intentional. Sosa construes beliefs as attempts and successful attempts as achievements. Pritchard writes that "It seems an essential part of achievements that they involve certain motivational states on the part of the subject with regard to the success in question – in particular, that the subject is actively seeking to bring this success about" (2010, 29). But it's implausible that necessarily whenever we form a belief about *p* we intend to believe the truth about *p*. This is doubtful for a number of reasons. First, we often form beliefs unconsciously. Any kind of belief formed unconsciously seems like a belief formed without intention. But further, we form beliefs unconsciously through wishful thinking and other processes based largely on self-deception. It seems even less reasonable to characterize this sort of belief formation as exhibiting an intention to believe the truth. Second, we often form true beliefs *despite* our desire to believe otherwise or despite utter indifference about their truth. Kelly gives the example of a "friend" who blurts out the ending to a movie. Kelly has no desire to learn the truth about the ending and would far rather not have heard the spoiler at all or believe something false about the ending (2003, 626). And Grimm gives an example of absent-mindedly counting the clumps of dust in the waiting room at the dentist. Despite having no curiosity about the correct number of clumps, you come to believe that there are 53 (Grimm 2008, 731 describing a case from Sosa 2003). Additionally, such an explanation

looks even less plausible when we consider that we often take the standards of belief attribution to be less demanding than complex attributions on which one believes that p only if one wants to believe that p only if p . It seems reasonable to think that animals and toddlers believe things, but less reasonable to think that they want to form beliefs only if those beliefs are true.

The Attempt Model needs for it to be the case that *every* belief involves an intention or desire to believe that p only if p . To overcome this challenge, proponents might try to construe the relevant desire or intention as a general or global desire to believe the truth. As Kvanvig puts it: “What we have is a general desire for the truth, and that interest attaches to particular truths in the manner of instantiation in predicate logic. The default position for any truth is that our general interest in the truth applies to it” (2003, 41). The claim that we have a general desire to believe the truth is worrisome, however. If we take a general desire to believe the truth as a single desire to always believe the truth that applies to each possible instantiation, this isn’t right. I know how valuable many kinds of non-truth-oriented belief formation processes are to my wellbeing, and so I for one don’t have the desire to always believe the truth. Hazlett (2013) points out that we might construe a general desire in a second way: a general desire to believe the truth is to generally desire to believe the truth. He compares it to a general desire to eat oysters. The problem, he points out, with this construal is that there are times in which we don’t want to eat oysters, say, when we’re really full. It doesn’t make sense to construe our desire to eat oysters as a *pro tanto* desire because it’s not as if the desire is always there but simply *overridden* in some circumstances. When we’re really full the desire isn’t there at all. So it would be more accurate to construe such a desire as a *prima facie* desire. But epistemic reasons aren’t *prima facie* reasons – if I have conclusive evidence that p , I don’t have a *prima facie* epistemic reason that could be absent in particular cases in which I have no desire to believe the truth. *Prima facie* reasons don’t satisfy the universality requirement on epistemic norms.

The Attempt Model does not look like a promising explanation of epistemic normativity. One might think, though, that a way in which it fails is instructive in pointing towards a different sort of strategy. A primary complaint about the Attempt Model is that it simply isn't true that believers always have intentions or desires to believe the truth. A *Biological Function* model, in contrast, does not require that believers themselves have intentions or desires that their beliefs be true. Instead, the Biological Function model posits the truth aim in belief itself. The biological function of the liver is to filter blood and it does so without any aim or desire or intention on the part of the person in whom the liver is housed. Just as the liver does not function properly if it fails to filter blood, beliefs only function properly if they are true. Velleman introduces an account of belief by saying that "...to believe a proposition is to accept it with the aim of thereby accepting a truth" (2000, 251). Velleman recognizes, though, first, that something can be a belief regardless of the agent's aims or intentions, and second, that there are creatures that have beliefs that don't possess the concept of belief. Thus his account of belief does not require the believer herself to have aims, but rather suggests that the aim belongs to the cognitive system itself. The cognitive system, he says, regulates cognition by "forming, revising, and extinguishing beliefs in response to evidence and argument" (Ibid., 253). It's not just that the cognitive system in fact does this, but importantly that it's part of its *design* to do this: "If a cognitive system isn't regulated by mechanisms designed to track the truth, then it isn't a belief: it's some other kind of cognition. That's why aiming at the truth is constitutive of belief" (Ibid., 17).

The Biological Model's predecessor, of course, is the idea that God designed our cognition to work in a certain way.¹⁰ On this view, truth is the aim of belief because God designed our beliefs to be true. Epistemic normativity thus gains its authority from the goodness of a state's functioning as it was intended to function. The Biological Model is more appealing to naturalistically minded philosophers

¹⁰ See for examples Descartes' Fourth Meditation in *Meditations on First Philosophy*, and Plantinga's proper function argument in *Warrant and Proper Function* (1993).

who worry about the fact that a view that beliefs were designed with an aim requires a designer to have designed that aim. I think, however, that the Biological Model suffers from similar problems to a design view in accounting for its normative authority.

Assumptions that the proper biological function of our mechanisms of belief production is production of true beliefs are very common. Alston writes that “[t]he function of sense perception is to provide us with true beliefs about the immediate physical environment, in the same sense in which the function of the heart is to pump blood around the body” (2005, 33). Quine claims that “creatures inveterately wrong in their inductions have a pathetic but praiseworthy tendency to die before reproducing their kind” (1994, 39). And Dennett writes that “the capacity to believe would have no survival value unless it were a capacity to believe truths” (1971, 101). Goldman’s work depends on the idea that cognition is oriented toward reliability, where reliability is spelled out in terms of production of true beliefs (1986). And Millikan (1993) and Sullivan-Bissett (2017) both make the explicit claim that truth is the biological proper function of belief.¹¹

What is it to ascribe a biological function to our belief-producing mechanisms? Both Millikan and Sullivan-Bissett adopt historical accounts. According to Sullivan-Bissett, it’s that “very roughly, the function of some trait is to do whatever ancestors of that trait did which got them selected. For it to be appropriate to ascribe a biological function to our mechanisms of belief-production, then, they need to have been selected for” (2017, S95). She goes on to note that the claim that true beliefs are adaptive has been “taken to be obvious by many philosophers”, for the reasons that “beliefs which are true are more likely to dispose a creature to act in ways which will satisfy its desires, be they for food, warmth, sex, and so on” (2017, S95). The strategy, then, is to claim that the fact that human ancestors had cognitive faculties that produced true beliefs conferred evolutionary advantage on them and that this fact explains why our beliefs aim at

¹¹ For further representation of views that assume the adaptive value of true belief, see Cowie 2014 and Street 2009.

truth. In order for this explanation to also be an explanation of the normative authority of our epistemic norms, we must assume that having properly functioning beliefs that were in this way selected for is itself good. In what follows, I'm going to describe two problems for the Biological Function account. First, it doesn't seem possible to derive necessary normativity from proper function, and second, we have reason to be worried even if we were to interpret the Biological Function account as grounding epistemic normativity in the idea that the proper function of belief is instrumentally good for us.

To begin, a note about the Biological Function model as a representative of epistemic constitutivism. In order for the Biological Function model to accurately give a constitutivist account of the grounding of epistemic normativity, it must be the case that beliefs have a proper function *and* that this proper function is what provides us with normative epistemic reasons. So the epistemic normativity must stem from the normativity of beliefs functioning as they ought to function. Think back to the beginning comparison with a properly functioning liver. We likely think that we have good reason to have organs that function properly. But our reason for thinking this does not seem to be grounded in the very fact that proper functioning is good. Rather it seems that we want our livers to function properly for instrumental reasons – because properly functioning livers help keep us alive and healthy. But as epistemic constitutivists, we must hold that the normativity is grounded in proper functioning alone, not a further benefit of proper functioning. I will argue that this route can't explain epistemic normativity. In the next section, I will also explore the (more intuitive and probably more widely represented) thought that belief's biological function is *good for* us, though we should understand that such a view is not a constitutivist view.

Embedded in the very concept of “proper” function is that functioning in the way that was intended is good. But we need a real explanation of the ensuing normativity, not just a notion that's definitionally or trivially true. Does a proponent of the organ analogy have an alternate explanation of the normativity of a liver's functioning as it ought to function aside from the health advantages it

confers? One might hold that the concept of a function does in fact entail facts about value or normative reasons. If the function of a liver is to screen blood impurities, then livers that do that are *good qua livers*. Analogously, if the function of belief is to represent accurately, then beliefs that are true are *good qua belief* and beliefs that are false are *bad qua belief*. The problem with this way of thinking is that if oughts that derive from proper function could entail necessary normative authority, then there would be a proliferation of unintuitive reasons to fulfill the purposes of various entities or activities. Torture seems to have a constitutive function – say inflicting severe pain in order to achieve some desired end – and so taking up the role of torturer would suffice to generate reasons for one to cause a lot of pain to one’s victim. But we don’t think that the mere fact that there exists some activity that is governed by an aim gives us good reasons to fulfill the purposes of that aim. Biological function is no different. There exist plenty of examples of biological functions that don’t seem in themselves valuable.

In addition, we should worry about the epistemic constitutivist’s goal to explain epistemic normativity through belief’s *characteristic* or *essential* nature, namely its connection to truth. Hazlett gives the following example (2013, 192). Consider the biological function of sexual intercourse. Plausibly the biological function of sex is reproduction. Having sex was advantageous to our ancestors. We can evaluate tokens of sexual intercourse in light of the standards implied by the biological function, namely whether a particular token results in reproduction. On this standard sex is good when it results in reproduction and bad otherwise. He calls this “reproductive evaluation of sex” (Ibid.). But we might think that though sometimes evaluation of sex via the reproductive standard is appropriate, often there are other standards that we deem more appropriate, such as whether or not the sex is pleasurable, whether or not it constitutes cheating on one’s partner, whether or not it was consensual, etc.. Though I think it’s plausible to think that epistemic evaluation of belief is but one metric along which beliefs can be evaluated, the Biological Model constitutivist sets out to explain epistemic norms through the fact that belief’s function is to be true.

One might object to the comparison of the function of torture and the function of belief. Torture is an activity whose function is generally bad, while one might hold that true beliefs are always either instrumentally or finally good. I think this idea is what theorists often have in mind when they refer to the adaptive benefit of belief formation processes that import true beliefs. Such a move, though, grounds an explanation of epistemic normativity in the value of true belief rather than the constitutive aim of belief. In the next section I will return to this idea. For now, I will turn to the last breed of constitutivist explanation.

The final model of epistemic constitutivism, the *Normativist Model* claims not that believers have any particular aims in believing, or that belief's biological function is to represent truly, but rather that belief is a state that has a standard of correctness - a constitutive norm that governs it. Because of belief's constitutive standard, epistemic norms apply to agents regardless of whether they have any desire or intention to conduct themselves in accordance with that standard.

Wedgwood writes:

It is often claimed that beliefs aim at the truth. Indeed this claim has often been thought to express an essential or constitutive feature of belief. But this claim is obviously not literally true. Beliefs are not little archers armed with little bows and arrows: they do not literally "aim" at anything. The claim must be interpreted as a metaphor. I propose to interpret this claim as a normative claim - roughly, as the claim that a belief is correct if and only if the proposition believed is true. (2002, 267)

And Boghossian:

I would maintain that the holding of this norm [For any p: One ought to believe that p only if p] is one of the defining features of the notion of belief: it's what captures the idea that it is constitutive of belief to

aim at the truth. The truth is what you ought to believe, whether or not you know how to go about it, and whether or not you know if you have attained it. That, in my view, is what makes it the state that it is. (2003, 39)¹²

The normativist interpretation of belief's truth aim is often articulated in terms of a standard of correctness for belief, namely that for any subject and proposition, a belief is correct if and only if it's true. In describing the sense of correctness in play, Wedwood writes:

To say that a mental state is 'correct' is to say that in having that mental state, one has got things 'right'; one's mental state is 'appropriate'. To say that a mental state is 'incorrect' is to say that in having that mental state, one has got things 'wrong' or made a 'mistake'; one's mental state is in a sense 'defective'. (2002, 267)

Normativist constitutivists think that the standard of correctness norm reveals the essential nature of belief, articulating "features that make belief the type of mental state that it is" (Wedgwood 2002, 270). Humberstone writes that "the very concept of belief imports its own criterion of success." (1992, 73) Shah thinks that even "exercising the concept of belief involves accepting that in some sense one *ought* to believe that *p* only if *p* is true", (2003, 449) and Wedgwood later writes that the "very nature" of belief "cannot be explained without mentioning normative properties and relations." (2007, 171)

Whereas the Attempt Model captured belief's truth aim in terms of the believer's attempt to believe the truth, the Normativist Model locates the truth aim in the normativity of belief itself, regardless of the believer's desires. At first this might seem to be a promising way to avoid the main criticism of the Attempt

¹² For other representative views: Shah 2003, Shah and Velleman 2005, and Hieronymi 2006.

Model – namely that because epistemic normativity arises from the satisfaction of agents’ desires to believe the truth, it struggles to account for the universality of epistemic norms. Since the normativist believes that all beliefs are subject to the truth norm there is no similar issue in accounting for universality. However the Attempt Model and the Biological Function Model have an advantage over the Normativist Model. Whereas for the Attempt Model and the Biological Function Model normativity can be reduced to desire satisfaction and evolutionary advantage and thus understood naturalistically, epistemic normativity on the Normativist Model appears to be grounded in a brute irreducible normative fact. When we ask what explains our epistemic reasons to believe the truth, the Normativist’s answer is that it’s an a priori conceptual truth about belief that it’s correct if and only if it’s true. It isn’t any further value conferred upon believers or anything about its biological advantage – the Normativist maintains that epistemic normativity simply flows from an irreducible and constitutive feature of belief.

In addition to the worry that the Normativist Model posits irreducible normative facts, it suffers from the same problem as the Biological Function model – it’s not clear that simply the fact that belief has a constitutive norm gives us reason to think that this norm confers necessary normative authority. There are many norms that arise in relation to the standards set by certain practices. But this doesn’t automatically mean that these norms confer necessary normative authority or reasons for compliance. Again, presumably there are norms that articulate standards for torture that don’t give us good reasons to torture. To put it in different words, in keeping with epistemic constitutivism, it must be the case on the Normativist Model that epistemic normativity arises from the fact that there is a constitutive feature of beliefs such that they are correct if and only if they’re true. But the fact that there is such a norm for belief doesn’t itself mean that there is necessary normative reason to have correct beliefs. Again, the normativist might want to argue that there *is* in fact good reason to have true beliefs because of some further fact about the value of true belief – this is a fact that I will explore in the next section. But that will be to depart from the constitutivist explanation.

I began this inquiry into epistemic constitutivism in order to find a way to make sense of the often-expressed commitment to evaluation of belief in terms of its necessary connection with the truth. The epistemic constitutivists try to express this connection in terms of constitutive facts about the nature of belief or believers and hold that these constitutive facts about belief explain why epistemic evaluation is essential evaluation of belief *qua* belief. I argued that none of the breeds of constitutivist explanation in the literature succeed in accounting for both the universal applicability of epistemic norms and epistemic normativity's supposed necessary normative force that always gives us at least some reason to follow epistemic norms. I turn now to accounts that explain epistemic normativity in terms of the value of true belief.

4 The Value of True Belief

Recall that a plausible fallback for the proponent of the Biological Function Model was to hold, not that epistemic normativity arises from the very fact that beliefs function as they biologically ought to function, but rather that the particular function that beliefs have, namely to represent their contents accurately, is good for us. Though this view is a departure from constitutivism in favor of an instrumental view of epistemic value, and thus is a *prima facie* less likely candidate to account for epistemic norms' universal applicability, it does seem to provide a natural explanation for the normative authority of epistemic norms. In this section I will consider the instrumental and the final value of true belief in turn.

As we saw in exploring the Biological Function Model, it is very common to assume that true belief has value. The question relevant for current purposes is *how* to understand the value of true belief. In order to explain what grounds our epistemic normativity, we must be able to explain how epistemic norms - like moral norms, our beginning comparison suggests - give rise to necessary

normative force. In other words, we will try to determine whether the value of true belief dictates that epistemic norms always give us at least some reason to follow them.

To begin, we can acknowledge that true beliefs are very often valuable. They help us get to the closest local bike shop, locate the wine opener, and satisfy our curiosity about Æthelstan. Sometimes, though, - like when we're leaping crevasses - it seems like it would clearly be better for us to have a false belief. So in precisely what way is true belief valuable? Lynch claims that (in addition to true belief's being intrinsically good, a view that we will visit later) true belief is *prima facie* good (2004, 46-47). This doesn't say much, though. *Prima facie* value is value "at first sight" - the implication, of course, that when we look closer we might see that that value doesn't obtain at all. Lynch also suggests that we can alternatively understand true belief's value as value "other things being equal" (Ibid., 47). Hazlett makes a point about the weakness of claims that X is better "other things being equal" (2013, 10). He compares a defense of the value of taking cocaine on the grounds that, other things being equal, it's always better to take cocaine. One can't object to this claim, he says, by pointing out that cocaine is unhealthy or addictive, because the claim is that cocaine is good *other things being equal*. The claim is supposed to tell us that for two people who are the same in terms of health and addiction, that *then* the one who takes cocaine is better off. But clearly this is not an adequate defense of the value of cocaine. Other things may rarely (or never) be equal. Lynch writes that

one thing people never tire of reminding me of whenever I claim that truth and its pursuit are good is that there are many times when they aren't. The truth can hurt and falsity can have its own rewards. But these exceptions hardly show that true belief isn't good and therefore not worth striving for. They are the exceptions that prove the rule. (2004, 46)

Foley writes that it's "highly improbable that this [false belief turns out to be practically rational] will happen frequently", that it will happen in unusual, but not

in “relatively normal situations” (1987, 224), and that “[t]he evidentialist is right to insist that ordinarily it is deeply irrational to undertake a project of worsening our epistemic situation in hopes of securing a belief that will generate pragmatic or long-term intellectual benefits” (1992, 27). Claims that true belief is usually or typically better for us abound.

It’s hard to know exactly what to make of these claims. On interpreting the claims that X is “normally” or “generally” or “ordinarily” or “typically” better, Hazlett points out that our background theoretical or normative commitments routinely affect our judgments about what is normal or typical (2013, 11). And it won’t suffice to defend a claim that X is “normally” better to show that in a statistical majority of cases X is better (Ibid., 13). Consider driving on the right side of the road. Though in a majority of cases, driving on the right side of the road is the better thing to do to avoid a crash, it doesn’t seem right to say that “normally, driving on the right side of the road is better.” Rather, *sometimes*, like when one’s in the UK, it’s better to drive on the left. That situation isn’t *abnormal* – after all, it happens in regular and predictable circumstances. Swerving into the left lane to avoid an animal is abnormal, but driving on the left in the UK is not.

Let’s consider, then, the claim that for any subject and any proposition, true belief is *normally* better than false belief. We saw that identifying regular and predictable patterns of driving on the left is sufficient for the falsity of the claim that driving on the right is normally better than driving on the left. Similarly, to determine whether true belief is normally better than false belief, we should look for clear and predictable patterns of belief formation in which it’s better for a subject to have false beliefs. In “Evidential Exclusivity, Correctness, and the Nature of Belief” I try to show the existence of multiple kinds of doxastic processes that clearly and predictably produce beneficial false beliefs. Here I will summarize those findings and suggest a few more.

Take a doxastic process to be a way of forming, maintaining, and revising one’s beliefs. Research in cognitive psychology makes clear that, though doxastic processes often seem oriented toward reflecting one’s environment accurately,

there are also many doxastic processes that are best seen as enhancing agential wellbeing. It seems likely that when epistemic constitutivists conclude that belief has built-in features that make it necessarily either conceptually or functionally truth-aimed, they might be concentrating on aspects of what it's like to deliberate about the truth about a proposition. But belief formation happens non-deliberatively too – automatically, subconsciously, and free from the apparent constraints of active deliberation – and so there's no reason to think that what often happens in non-deliberative contexts is any less indicative of the true nature of belief. In non-deliberative contexts, doxastic processes produce beliefs that cause agents to view themselves in unrealistically positive terms, think that they have more causal control over their environments than they actually have, and hold overly optimistic views about the future.¹³ Cognitive filters of selective attention, representation, and recall interact with our doxastic processes to influence our beliefs about ourselves and the trajectory of our lives. We can maintain these overly positive beliefs via selective inquiry and selective attention to evidence (Buunk et al. 1990; Sedikides and Gregg 2008).

In addition to overly positive beliefs about ourselves and our lives, we manifest other sorts of doxastic biases that seem to have positive effects on our wellbeing. Stroud, for example, argues that friendship constitutively involves epistemic partiality (2006, 499). She writes that “friendship places demands not just on our feelings or our motivations, but on our beliefs and our methods of forming beliefs” and that “this epistemic partiality is contrary to the standards of epistemic responsibility and justification held up by mainstream epistemological theories” (Ibid.). When we hear unfavorable things said about a friend, she thinks that it's incumbent on us, qua friend, not just to defend our friend outwardly, but also inwardly in terms of our beliefs. So friendship, she thinks

¹³ Tiberius 2008, summarizing Taylor and Brown (1994, 21).

yields different doxastic outcomes: different beliefs. I want to claim that not just the cognitive procedures we deploy in assessing evidence but the set of beliefs we end up with will systematically differ when the subject of the story is our friend, and in particular that the latter will be more favorable than they otherwise would be. This is the second locus of the differential epistemic practices connected with friendship: where our friends are concerned, we draw different conclusions and make different inferences than we otherwise would (or than a detached observer would). I noted above that we spend more energy generating, that is, coming up with, alternative—and less damning explanations of the reported conduct when the story we're told is about a friend. But that isn't all, for we are also likely to give such alternative constructions greater credence than we would for a nonfriend. And at the end of the day we are simply less likely to conclude that our friend acted disreputably, or that he is a bad person, than we would be in the case of a nonfriend. Friendship seems to alter not just the procedures we use to process new information but the conclusions we end up drawing. (2006, 506)

Stroud thinks, then, that friendship requires a biased doxastic process – a process that systematically produces false beliefs that help sustain it. Though she argues that we owe our friends a kind of epistemic partiality beyond that to which we owe strangers, we might think that in fact we would be better off if our thinking in regard to the rest of the population also showed a kind of biased deference. Pace writes that

[t]hinking charitably of others...may in fact be a prima facie moral obligation...regarding all people. The disposition to think charitably of others is often considered a moral and intellectual virtue. Doing so arguably involves requiring less evidence to think well of them

and requiring more evidence before one thinks ill of them. (2011, 258)

It's less straightforward that being widely charitable towards others has a direct impact on our wellbeing in the way that friendship does. Whereas people often think that having friends is partially constitutive of living well, it might be more reasonable to argue for the value of believing charitably in a less direct fashion: it does seem, however, that thinking charitably of others is morally virtuous and it's common to think that moral virtue is partially constitutive of wellbeing.

I've suggested that there's a wide range of kinds of doxastic process that produce false beliefs. Further, these doxastic processes have a host of effects on our wellbeing. Unrealistically positive beliefs about ourselves, inflated beliefs about our capacity to control our environment, optimism about the future, and having close friendships are all significantly correlated with not being depressed. Non-depression, in turn, is linked to higher subjective wellbeing, greater motivation, and superior coping abilities. Collectively, this gives us strong evidence that there's a clear pattern of non-truth-directed belief formation that makes our lives go better.

It looks then, like believing truly can't be thought to *always* or *normally* be better for us than believing falsely. One might think, though, that to focus exclusively on the instrumental value of true belief is to ignore another way that true belief might be valuable. It's often claimed that true belief has final value. We should interpret this claim to mean that truth always confers value on our beliefs. This value is to be understood not in terms of some further benefit that makes a belief better for us in virtue of being true, like making it more useful, but rather in terms of fundamental final value. Though it's notoriously difficult to argue that some candidate value does or doesn't have final value, I will try to argue that we have very good reason to think that it does not.

In order to make sure that the intuitions we elicit are actually about the final value of true belief and not the instrumental value of true belief, it's best to

focus the investigation on pointless true beliefs. Pointless beliefs are beliefs that have no impact on our lives. A common example is the number of grains of sand in some stretch of desert or whether the number of blades of grass in some stretch of field is even or odd. We have nothing to gain or lose by having a true belief, a false belief, or no belief at all about the number of grains of sand in a quadrant of desert. If we think we have good reason to believe that it's not any better to have a true belief than a false belief or no belief about the number of grains of sand, then we should conclude that true belief likely does not have final value after all.

Kvanvig argues that cases like this are misleading. Because we are not gods with limitless resources, we have reasons not to expend energy finding out the truth about the number of grains of sand. But this is not to say that believing the truth would not be good. He appeals to our intuitions in cases in which resources are not limited:

We should ask ourselves, regarding possible individuals in such a cost-free environment, what the cognitive ideal would involve. Here [supporters of the view that truth has final value] have millennia of theological reflection on their side. Part of the cognitive ideal, whatever else it may involve, is knowledge of all truths; omniscience for short. But for omniscience to be part of the ideal, no truth can be pointless enough to play no role at all in the story of what it takes to be cognitively ideal. (2008, 209-210)

On Kvanvig's view, then, though acquiring the truths might not be worth it, omniscience realizes a sort of cognitive ideal. It would be better, he thinks, to have every true belief than to have only the true beliefs that are not pointless. I don't share the view that having every true belief is better than having only important true beliefs. But that's not an easy debate to have. What we can extract from Kvanvig's view is two thoughts that are worth pursuing. First, in cases in which we have no other interests or needs we should be motivated to learn the truth about

pointless propositions, and second, that there will be plenty of cases in which the value of true belief is overridden by other things that are valuable. Wrenn, in discussing this first thought, writes

Truth...is so good we should want it when it is 'cost-free.' This is the merest quantum of value. It is distinguishable from neutrality only in the fantastical case of someone without cognitive limitations, without other concerns, interests or needs, or who somehow gains benefits from inquiry that precisely offset the costs involved in pursuing it. We are not like that. We have cognitive limitations. We have concerns and interests other than truth, and we have needs that can conflict with the conduct of inquiry into all questions. If truth's intrinsic value is invisible except when nothing else is at stake, then truth's intrinsic value might as well *always* be invisible, and we have no reason to think truth is more important than anything else we value. (2017, 114)

Further, even in cases in which there really are no competing concerns, I don't share the intuition that one ought to be motivated to learn the truth. It seems that true beliefs are better thought of as something that is permissible to value, but not necessary to value. In cases in which someone with no competing concerns were unwilling to learn a pointless truth, I see no reason to claim that they are undervaluing truth. This suggests that there is no degree to which we *ought* to value it.

On the second thought, Kvanvig acknowledges that there are many cases in which we will not pursue true beliefs because of competing values. It's worth thinking about how true belief does stack up against other values. Consider other candidate final values, such as moral value or pleasure or beauty. Though it's impossible to strictly quantify exchange rates between values, it looks like pointless true beliefs don't fare very well in comparisons. It seems straightforwardly irrational to prefer even a lot of pointless true beliefs over really

any amount of moral value. I would definitely choose anyone's not suffering from a five minute headache over one thousand true beliefs about grains of sand in desert quadrants. The same goes for units of pleasure or beauty value. It seems difficult to claim that someone who preferred any other final value to huge amounts of pointless true belief has done anything improper in their valuing. Though I began with a concession as to the difficulty of determining what is or is not a final value, I take these to be strong reasons not to view true belief as a final value.

5 An axial account of the epistemic and the norms of belief

I've embarked on an investigation into epistemic constitutivism and the value of true belief in order to better understand the source of epistemic normativity. What I tried to show in my examination of epistemic constitutivism is that each type of constitutivist view holds that epistemic evaluation of beliefs is *essential* evaluation of belief – in other words evaluation in accordance with belief's essence. This is the Doxastic Account of the Epistemic. It makes "epistemic" out to be about belief, and it holds the belief state to be properly evaluated in accordance with its nature, as a state that is aimed in some way at the truth. The Attempt Model holds that in believing, I necessarily attempt to believe the truth, so essential evaluation of belief is evaluation in accordance with the truth norm. Since the Biological Function model holds that truth is the biological proper function of belief it too holds that essential evaluation of belief is evaluation with respect to the truth norm. And the Normativist Model builds the truth norm into the very concept of belief, thus making epistemic norm-directedness part of what it is to be a belief.

We began by noting that epistemologists often claim that epistemology addresses the question of what we should believe and that epistemologists take this question to be a normative question in much the same sense as the ethicist's normative question about what we ought to do. The "epistemic perspective" that

guides epistemology, however, is a perspective that is marked by its direction toward truth. That gives us reason to resist thinking of epistemology's guiding question as "What should I believe?". For we have seen that although the epistemic perspective privileges a truth-directed answer to that question, neither epistemic constitutivist arguments nor arguments about the instrumental or final value of true belief should lead us to conclude that belief norms are similarly truth-directed.

Neither should we claim that epistemic norms are normative in the very same sense as moral norms. Whereas we can see moral norms as making necessary claims on us – as generating necessary normative force and always providing us with at least some reason to follow them – the views of epistemic normativity that we have investigated fail to account for epistemic norms' necessary normative force. The Attempt Model does not give us a normative reason, for each proposition, to believe the truth, because it's not the case that every belief is an attempt by the believer to believe the truth. The Biological Function Model does not give us a normative reason, for each proposition, to believe the truth because proper functions don't deliver necessary normative force. And the Normativist Model also fails to necessarily provide us with reasons because "correctness qua belief" does not produce necessary normative force either. Neither, I've tried to argue, is it the case that the instrumental or final value of true belief always provides normative force and always provides us with some reason to follow epistemic norms.

I've argued that we should resist both commonly made assumptions about the epistemic domain. Where does this leave us with respect to understanding the nature of belief and what *does* properly demarcate the epistemic domain? The Doxastic Account of the Epistemic defines the epistemic in terms of belief. I've tried to suggest that such a conception is very common in epistemology. We saw it in play in Schroeder's distinction between epistemic reasons, which he treated as reasons for belief, and practical reasons, which were treated as reasons for action. We should resist such a conception because it bars the very possibility of non-

truth-directed answers to the question of what we should believe. It also serves to bar other conceptions of what can be a proper subject of epistemic reasons. On the Doxastic Account of the Epistemic, epistemic reasons just are reasons for belief, so we are less inclined to think broadly about epistemic reasons for mental attitudes other than belief or epistemic reasons for actions.

A better alternative, I think is an axial account of the epistemic. Such an account understands the epistemic domain in terms of value. We think of truth and knowledge and understanding as valuable epistemic goods. On this conception, epistemic norms are directed toward the acquisition of such goods. An axial account does not posit a fundamental contrast between the epistemic and the practical so as to allow for epistemic reasons for action – “How should we conduct our inquiry?” and practical reasons for belief – “What would it be better for me to believe about the cause of my son’s death?” Instead, the epistemic domain is more informatively contrasted with other domains of value, like the moral or aesthetic domains.

Our inquiry into the value of true belief highlighted the existence of many types of doxastic process that, when they are functioning in ways that reliably contribute to agential wellbeing, predictably produce false beliefs. Though we may well ultimately find it better to eschew all talk of the “aim of belief”, we can at least recognize that these processes don’t seem to be truth-aimed. The correlation between these doxastic processes and significant effects on motivation and subjective wellbeing suggest that we must abandon any conception of the normativity of belief that holds that necessarily or constitutively beliefs are to be categorized or evaluated exclusively in accordance with a truth connection.

What, then, should we take to be normative for the state of belief? In one sense, I think recognizing belief’s full functional role can help steer us away from incorrect views of belief’s normativity. Though a neat account of the necessary and sufficient conditions of belief likely remains outside our grasp, we can at least identify a collection of central functional properties. These properties include belief’s inferential relationship to other beliefs, its relationship to perception, its

role in assertion, and importantly, I think, its role in motivation. I take William James' idea that *thinking is designed for doing* to be a helpful backdrop in constructing a theory of belief, and I think part of what cognitive psychology has shown us is that non-truth-aimed processes very often help us *do* better. This is in large part due to their role in motivation. By seeing ourselves in more positive lights, thinking that we have greater causal impact on our trajectories, and remaining optimistic about the future – despite the fact that such positive beliefs are often unwarranted - we are motivated to remain engaged in our lives and with our friends. This, then, could lead us to an account of belief's normativity that is directed not toward truth, but toward informing our practical cognition in ways that stably contribute to wellbeing. I think such a conception is helpful in that it highlights a contrast with the received view that beliefs are necessarily truth-aimed. On the other hand, we might side with David Papineau who thinks that there is no distinctive species of normativity that attaches to belief. I think in some sense this view is also importantly right. He writes

I do not of course wish to deny that there are some valid prescriptions that apply to the adoption of beliefs. However...these are always prescriptions of a kind that arise in connection with other human activities as well as belief-formation. More specifically,...all such prescriptions arise from considerations of moral value, or personal value, or possibly aesthetic value, and not from any distinct species of doxastic value. (2013, 64)

So although understanding the belief state in both theoretical and practical contexts might require that we actively attempt to shift our understanding of its normativity away from truth exclusivity and towards its contribution to practical cognition, it also seems true that it might not always be helpful to consider a domain of uniquely doxastic value. In this sense our theorizing about belief will mirror our theorizing about action. Though we

often want to ask about norms of action in the moral domain or prudential domain and give answers that depend on moral or prudential value, we don't often think that such answers will be captured in terms of a unique domain of "action value". The domain of belief seems similar. Regardless of whether we opt for a specific characterization of belief norms, I hope to have shown the problematic nature of the Doxastic Account of the Epistemic. Epistemic norms are not the norms of belief, and belief norms cannot be taken to be necessarily truth-directed.

References

- Alston, William. 2005. *Beyond Justification: Dimensions of Epistemic Evaluation*. Ithaca: Cornell University Press.
- . 1988. "The Deontological Conception of Epistemic Justification." *Philosophical Perspectives* 2: 257-99.
- Boghossian, Paul A. 2003. "The Normativity of Content." *Philosophical Issues* 13 (1): 31-45.
- BonJour, Laurence. 1985. *The Structure of Empirical Knowledge*. Harvard University Press.
- Broome, John. 2013. *Rationality Through Reasoning*. Wiley-Blackwell.
- Buunk, B.P., R.L. Collins, S.E. Taylor, N.W. VanYperen, and G.A. Dakof. 1990. "The Affective Consequences of Social Comparison: Either Direction has its Ups and Downs." *Journal of Personality and Social Psychology* 59 (6): 1238-49.
- Côté-Bouchard, Charles. 2016. "Can the Aim of Belief Ground Epistemic Normativity?" *Philosophical Studies*, 173 (12): 3181-98.
- Cowie, Christopher. 2014. "In Defence of Instrumentalism about Epistemic Normativity." *Synthese* 191 (16): 4003-17.
- David, Marian. 2001. "Truth as the Epistemic Goal." In *Knowledge, Truth, and Duty*, edited by M. Steup, 151-169. New York NY: Oxford University Press.
- Dennett, Daniel C. 1971. "Intentional Systems." *Journal of Philosophy* 68 (February): 87-106.
- Descartes, R. and Cress, D. (1993). *Meditations on First Philosophy*. Indianapolis: Hackett Pub. Co.
- Feldman, Richard. 2000. "The Ethics of Belief." *Philosophy and Phenomenological Research* 60 (3): 667-95.
- . 1988. "Epistemic Obligations." *Philosophical Perspectives* 2: 235-56.

- Foley, Richard. 1992. *Working Without a Net: A Study of Egocentric Epistemology*. New York NY: Oxford University Press.
- . 1987. *The Theory of Epistemic Rationality*. Cambridge, MA: Harvard University Press.
- Gibbard, Allan. 2005. "Truth and Correct Belief." *Philosophical Issues* 15 (1): 338-50.
- Goldman, Alvin I. 1986. *Epistemology and Cognition*. Harvard University Press.
- Grimm, Stephen R. 2009. "Epistemic Normativity." In *Epistemic Value*, edited by Adrian Haddock, Alan Millar and Duncan Pritchard, 243-264. Oxford: Oxford University Press.
- . 2008. "Epistemic Goals and Epistemic Values." *Philosophy and Phenomenological Research* 77 (3): 725-44.
- Hazlett, Allan. 2013. *A Luxury of the Understanding: On the Value of True Belief*. Oxford University Press.
- Hieronymi, Pamela. 2006. "Controlling Attitudes." *Pacific Philosophical Quarterly* 87 (1): 45-74.
- Humberstone, I. L. 1992. "Direction of Fit." *Mind* 101 (401): 59-83.
- Kelly, Thomas. 2003. "Epistemic Rationality as Instrumental Rationality: A Critique." *Philosophy and Phenomenological Research* 66 (3): 612-40.
- Kelly, Thomas, and Pryor, James. *What Should We Believe? The Fundamentals of Epistemology*. 2018. Available from <https://www.amazon.co.uk/What-Should-Believe-Fundamentals-Epistemology/dp/140512363X>.
- Kim, Jaegwon. 1988. "What is Naturalized Epistemology?" *Philosophical Perspectives* 2: 381-405.
- Kornblith, Hilary. 1993. "Epistemic Normativity." *Synthese* 94 (3): 357-76.
- Kvanvig, Jonathan L. 2008. "Pointless Truth." *Midwest Studies in Philosophy* 32 (1): 199-212.
- . 2003. *The Value of Knowledge and the Pursuit of Understanding*. Cambridge University Press.
- Lynch, Michael. 2004. *True to Life: Why Truth Matters*. Cambridge MA: MIT Press.

- Millikan, Ruth G. 1993. "Explanation in Biopsychology." In *White Queen Psychology and Other Essays for Alice*, 171-192. Cambridge MA: MIT Press.
- Pace, Michael. 2011. "The Epistemic Value of Moral Considerations: Justification, Moral Encroachment, and James' 'Will to Believe.'" *Noûs* 45 (2): 239-268.
- Papineau, David. 2013. "There Are No Norms of Belief." In *The Aim of Belief* edited by Timothy Chan, 32-63. Oxford University Press.
- Parfit, Derek. 2001. "Rationality and Reasons." In *Exploring Practical Philosophy: From Action to Values*, edited by Dan Egonsson, and Ingmar Persson, 17-39. Burlington, VT: Ashgate.
- Plantinga, Alvin. 1993. *Warrant and Proper Function*. Oxford University Press.
- Pritchard, Duncan. 2010. "Knowledge and Understanding." In *The Nature and Value of Knowledge*, edited by Duncan Pritchard, Alan Millar and Adrian Haddock, 1-88. Oxford: Oxford University Press.
- Quine, W. V. O. 1994. "Natural Kinds." In *Naturalizing Epistemology*, 2nd ed., edited by Hilary Kornblith, 57-76. Cambridge, MA: Bradford of MIT press.
- Railton, Peter. 1997. "On the Hypothetical and Non-Hypothetical in Reasoning about Belief and Action." In *Ethics and Practical Reason*, edited by Garrett Cullity, and Berys Nigel Gaut, 53-79. Oxford University Press.
- Schroeder, Mark. 2012. "The Ubiquity of State-Given Reasons." *Ethics* 122 (3): 457-88.
- Sedikides, Constantine, and Aiden P. Gregg. 2008. "Self-Enhancement: Food for Thought." *Perspectives on Psychological Science* 3: 102-116.
- Shah, Nishi. 2003. "How Truth Governs Belief." *Philosophical Review* 112 (4): 447-82.
- Shah, Nishi, and J. David Velleman. 2005. "Doxastic Deliberation." *Philosophical Review* 114 (4): 497-534.
- Sharadin, Nathaniel. 2016. "Epistemic Instrumentalism and the Reason to Believe in Accord with the Evidence." *Synthese*: 1-19.
- Sosa, Ernest. 2009. "Knowing Full Well: The Normativity of Beliefs as Performances." *Philosophical Studies* 142 (1): 5-15.

- . 2007. *A Virtue Epistemology: Apt Belief and Reflective Knowledge, Volume I*. Oxford University Press.
- . 2003. "The Place of Truth in Epistemology." In *Intellectual Virtue: Perspectives from Ethics and Epistemology*, edited by Linda Zagzebski, and Michael DePaul, 155-180. New York NY: Oxford University Press.
- Steglich-Petersen, Asbjørn. 2009. "Weighing the Aim of Belief." *Philosophical Studies* 145 (3): 395-405.
- . 2006. "No Norm Needed: On the Aim of Belief." *Philosophical Quarterly* 56 (225): 499-516.
- Street, Sharon. 2009. "Evolution and the Normativity of Epistemic Reasons." *Canadian Journal of Philosophy* 39 (S1): 213-48.
- Stroud, Sarah. 2006. "Epistemic Partiality in Friendship." *Ethics* 116 (3): 498-524.
- Sullivan-Bissett, Ema. 2017. "Biological Function and Epistemic Normativity." *Philosophical Explorations* 20 (1): 94-110.
- Talbot, Brian. 2014. "Truth Promoting Non-Evidential Reasons for Belief." *Philosophical Studies* 168 (3): 599-618.
- Tiberius, Valerie. 2008. *The Reflective Life: Living Wisely with our Limits*. Oxford University Press.
- Velleman, David. 2000. *The Possibility of Practical Reason*. Oxford University Press.
- Wedgwood, Ralph. 2007. *The Nature of Normativity*. Oxford University Press.
- . 2002. "The Aim of Belief." *Philosophical Perspectives* 36 (s16): 267-97.
- Wiland, Eric. 2012. *Reasons*. New York NY: Continuum International Publishing Group.
- Williams, Bernard. 1973. "Deciding to Believe." In *Problems of the Self*, 136-151. Cambridge: Cambridge University Press.
- Wrenn, Chase B. 2017. "Truth is Not (Very) Intrinsically Valuable." *Pacific Philosophical Quarterly* 98 (1): 108-128.

Zagzebski, Linda. 2003. Intellectual motivation and the good of truth. In *Intellectual virtue: Perspectives from ethics and epistemology.*, eds. Linda Zagzebski, Michael DePaul, 135-154 New York: Oxford University Press.

Chapter 2 - Evidential Exclusivity, Correctness, and the Nature of Belief

Introduction

Recently epistemologists have focused on a pair of claims that seem to be true about doxastic deliberation. The first, Evidential Exclusivity, says that in doxastic deliberation, only what are taken to be evidential considerations can play a role as motivating reasons to have a doxastic attitude toward a proposition p . The second, Transparency, says that in doxastic deliberation, the first-personal question “whether to belief that p ” immediately and non-inferentially gives way to the question “whether p ”. These claims have figured prominently in the literature that attempts to elucidate the “aim of belief”, in the literature that describes our epistemic deontic obligations and how they arise, and has even been used to argue for the necessity of metaethical realism. In this paper, I’m going to focus on the way in which Evidential Exclusivity and Transparency have been taken to be evidence for the truth of a certain proposal about the nature of belief, namely that it has a standard of correctness that holds that a belief that p is correct if and only if p is true. Epistemologists have claimed that the best explanation for why Evidential Exclusivity and Transparency hold for agents engaged in doxastic deliberation is that during deliberation, agents deploy the belief concept and that belief’s constitutive standard is that it’s correct iff it’s true. The plan for the paper is as follows: In section 1, I’ll explain Evidential Exclusivity and Transparency, clarify their relationship to one another, and show how they’re taken to arise from and be explained by belief’s standard of correctness. In section 2, I’ll highlight two

observations about Evidential Exclusivity and Transparency: first, the apparent possibility of deliberate suspension in cases of roughly balanced evidence, and second, the apparent possibility of reasonable divergence in synchronic and diachronic judgments about the normative status of a belief. Whereas the first observation can be taken to be a possible counterexample to EE/T, the second can be thought of as a feature of belief formation that the Standard of Correctness proposal leaves unexplained. Though my aim is not to deny EE/T, I want to draw attention to these two points because I think that the alternative explanation that I ultimately endorse will do a better job accounting for these data than the proposed standard of correctness. In Section 3 I will take a closer look at the Standard of Correctness proposal and propose a distinction between two ways of interpreting the standard's normative force. I will show how epistemologists have assumed that the weaker of the two distinguished normative senses acts as a kind of fallback position – even if sense can't be made of how the Standard of Correctness gives rise to a stronger breed of normative force, belief seems at the very least to be subject to a standard of correctness that holds true in the weaker sense. In Section 4 I will attempt to make out what the weaker interpretation of the standard of correctness amounts to via an analogy to norms of etiquette and I'll argue that we have reason to doubt even this weaker interpretation. Finally, in Section 5, we'll put together the pieces: I'll suggest that we can come up with an alternative explanation of how EE/T might arise that will account for both of the observations that the Correctness proposal left unexplained and will succeed where the Correctness proposal failed - in accounting for the core features of belief's functional profile, not just its role in deliberation.

1 A Pair of Claims about Doxastic Deliberation

It's a well-appreciated fact that the beliefs we form don't always reflect our evidence. Multiple psychological mechanisms are implicated in belief formation and some are truth-independent. Often we form beliefs that we wish were true,

but for which we have very little evidence. Yet even though this kind of causal route to belief formation is widely acknowledged, it also seems clear that only certain kinds of *reflections* on what to believe can succeed in directly bringing it about that we believe. So although we recognize that it often *happens* that we, say, believe in accordance with what we desire to be true even though we have insufficient evidence for that belief, reflecting on facts about what we desire alone cannot be the reason *for which* we form a belief. Rather, it seems that the reflections that are able to directly give rise to beliefs during deliberation are exclusively evidential considerations. Call this claim *Evidential Exclusivity*.

Evidential Exclusivity: In doxastic deliberation, only what are taken to be evidential considerations can play a role as motivating reasons to have a doxastic attitude toward a proposition *p*.¹⁴

A closely related claim, *Transparency*, tells us something about the *phenomenology* of conscious deliberation about whether to believe:

Transparency: In doxastic deliberation, the first-personal question “whether to believe that *p* immediately, non-inferentially gives way to the question “whether *p*.”¹⁵

Whenever we ask ourselves whether or not to believe some particular proposition, we immediately recognize that this question is settled by, and only by, answering the question whether that proposition is true. Nishi Shah’s proposed explanation of the phenomenon of transparency is to suggest that it is part of the concept of belief that it has a certain *standard of correctness*:

¹⁴ Steglich-Petersen 2009; McHugh 2015; Sharadin 2016.

¹⁵ Transparency was introduced by Evans 1982 and revisited by Moran 2001. This formulation comes from Shah and Velleman 2005.

(SC) Belief that p is correct iff¹⁶ it is true that p , and incorrect iff it is untrue that p . (Shah 2013, 3)

In other words, it is a *conceptual truth* that the standard of correctness for belief is truth. Here's how Shah thinks (SC) explains Transparency.

1. During doxastic deliberation, one's thinking is guided by the question "whether to believe that p ".
2. So, during doxastic deliberation, one necessarily employs the concept of belief.
3. So, during doxastic deliberation, one necessarily takes one's potential belief that p to be correct iff p is true.
4. So, during doxastic deliberation, one necessarily takes the question "whether to believe that p " to give way to the question "whether p is true".

Why is it that reflection on the evidence that p but not on my desire that p be true can bring me to believe that p ? Because, according to Shah, our reflection about what to believe is structured by our acceptance of the norm (SC). The (SC) norm explains the evidential constraint on deliberation about what to believe because that norm is contained in the very concept of belief.

If Transparency is a claim about the phenomenology of deliberation and Exclusivity is a psychological claim about what kinds of considerations can play the role of motivating reasons during deliberation, how, precisely, are the two claims related? If one takes evidence to consist in considerations that are relevant to determining the truth of p , or in other words, considerations that are relevant to

¹⁶ There's a lot of debate in the literature about the content of (SC), in particular whether the norm should say "iff" or "only if". There are reasons to prefer "only if" because it seems like we shouldn't believe everything that's true, for example all trivial entailments of our beliefs that no finite cognitive machine could ever support. I'm not going to get into these debates in this paper and will go with Shah's "iff".

answering the question “whether p is true”, then one can derive Exclusivity: In the context of first-personal doxastic deliberation, the question “whether to believe that p ” is always settled by evidential considerations alone. Since only evidential considerations can settle “whether p ”, then only evidential considerations can settle deliberation over “whether to believe p ”. Transparency thus entails Exclusivity.

Transparency, for Shah, is a necessary truth. He writes that “so long as one is considering the deliberative question of what to believe, these two questions must be answered by, and answerable to, the same set of considerations.” He says that the “seamless shift in focus from belief to truth is not a quirky feature of human psychology, but something that is demanded by the nature of first-personal doxastic deliberation” (2003, 447). Shah goes on to argue that belief’s standard of correctness provides an argument for Evidentialism, namely a position he takes to hold that “only evidence can be a reason for belief” and claims that the best way to gloss the correctness standard is via the deontic modal: “One ought to believe that p if and only if p is true.” In Section 3 I will take up the issue of how we should interpret (SC)’s normative import. For now, though, I want to make two observations about Evidential Exclusivity and Transparency that I think merit further attention.

2 Two observations about Evidential Exclusivity and Transparency

1. Recall that Evidential Exclusivity claims that only what are taken to be evidential considerations can play a role as motivating reasons to have *any doxastic attitude* towards a proposition. Whereas this likely strikes us as plausible when we consider the reasons that motivate us to form a belief that p or to reject a belief that p (form a belief that not- p), it seems that one often has some discretion in whether or not to have the doxastic attitude of *suspension* on p . In other words, it seems that in some cases of reflection on whether p , we can deliberately choose between suspending belief or believing that p for presumably non-evidential

reasons. Conor McHugh, in his 2015 paper “The Illusion of Exclusivity”, suggests a few kinds of cases which he thinks serve as counterexamples to Exclusivity. He takes them to share an important feature, namely that they’re all cases in which *the evidence at a deliberator’s disposal is sufficient to permit belief*. Nate Sharadin (2016), in another paper that aims to provide a counter-example to Exclusivity, argues that non-evidential considerations can play a role in certain cases of *roughly balanced evidence*.¹⁷ The proposal up for consideration, then, is whether in cases of sufficient or roughly balanced evidence one sometimes has the capacity to suspend belief on the basis of non-evidential considerations. Consider the following two cases:

Testimony: When we receive testimony that p , it seems as though we often have a choice about whether to accept that testimony and believe that p or instead raise for ourselves the possibility that the testifier, even one whom we might generally take to be highly reliable, is mistaken or deceiving us. If we elect to concentrate on the latter possibilities, we can deliberately suspend belief about p until we’re able to rule these possibilities out. Otherwise we can decide to ignore the possibility that the testimony is incorrect and go on to believe that p .

Inductive generalizations: When we set out to gather evidence, it can seem like a choice whether or not we’ve witnessed enough instances of a certain type to believe the truth of an inductive generalization or whether to go on seeking further confirmation that there isn’t an exception to the generalization. As a matter of psychological fact, it often seems to us that we have discretion in whether to believe some inductive proposition or to suspend belief on that proposition and

¹⁷ Sharadin’s proposed counterexample in his “Nothing but the Evidential Considerations” (2016) involves auto-alethic beliefs.

search for further proof. A particular breed of this kind of case is familiar to us from the contextualism and pragmatic encroachment literature. Whether or not we believe that the bank is open at noon on Saturday or we suspend belief about its Saturday hours seems to vary with the practical stakes.

The reason I raise the possibility of these cases of apparent deliberate suspension is not because I'm interested in denying Exclusivity and Transparency, but rather because I want to note carefully the situations in which they appear to falter in order to gain insight into what might be going on behind the scenes in the doxastic processes that produce these beliefs – information that will help us go on to construct an alternative hypothesis. After all, the supporters of EE/T are surely right that these claims obtain sufficiently frequently and widely so as to warrant an explanation – even if they're not correct to hold that they're necessarily true for all doxastic attitudes. In addition, I would be happy to interpret the Standard of Correctness explanation as a conditional claim: *If* Evidential Exclusivity and Transparency are true, the best explanation for their truth is that belief has a standard of correctness, such that a belief that *p* is correct iff *p* is true.

2. Above I noted that in cases of apparently sufficient or roughly balanced evidence, it can seem to us as if our evidence *licenses* but does not *compel* belief, thus allowing us to deliberately suspend on *p* or believe that *p*. For the second observation, I want to draw attention to an aspect of our diachronic doxastic judgment that I think must be appreciated alongside our analysis of synchronic deliberation if we are to properly evaluate the (SC) proposal. This is the apparent divergence in synchronic and diachronic judgments about the normative status of our beliefs. To begin, note that Shah seems to vacillate between descriptive and normative formulations of Transparency. In some cases he claims that the two questions (“whether to believe *p*” and “whether *p* is true”) *are* always taken to be answered or settled by the same set of considerations (2003, 447, 463). In other

instances the claim seems to be normative: The two questions “must” be seen as being answered by the same set of considerations. (Ibid., 447, 453). Shah is quick to acknowledge that the way in which evidential considerations are decisive is unique to doxastic deliberation – when we’re not consciously reflecting on what to believe, we sometimes believe for non-evidential reasons. But in keeping with his normative formulation of transparency and his judgment that (SC) is a *norm* of belief, he, in familiar epistemological company, follows the tradition of taking beliefs that are formed in accordance with non-evidential reasons to be *defective* if predictable. I want to highlight that it’s unclear that we *do* uniformly resist normative endorsement of beliefs that aren’t true once we widen our lens from the synchronic judgment of beliefs during deliberation to the diachronic judgment of beliefs as embedded in certain doxastic processes that are parts of a complete doxastic economy.

To make this thought clear, we might think about the following kind of thought experiment. Consider an agent who is asked about the normative status of a belief *p* about which she deliberates at time *t*. Presumably, according to Shah and in accordance with EE/T, her verdict is that she will/ought endorse that belief *p* iff she takes it to accord with the evidence at her disposal. Now in contrast to the synchronic judgment that we elicit, I want to think about what we might take to be a reasonable judgment about the normative status of some kinds of beliefs that are embedded in a diachronic picture. To do this, we should examine the role that certain doxastic processes seem to routinely play in our cognitive lives.

To begin, think of a doxastic process as a way of forming, maintaining, and revising one’s beliefs. As EE/T make clear, deliberation appears to be a kind of doxastic process that focuses agents on evidential considerations – considerations that are relevant to adopting true beliefs. When we survey a wider array of doxastic processes, however, we notice that there are many that seem to be directed toward a different end – an end that we might broadly characterize as enhancing that agent’s wellbeing.

Valerie Tiberius' work treats the role of positive illusions in human psychology (2008, 113). Summarizing Taylor and Brown's studies (1994, 21), she writes that the doxastic processes that produce these positive illusions cause agents to

- a) view themselves in unrealistically positive terms;
- b) believe that they have greater control over their environmental events than is actually the case; and
- c) hold views of the future that are more optimistic than base-rate data can justify.

Let's look more closely at each of these claims. a) is known as the familiar *self-enhancement bias*. That most people exhibit self-enhancement bias is an uncontroversial claim in psychology.¹⁸ Self-enhancement bias refers to the way in which humans' beliefs about themselves manifest bias in favor of a positive self-conception. Studies find that most peoples' self-evaluations fail to correspond with the evaluations of other people (Gosling et al. 1998), with experts (Helgeson and Taylor 1993), and with objective measures (Vazire 2010). People rate themselves more positively and less negatively than they rate most other people in terms of (to highlight just a few) physical attractiveness, intelligence, athletic ability, social competence, sincerity, being considerate, being rude, and being superficial.¹⁹

B) targets a kind of causal attribution bias that people routinely display. Most people take themselves to possess a higher degree of control over events in the world than they actually possess. In a group of studies, subjects are asked to predict the outcomes of a series of coin tosses, to estimate their ability to predict the outcome of coin tosses in general, and to indicate the degree to which they believed that they were able to improve their ability to predict the outcomes with

¹⁸ Hazlett's Chapter 2 "Greatness of Mind" (2013) is a great overview of the evidence for and role of self-enhancement bias that I've relied on heavily. See Taylor and Brown (1994), Armor and Taylor (2002), Tiberius (2008) and McKay and Dennett (2009) for others.

¹⁹ Dunning (1995), Trafimow et al. (2004), Hazlett (2013).

practice. Subjects' perceived outcomes were manipulated so that their beliefs about their success trajectories varied from seeing themselves as initially successful but later unsuccessful, to initially unsuccessful and later successful, or neither of these. All subjects were "correct" in their prediction 50% of the time. Initially successful candidates were far likelier to rate themselves as significantly better at predicting than were initially unsuccessful candidates, who were far likelier to claim that the ability to predict could not be improved with practice. Across all three groups, 25% claimed their predictions could be hampered with distractions and 40% felt that their ability to predict could be improved with practice. These beliefs also manifested in peoples' related betting behavior (Hazlett 2013, Langer and Roth 1975). A closely related phenomenon is the familiar manner in which we over-estimate the degree of control we have over our successes and under-estimate the control we have over our failures. We're likelier to attribute our failures to land a job to the randomness of the job market or the fluke glass that we dropped during the interview than to our philosophical incompetence or the weakness of our writing sample.

C) picks out the way in which we have overly optimistic views about the future. Studies find that our comparative estimates of our chances of experiencing particular life events are strongly correlated with the valence of the event in question. For example, college-aged students estimate their chances of owning a home to be 44% higher than the average chance of their classmates and their chance of getting divorced as 48% lower than their average classmate (Epley and Dunning 2006). These differences are significantly correlated neither with the general probability of the event in question, the perceived general probability of the event, nor the subject's "personal experience with the type of event in question" (Hazlett 2013, 48).

Buunk et al. (1990) and Sedikides and Gregg (2008) note that we can look to the way in which cognitive filters of selective attention, representation, and recall interact with our doxastic processes to influence the doxastic outputs of the believer. Overly positive beliefs about ourselves can be maintained via selective

inquiry and selective attention to evidence, in which we seek out positive information and avoid negative information, and dwell on evidence that positively represents us and avoid attending to evidence that reflects negatively on us. And studies have found that most people have “differential processing and recall speeds” for positive and negative information – in other words, we recall and process positive information more rapidly than negative information (Mele 2001).

One may respond to the above facts about our doxastic biases by concluding that peoples’ lives, or at least their cognitive lives, are going pretty poorly. We might think that we’re all far from being ideal thinkers and most of us are far from any reasonable conception of flourishing, so it should come as no surprise that we have a lot of false beliefs. Or we might think, more modestly, that even if our lives aren’t so bad overall, that the false beliefs that we have constitute ways in which our beliefs are *defective*, so at least in this sense they contribute to our lives going worse.

These intuitions gain support by focusing on the value of self-awareness. Self-awareness helps give us a sense of what would makes our lives go well. It can help us cultivate the habits required to make the choices that are more likely to be satisfying and to be more effective in the pursuit of our goals. Tiberius recognizes the value of self-awareness and sees it as an essential part of a general conception of wisdom. However she describes self-awareness as including “the skills and commitments that are necessary for acquiring self-knowledge and for constructing a self-conception that befits one’s own view about how to live *without destroying the harmless illusions we have about ourselves that make life better*” (2008, 111, emphasis mine). Because Tiberius recognizes that having a systematically elevated conception of our abilities is “correlated with contentment...the ability to care for and about others, openness to new ideas and people, creativity, the ability to perform creative and productive work, and the ability to grow, develop, and self-actualize, especially in response to stressful events”, she concludes that “too much...self-knowledge in many contexts seems to have some serious disadvantages” and that reflective wisdom requires the capacity to know when it

makes sense to adopt a perspective that focuses on the “importance of close attention to one’s motivations and a commitment to get at the truth of oneself” (Ibid., 115).

The thought mentioned above, then, - namely that false beliefs necessarily contribute to our lives going worse, - is implausible. To see this more clearly, consider the correlation between depression and the doxastic biases that we’ve just considered. Depression includes a variety of symptoms, some cognitive (negative self-conception, negative expectations), some emotional (dejected mood, self-dislike, loss of gratification), some motivational (loss of motivation, suicidal desires), and some somatic (loss of appetite, loss of libido) (Hazlett 2013, Beck and Alford 2009). Self-enhancement bias (Brown 1993), inflated beliefs about one’s control over one’s environment (Taylor et al. 1994), and unrealistic optimism (Helweg-Larsen and Shepperd 2001) are all highly significantly correlated with non-depression. And whereas non-depressed people recall and process positive information more rapidly than they recall and process negative information, people with depression do not show this trait. Hazlett (2013) calls the thesis that “non-depressed people’s beliefs about themselves tend to be inaccurate and manifest self-enhancement bias” *non-depressive unrealism* (53). He goes on to cite studies that claim that the higher self-esteem characteristic of non-depressive unrealism is correlated with higher subjective wellbeing,²⁰ greater motivation,²¹ and superior coping abilities.²²

It wouldn’t be interesting for the (SC) proposal that this paper considers to point out that some false beliefs are beneficial. But I think what would be consequential is if we could show that there exist *patterns* of belief formation that

²⁰ There seems to be a causal connection between SEB and subjective well-being. Consider studies that show that subjects’ mood states vary based on whether they engage in self-serving causal attributions following the performance of a task (McFarland and Ross 1982), subjects’ levels of optimism predicted their levels of postpartum depression (Carver and Gaines 1987) and subjects’ moods could be elevated by engaging in positively biased social comparisons (Gibbons 1986).

²¹ Subjects with overly positive self-images work harder and longer on tasks and are more able to complete tasks.

²² Buunk et al. (1990) finds biased social comparison to be correlated with diminished stress, anxiety, hopelessness and depression in response to negative life events.

systematically, predictably, and beneficially produce false beliefs. For then we would need to proceed carefully in order to spell out the notion of correctness in play when we claim that necessarily, correct beliefs are true beliefs.

I tried in this section to show that this state of doxastic affairs is indeed plausible. Recall that our aim in setting off on this detour through the literature on belief formation was to set ourselves up to consider a particular question: namely whether there might be a difference in our normative judgments in cases of doxastic deliberation and in cases in which we reflect on the role that beliefs that are the output of some particular doxastic processes seem to play in our overall wellbeing. Here's my claim: Imagine that we point to some belief that's embedded in a doxastic process that's systematically correlated with enhanced subjective wellbeing, motivational benefits, improved coping capacities, and non-depression. If we ask someone whether or not they normatively endorse that belief, I maintain that, at the very least, it would be reasonable for that person to respond affirmatively. EE/T take as their data our judgments in instances of deliberation to draw conclusions about the norm to which belief is subject. I think we must include our observations about divergent diachronic normative judgments in our assessment of the data we will use to inform ourselves about belief's features and standards. We will keep this data point in mind as we go on to consider belief's Standard of Correctness proposal and formulate an alternative story.

3 The Standard of Correctness Proposal

Recall that Shah and Velleman think that the best explanation for EE/T is that when we deliberate we deploy the concept of belief, which has a constitutive standard of correctness such that a belief is correct iff it's true. How should we interpret the normativity of such a standard? As mentioned in Section 1, Shah thinks this standard can be understood as entailing a corresponding claim about what we ought to believe – we ought to believe p iff p is true. Here are some other takes on the normativity of belief's standard of correctness:

Gibbard: For belief, correctness is truth. Correct belief is true belief. My belief that snow is white is correct just in case the belief is true, just in case snow is white. Correctness, now, seems normative...The correct belief seems to be the one a subject ought...in this sense, to have. (2005, 338)

McHugh: When we say that someone has an incorrect belief, we do not seem to be merely describing some feature of her, or of her belief – a feature whose presence she could sensibly remain unconcerned by. We are saying that she believes *wrongly*. (2017, 1452)

Wedgwood: According to the assumption...that “correctness” is a normative concept, if your belief is incorrect, it has a certain sort of *defect* – while if your belief is correct, it is wholly free from that defect. (2013, 218)

Engel: The term “correctness” here is meant to characterize normativity in the broad sense...According to (CT),²³ the basic dimension of evaluation of the correctness of belief is truth...When our beliefs are false, they do not simply lack the property of being true or of being fitted in the proper way to the world: they are incorrect and in need of correction...(CT) not only tells us when a belief is true or false: it tells us that it is *right* to believe what is true, and that to fail at having true belief is in some sense *wrong*. (2013, 33)

When Gibbard tells us that the correct beliefs are the ones a subject *ought to have*, McHugh and Engel tell us that *believers who believe falsely believe wrongly* in a

²³ Engel's (CT): For any *P*, a belief that is *P* is correct if and only if *P* is true (2013).

sense in which they *couldn't sensibly fail to care about it* and Wedgwood tells us that false beliefs are *necessarily defective*, a strong view about the normative implications of correctness begins to emerge. This might seem to contrast, though, with other applications of the generic notion of correctness. Consider the following:

- The correctness or incorrectness of an artistic rendering of a skeleton based on its anatomical features
- The correctness of a photographic representation based on its likeness to its subject
- The correctness of an orchestral score as fixed by a set of notes
- The correctness of a map as fixed by the similarity between the map and the terrain it represents

The properties these standards target at least appear to be descriptive, not normative. We might think that Dretske has a similar idea in mind about belief's relationship with truth. He writes

Beliefs and judgments must either be true or false, yes, but there is nothing normative about truth and falsity. What makes a judgment false or true is the fact that it fails or succeeds in corresponding to the facts, and failing or succeeding in corresponding to the facts is, as far as I can see, a straightforward factual matter. Nothing normative about it. (2001, 247)

According to Dretske, the fact that a belief can be true or false is a general descriptive principle of belief, but one that carries no normative force – presumably he thinks that it carries no weight in terms of how a believer should believe. Alternatively, we might acknowledge that on its face belief's standard of correctness appears to express a norm, but think in the spirit of Dretske's claim

that this norm carries no normative force. One way to put this is that belief's standard of correctness is a norm that doesn't give rise to reasons. How might we articulate such a view?

To start, notice that norms are born easily. I can create a norm that holds that all trips to the bathroom during paper-writing sessions must be completed on one foot. We should think, though, that the mere fact that I've decreed such a rule doesn't yet give anyone any good reason to act in accordance with it.²⁴ Etiquette, the mafia, hipsterdom, and serial killing all plausibly have associated bodies of norms or standards, yet don't necessarily give rise to good reasons to follow them. So we can recognize that all norms trivially set standards according to which things are permissible, prohibited, good, bad, and the like and simultaneously acknowledge that these standards do not necessarily give rise to good reasons to act in accordance with them. As Côté-Bouchard puts it,

If a norm N forbids ϕ -ing under conditions C, then trivially, ϕ -ing under C is forbidden, incorrect, wrong, or bad *relative to* the standard set by N. However, the descriptive *norm-relative* question of what is permitted or required according to N is distinct from the *normative* question of what there is *good* or *genuinely normative* reasons to do. (2016, 3182, emphasis his)

On the picture I'm endorsing, reasons come about in accordance with value.²⁵ So,

²⁴ If we feel that some such norm does exert at least a *small* pull on us, we should consider a diagnosis that looks at a kind of *derivative* force that comes about from the role that norms tend to play in the societies we inhabit. But the important difference is that the normativity at issue in such a case would be *merely derivative*, not arising from the *content* of the norm itself.

²⁵ McGuire (2016) gives a comprehensive overview of a value-based theory of reasons and takes the following claim to be an *analysis* of what it is to be a reason. I take McGuire's view to be a *substantive* view about reasons, but one that I think is independently preferable for reasons I can't explore in this paper.

For any fact x that holds that some action ϕ would promote some state of affairs S , x is a reason to ϕ iff, and due to the facts that, ϕ would promote S and S is valuable.

This gives a nice explanation of why not all norms necessarily give rise to reasons to follow them. Presumably we think norms of serial killing set standards about what is encouraged or discouraged relative to those norms, but that those norms don't give rise to reasons because serial killing is not valuable (/is positively disvaluable).

To help us think further about the normativity that undergirds belief's standard of correctness, I want to combine our value-based picture of reasons with a related distinction between two kinds of normativity. A number of normativity theorists have made use of a breed of the distinction I'm targeting. Schroeder (2003), for one, distinguishes two senses of norms. In the first sense there are norms of categorization or classification schemes that divide actions or events into distinct types, or that act as general principles of description. In a second sense, are norms as what he calls *force-makers*. Norms in this second sense act as prescriptions or governance principles that guide our actions or give us aims to follow (2003, 3).

Parfit makes use of a similar distinction between norms that he calls *rule-implying* and norms that he calls *reason-implying*. For Parfit, "facts are normative in the *rule-implying* sense when these facts are about what is correct or incorrect, or allowed or disallowed, by some rule or requirement in some practice or institution" (2011, 309). He takes as examples of facts that are normative in the rule-implying sense that certain acts are illegal, that some act would be bad etiquette, or the wrong use of a word, or an impermissible move in chess or some other game. He goes on to say that "we can describe in non-normative terms what is involved when such social rules and practices are established...and when certain acts break these rules, that's *what it is* for these acts to be disallowed or incorrect in these ways." Parfit thinks we can't make similar claims about *what it is* for there

to be a norm that gives us a reason. The reason-implying sense, in contrast, tells us what we “should or ought to do” (Ibid.). Parfit’s rule-implying sense of norms maps nicely onto Schroeder’s first sense of norms, and his reason-implying sense of norms maps nicely onto Schroeder’s second sense of norms as force-makers.

We have in place, then, a distinction between norms that are individuated by the existence of a rule or practice that don’t necessarily give rise to reasons to follow them, and norms that necessarily give us some reason to follow them. Call the first kind the “weak sense” and the second kind the “strong sense”. Of course norms of the two types often overlap. Good examples are many legal norms and moral norms that both give us good reasons to act and involve or depend on moral rules or requirements. We might even hold a meta-normative view on which all norms in the strong sense entail normativity in the weak sense. But the important thing to note for our present purposes is that there can exist norms in the weak sense that don’t necessarily give rise to normativity in the strong sense.

What, then, is the sense of normativity at stake in (SC)? The above distinction gives us a way to think about the notion of correctness in play in belief’s standard of correctness more clearly. The theorists quoted at the beginning of the section (Gibbard, McHugh, Wedgwood, Engel and in addition, Shah and Velleman) all endorse breeds of force-making, action-directing, reason-implying normativity of the strong sense. But although Shah and these other normativists take (SC) to apply to belief necessarily, I think when we combine our understanding of norms as force-makers and reason-givers only when they promote states of affairs with value, then we can see our way towards an argument that correctness can’t in fact pick out the strong sense of normativity. For the considerations about doxastic processes from the second part of Section 2 give us good reason to think that the states of affairs that have most value are often ones in which we do not believe truly. Since the value in ϕ -ing explains what reasons we have to ϕ , and in some cases it appears as if believing truly does not give rise to a valuable state of affairs, then it can’t be the case that the correctness in play in (SC) is normative in the strong sense that necessarily gives rise to reasons to follow it.

One might think that this way of putting things commits the strong sense normativists about correctness to too much. Plausibly they take themselves to be committed merely to the claim that (SC) necessarily bestows *some* action-guiding force, or *some* reason to believe p iff p is true. Surely we shouldn't commit them to the view that they take (SC) to necessarily give rise to an *all-things-considered* ought or to *overriding* reason to believe p iff p is true. But what they are committed to is that (SC) necessarily gives rise to *pro tanto* normative force or a reason to follow it. I think this view is also implausible. It would require that we take true beliefs to have necessary value. Since we have already seen how it's doubtful that they always have instrumental value, we would have to think that every true belief is necessarily finally valuable. Though it's hard to dissuade staunch fans, I argue elsewhere that this position has little to recommend it. Here I'll say just that the view looks extremely implausible when we consider beliefs about propositions concerning the number of grains of sand in a particular square acre of Namibian desert.

I've argued that we shouldn't think that belief's standard of correctness is normative in the strong sense. In the next section I want to turn to uncovering what an interpretation of the normativity of correctness in the weaker sense might look like.

4 The Standard of Correctness: The Weak Version

It's not uncommon for theorists who doubt the possibility of a strong construal of (SC)'s normativity to settle on a kind of fallback position: though the standard of correctness might not give rise to any necessary principles of active governance, it's a basic fact about belief that it's a state that's correct iff it's true. But what does such a claim amount to? When we drew the distinction between a strong and weak sense of normativity, we followed Schroeder's characterization of the weak sense as consisting in "categorization or classification schemes that divide actions or events into distinct types", or acting as "general principles of description" and

Parfit's characterization of facts whose correct-making features are exhausted by what is "allowed or disallowed by some rule or requirement in some practice or institution." Both gave norms of etiquette and rules of a game as examples that fell into the weak sense category. We might then try to understand belief's standard of correctness by thinking of it as similar to an etiquette norm or a game rule. On this picture we could say that *what it is* for it to be the case that, in the etiquette domain, putting your napkin in your lap is correct, is for it to be endorsed by a corresponding norm of etiquette. Similarly, what it is to be a correct move in Tic-Tac-Toe is for that move to be permissible by the rules of Tic-Tac-Toe. So on our comparison with the correctness standard for belief, what it is to be correct *qua belief* is to be correct iff true.

This comparison has in fact been made in the literature on the norms of belief. Côté-Bouchard suggests that perhaps belief norms "have the same kind of normative authority as e.g. norms of etiquette, fashion, games, and the like." He thinks that on this picture "there can still be facts about what we should...believe. But just like facts about what we should do according to etiquette, they are not genuinely normative facts" (2016, 3196). Hazlett makes a similar suggestion (2013). So on this view, it's good *relative to the believing domain* that we believe propositions iff they're true, but that's not to say that the believing domain issues norms that have force in the strong sense. I propose, then, that we consider more carefully whether an apt comparison can be made between belief's standard of correctness and a norm of etiquette or a game norm.

Consider two domains: etiquette and Tic-Tac-Toe. And take a token norm from each domain. Let's use as examples:

domain: etiquette

token norm: When seated at the dinner table, it's correct to place one's napkin in one's lap.

domain: Tic-Tac-Toe

token norm: When playing Tic-Tac-Toe it's correct to try to configure the board so that three consecutive squares are yours.²⁶

What's important to notice when we consider token norms in these domains, is that they avail themselves of no *further* kinds of consideration *qua* norm of that domain. There are no available further considerations about a "Put your napkin in your lap" norm *qua* norm of etiquette, and there are no available further considerations about a "Mark three squares in a row to win" norm *qua* norm of Tic-Tac-Toe. One way we might put this thought is that our concept of etiquette can be fully cashed out in terms of, say, requirements of politeness or propriety. And our concept of Tic-Tac-Toe is one that can be fully cashed out in terms of the rules that constitute a game win. This is not to say, of course, that we can't imagine reasons for which we might ignore or flout these norms. I might have overriding reason not to put my napkin in my lap because we're having a contest to see whose napkin can remain the cleanest, and we've agreed that our napkins must be on full display at all times. And I might have all-things-considered reason to let you win at Tic-Tac-Toe because I fear your petulance when you lose. But notice that when we deploy such considerations, we've clearly departed from evaluation from *within the domain of etiquette* or from *within the domain of Tic-Tac-Toe*. In other words, there's a sense in which the *standard of the norm exhausts the conceptual role*.

Do the same considerations hold when we compare the standard of correctness and the concept of belief? I think it's very different in the belief domain and the difference gives us good reason to reject the thought that there's any sense to be made of the weak normative interpretation that endorses the mere fact that beliefs are correct iff true. When we think back to the considerations about doxastic deliberation with which we began, presumably we're supposed to think that Evidential Exclusivity and Transparency are features of deliberation that *suggest to us* belief's standard of correctness. And Shah tells us that it's the fact

²⁶ Assuming you've followed the other rules of play.

that belief has a standard of correctness such that a belief is correct iff only if it's true that is supposed to *explain* EE/T. But it seems perfectly in keeping with the concept of belief to consider beliefs formed and regulated in accordance with truth-independent reasons. And when we think about false beliefs that are the results of beneficial doxastic biases, these considerations can seem powerful, authoritative, or dominant. Unlike etiquette or Tic-Tac-Toe, whose conceptual domains seem to be exhausted by their sets of norms, we seem to remain clearly within the conceptual domain of belief when we consider or evaluate beliefs held for truth-independent reasons. It's simply that during doxastic deliberation we can't/don't endorse them. Because of this, I think it's neither true to say that it's our concept of belief that moves us to settle deliberation about what to believe only in accordance with truth-related considerations or that belief has a standard of correctness such that a belief is correct iff it's true. It's neither the case that we can interpret the correctness in play in that standard as giving rise to certain normative recommendations or that it's simply true in a sort of categorizational sense that it falls out of the concept of belief that the correct ones are the true ones.

Shah's and Velleman's explanation of EE/T hold that it's the deployment of the concept <BELIEF>, or in other words, the awareness that it's the content of a *belief* that we're deliberating over, that triggers us to treat only evidential considerations as relevant. But this explanation seems unsatisfactory because it can very well seem to be an open question – even during deliberation – about whether it might be better if we were to believe in accordance with non-truth-related considerations. Shah writes that the “deliberative question whether to believe that p inevitably gives way to the factual question whether p ”. But it would seem to require no conceptual confusion to pose to oneself the question whether to believe that p , and then to go on to consider what the effect on one's wellbeing might be if one were to believe that p . Rather what's at stake is the fact that one can't *transition directly* from non-evidential considerations to belief. But that fact doesn't need to be explained by positing a norm to which the belief state is necessarily subject.

To see that it doesn't seem to be the concept of belief that's doing the work, notice that many (most?) of the instances in which a deliberator might be described as "deliberating about whether to believe p " she could just as well be said to be asking herself "whether p " without any reference to belief at all. Shah might take this as a point in favor of his claim that the first question immediately reduces to the second. But we might instead take it as a reason to question whether we can derive any conclusions about the normative nature of *the concept of belief* at all. Consider the difference between two questions:

1. Stella is trying to find out whether Hillary will win. She first consults fivethirtyeight.com, which gives Hillary a 64% chance, and then she consults electionbettingodds.com, which gives her an 82% chance. She then asks herself "Should I believe that Hillary will win?" and begins to conduct an investigation into the reliability of the websites.
2. (3 days later) Stella is thinking about the tragedy of life on Earth. She thinks the world seems overly cruel and that, contrary to what her grandmother thinks, it's not the case that "everything happens for a reason". She asks herself if maybe she should believe in an afterlife, despite her lack of evidence.

The first question I think reduces straightforwardly to the question "Is p true?". The second question doesn't reduce so neatly. Stella's not wondering whether p is true, but rather whether or not it would be to her advantage to believe p . She's not conceptually confused, she is deliberating about whether to believe p ,²⁷ and, contrary to what Shah says, her deliberative question did not immediately give way to the question "whether p is true." It's interesting that it's the second

²⁷ If we build into deliberation that considerations are necessarily truth-relevant, then we've begged the question.

question that seems to be more strongly “framed by the concept of belief” than the first, in which we could more easily excise all reference to belief.

I’ve tried to put us in a position to think that the (SC) proposal is an unsatisfying explanation of EE/T. Instead of licensing us to derive normative conclusions about the state of belief, I’ve suggested that we should think that EE/T can be used only to highlight a feature of deliberation (and as suggested in the preceding paragraph, maybe only *some* kinds of deliberation). In the final section, I want to sketch the beginnings of a picture that could serve as an alternative explanation of EE/T – one that not only takes into account features of deliberation, but also tries to respect other core features of belief’s functional profile.

5 An Alternate Explanation

We should hope that an alternate explanation for Evidential Exclusivity and Transparency will

- a) make sense of those claims in the way they in fact seem to occur for us;
- b) make sense of our two observations about deliberation with roughly balanced evidence and normative synchronic and diachronic divergence; and
- c) (related to b’s second observation), do a better job accounting for core features of belief’s functional profile.

Before I start, I want to say something about the apparent dialectical situation as I see it and about the kind of alternate suggestion I’m going to make. To my knowledge there aren’t any other explanations for Evidential Exclusivity and Transparency that exist in the literature. Though Transparency, and more recently, Evidential Exclusivity, have been topics of frequent discussion and have been taken to license a wide variety of conclusions about belief, Shah’s (and later

Shah's and Velleman's) argument looks to be the only explicit attempt to posit an explanation for their existence. There are, as we've seen, many interpretations of the *breed* of normativity that arises from belief's standard of correctness. But I don't think there are other options up for consideration regarding what could explain EE/T other than a standard of correctness. This naturally limits what we consider to be genuine options, and I think it's played a role in shaping our understanding of the normativity of belief. The alternative picture I'm going to sketch is a bit of an evolutionary just-so story. Evolutionary just-so stories can be frustrating because they're often marked by evidential under-determination. That's true of this one as well, but I think it's substantially more plausible than the existing alternative, and, especially given the lack of considered options, should be on the table.

How do we learn about belief's functional profile? I suggest that the best method is to combine examination of our deeply held commitments with an exploration of creature construction. An examination of our deep commitments should elucidate for us the target of explanation. What features of belief are we committed to talking about, to reasoning with, to positing in our best explanations and predictions of action guidance and behavior? What must the mental state be like to satisfy such roles? And how might a creature be constructed and evolve so as to function in accordance with these roles?

Though we should think that the answers to these questions leave us with a bundle of candidate traits, I want to focus on what I think we can take to be the core of these traits. Eric Mandelbaum writes

I take it that, whatever beliefs are, they must have at least the following properties: they must be semantically evaluable (have conditions of satisfaction), be able to be acquired by perception, interact with desires and other motivational states to cause behavior...and be able to serve as the premises in inference. (2014, 82)

Whereas we might take Shah's explanation to be sensitive to the first and third of these conditions, I think his proposal neglects belief's role in motivation. In characterizing belief's functional role, Glüer and Wikforss take the core components to involve, on the input side, a distinctive sort of evidence-sensitivity, and on the output side, action guidance - both inferentially, by providing reasons for further beliefs and actions, and non-inferentially (2013, 138). I want to look more closely at the difference between inferential and non-inferential guidance, and also at what they mean by the "distinctive sort of evidence-sensitivity".

It's important to remind ourselves that much of belief's action-guidance happens non-inferentially. As Railton notes, "For the most part, the regulation of action by belief 'takes care of itself' without need for supervision or intervention" (2014, 139). When we focus on deliberation, we can lose track of this fact. Though it might seem to us a mundane or obvious feature, I think it bears noting because it seems as if deliberative and non-deliberative guidance might happen in importantly different ways.

After Tiberius notes that reflective wisdom consists in the capacity to alternate between a perspective of rigorous self-awareness and the abandonment of this perspective, she cites studies claiming that this kind of self-awareness is indeed possible: people's illusions are more modest when they are in the deliberative mindset, contemplating a decision (2008, 126).²⁸ Hazlett supports this claim. He notes that doxastic activity manifests patterns of selective self-enhancement. People with serious illnesses compare themselves with those who are worse-off for self-enhancement purposes, but occupy a deliberative mind-set aimed at acquiring accurate information when seeking out treatment plans. In addition, he cites multiple studies claiming that "unrealistic optimism and illusions of control are manifested to a greater degree in situations of implementation and

²⁸ See for example Gollwitzer and Kinney 1989.

action than in situations of deliberation and decision-making.”²⁹ Psychologists who see positive illusions as a significant mark of good mental health posit an “optimal margin of illusion” that aims to track the sweet spot between complete accuracy (which correlates highly with depression) and delusional inaccuracy (which leads to imprudent behavior) (Hazlett 2013, 56). For this reason, it seems like we should consider as plausible the thought that deliberation, far from exposing general features of the belief state, is a special kind of doxastic process that enables us to actively consider factors relevant to accuracy.

As we put these facts together, a picture of belief’s functional role begins to emerge. On the input side we think sensitivity to evidence is important in demarcating beliefs from other mental states. But we noted that it was a *special kind* of sensitivity to evidence. The output side of belief’s functional role has as an important component belief’s disposition to guide action. As we’ve seen, in order to guide actions that result in a creature’s enhanced wellbeing, the doxastic representations that are required are ones that often aren’t accuracy-regulated. The degree of evidence-sensitivity, then, should be coordinated with the “optimal margin of illusion.”

Consider, then, a perspective from creature construction. One model we could consider is a creature who seeks to maximize expected utility. But let’s suppose it turns out that that creature’s planning capacities are far from the ideal. It’s exceedingly poor at forgoing current pleasures or satisfactions in favor of increasing utility down the road – even in cases when the long-term utility far outweighs the sacrifice in short-term utility. In addition, as John Heil notes, it may well be “on the whole, ill-equipped to produce dispassionate reckonings of the long-range consequences of actions even in settings of only moderate complexity” (1992, 54). So the first creature loses out to a creature that is programmed

²⁹ In one, for example, subjects were put into a “deliberative mindset” in which they were told that their task was to find out which of two lights they have more control over, while other subjects were put into an “implemental mindset” and told that their task was to turn on a light of their choosing as many times as possible. Subjects in the implemental group overestimated the number of lights they could turn on to significantly higher degree than the deliberative group.

according to a different model. This second creature operates during deliberation in accordance with a simple, easily implementable norm: “Believe the truth.” Though it’s not always possible for the creature to follow that norm directly, it can seek to operate in accordance with straightforward bridge principles like “Believe in accordance with your evidence.”

Heil claims that a “finite doxastic system in which non-epistemic elements functioned as constant backseat drivers threatens stability” (Ibid., 53). Not only would it threaten stability, I claim, but it would suffer from significant rule erosion. If a system has a ready capacity to manually override its operational rules, those rules cease to have efficacious force. This, I think, is significantly related to our first observation of Section 2. It’s interesting to consider why we seem only to be able to exert deliberative discretion in cases of *roughly balanced evidence*. We could imagine things going otherwise – for example, why isn’t it the case that we have discretion only in cases in which what we stand to gain prudentially surpasses some threshold? Philosophers often begin their cases for EE/T or against doxastic voluntarism by noting our inability to believe that there’s an elephant on our desks, even for massive monetary reward. It’s these cases, though, that would result in the most significant system damage – we would have to see ourselves as flouting the norms in the most obvious way, thus opening wide the door for recurrent overrides. It’s no surprise, then, that a creature would be constructed so as to maintain relative immunity to this kind of damage.

We might liken this situation to a distinction familiar in the moral domain. We can compare the substantive principles that a system of norms mandates with its regulative principles, the principles agents must follow in order to satisfy the substantive norms. Thus, utilitarianism might require agents to cultivate non-consequentialist motives to best achieve substantive utilitarian aims.

The alternate explanation, then, suggests that we’re left with a picture on which the best way to realize an overall more valuable outcome is a system on which we deliberate in accordance with a norm such as (SC), yet have a doxastic system that is directed not towards realizing belief’s standard of correctness, but

rather towards increasing wellbeing. In this way, there isn't a true interpretation of belief's standard of correctness. Though it seems we may be built to deliberate in accordance with such a norm, that itself isn't an indication that, in any meaningful sense, a belief is correct iff it's true.

References

- Armor, David A. and Taylor, Shelley E. 2002. "When Predictions Fail: The Dilemma of Unrealistic Optimism." In *Heuristics and Biases: The Psychology of Intuitive Judgment*, edited by T. Gilovich, D.W. Griffin, and D. Kahneman, 334-47. Cambridge University Press.
- Beck, Aaron T., and Alford, Brad A. 2009. *Depression: Causes and Treatment*. Second edition. University of Pennsylvania Press.
- Brown, Jonathon D. 1993. "Self-Esteem and Self-Evaluation: Feeling is Believing." In *Psychological Perspectives on the Self, Volume 4 The Self in Social Perspective*, edited by J.M. Suls, 27-58. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Buunk, B.P., R.L. Collins, S.E. Taylor, N.W. VanYperen, and G.A. Dakof. 1990. "The Affective Consequences of Social Comparison: Either Direction has its Ups and Downs." *Journal of Personality and Social Psychology* 59 (6): 1238-49.
- Carver, C.S., and J.G. Gaines. 1987. "Optimism, Pessimism, and Post-Partum Depression." *Cognitive Therapy and Research* 11, 449-62.
- Côté-Bouchard, Charles. 2016. "Can the Aim of Belief Ground Epistemic Normativity?" *Philosophical Studies*, 173 (12): 3181-98.
- Dretske, Fred. 2001. "Norms, History, and the Mental." In *Naturalism, Mind, and Evolution* edited by Denis Walsh. Cambridge: Cambridge University Press.
- Dunning, David. 1995. "Trait Importance and Modifiability as Factors Influencing Self-Assessment and Self-Enhancement Motives." *Personality and Social Psychology Bulletin* 21(12): 1297-306.
- Engel, Pascal. 2013. "In Defense of Normativism about the Aim of Belief." In *The Aim of Belief* edited by Timothy Chan, 32-63. Oxford University Press.
- Epley, Nicholas and David Dunning. 2006. "The Mixed-Blessings of Self-Knowledge in Behavioral Prediction: Enhanced Discrimination but Exacerbated Bias." *Personality and Social Psychology Bulletin* 32(5): 641-55.
- Evans, Gareth. 1982. *The Varieties of Reference*. Oxford University Press.
- Gibbard, Allan. 2005. "Truth and Correct Belief." *Philosophical Issues*, 15 (1): 338-50.

- Gibbons, F.X. 1986. "Social Comparison and Depression: Company's Effect on Misery." *Journal of Personality and Social Psychology* 51(1): 140-148.
- Glüer, Kathrin. and Åsa Wikforss. 2013. "Aiming at Truth: On the Role of Belief." *teorema* 32 (3): 137-162.
- Gollwitzer, Peter M., and Ronald F. Kinney. 1989. "Effects of Deliberative and Implemental Mind-Sets on Illusion of Control." *Journal of Personality and Social Psychology* 56 (4): 531-42.
- Gosling, S.D., O.P. John, K.H. Craik, and R.W. Robins. 1998. "Do People Know How They Behave? Self-Reported Act Frequencies Compared with On-Line Codings by Observers." *Journal of Personality and Social Psychology* 74 (5): 1337-49.
- Hazlett, Allan. 2013. *A Luxury of the Understanding: On the Value of True Belief*. Oxford University Press.
- Heil, John. 1992. "Believing Reasonably." *Noûs* 26, no. 1 (March): 47-61.
- Helgeson, Vicki S. and Taylor, Shelley E. 1993. "Social Comparisons and Adjustment Among Cardiac Patients." *Journal of Applied Social Psychology* 23 (15): 1171-95.
- Helweg-Larsen, Marie. and James A. Shepperd. 2001. "Do Moderators of the Optimistic Bias Affect Personal or Target Risk Estimates? A Review of the Literature." *Personality and Social Psychology Review* 5 (1): 74-95.
- Langer, Ellen J., and Jane Roth. 1975. "Heads I Win, Tails It's Chance." *Journal of Personality and Social Psychology* 32 (6): 951-5.
- Mandelbaum, Eric. 2014. "Thinking is Believing." *Inquiry*, 57 (1): 55-96.
- McGuire, Barry. 2016. "The Value-Based Theory of Reasons." *Ergo*, 3 (9): 233-262.
- McHugh, Conor. 2017. "Engel on Doxastic Correctness." *Synthese* 194 (5): 1451-1462.
- . 2015. "The Illusion of Exclusivity." *European Journal of Philosophy* 23:4, 1117-36.
- McKay, Ryan T., and Daniel Dennett. 2009. "The Evolution of Misbelief." *Behavioral and Brain Sciences* 32: 493-561.
- McFarland, Cathy, and Michael Ross. 1982. "Impact of Causal Attributions on Affective Reactions to Success and Failure." *Journal of Personality and Social Psychology* 45, 937-46.

- Mele, Alfred R. 2001. *Self-Deception Unmasked*. Princeton University Press.
- Moran, Richard. 2001. *Authority and Estrangement: An Essay on Self-Knowledge*. Princeton University Press.
- Parfit, Derek. 2011. *On What Matters*. Oxford University Press.
- Railton, Peter. 2014. "Reliance, Trust, and Belief." *Inquiry* 57 (1): 122-150.
- Schroeder, Timothy. 2003. "Donald Davidson's Theory of Mind is Non-Normative." *Philosophers' Imprint* 3 (1): 1-14.
- Sedikides, Constantine, and J.D. Green. 2000. "On the Self-Protective Nature of Inconsistency-Negativity Management: Using the Person Memory Paradigm to Examine Self-Referent Memory." *Journal of Personality and Social Psychology* 79 (6): 906-22.
- Sedikides, Constantine, and Aiden P. Gregg. 2008. "Self-Enhancement: Food for Thought." *Perspectives on Psychological Science* 3: 102-116.
- Shah, Nishi. 2003. "How Truth Governs Belief." *Philosophical Review* 112: 447-82.
- . 2013. "Why we Reason the Way We Do." *Philosophical Issues* 23 (1): 311-325.
- Shah, Nishi, and J. David Velleman. 2005. "Doxastic Deliberation." *Philosophical Review* 114 (4): 497-534.
- Sharadin, Nathaniel P. 2016. "Nothing but the Evidential Considerations?" *Australasian Journal of Philosophy* 94 (2): 343-61.
- Steglich-Petersen, Asbjørn. 2009. "Weighing the Aim of Belief." *Philosophical Studies* 145 (3): 395-405.
- Taylor, Shelley E. 1991. *Positive Illusions: Creative Self-Deception and the Healthy Mind*. New York: Basic Books.
- Taylor, Shelley E., and Jonathon D. Brown. 1994. "Positive Illusions and Well-Being Revisited: Separating Fact from Fiction." *Psychological Bulletin* 116 (1): 21-7.
- Tiberius, Valerie. 2008. *The Reflective Life: Living Wisely with our Limits*. Oxford University Press.

Trafimow, D., M.L. Armendariz, and L. Madson. 2004. "A Test for Whether Attributions Provide for Self-Enhancement or Self-Defense." *Journal of Social Psychology* 144 (5): 453-63.

Vazire, S. 2010. "Who Knows About a Person? The Self-Other Knowledge Asymmetry (SOKA) Model." *Journal of Personality and Social Psychology*. 98 (2): 281-300.

Wedgwood, Ralph. 2013. "Doxastic Correctness." *Aristotelian Society Supplementary Volume* 87 (1): 217-234.

Chapter 3 - Tradeoffs and Epistemic Value

Introduction

One of the most active contemporary debates in epistemology concerns the nature of epistemic norms – what explains their content and how they instruct believers. In particular, these debates have recently centered on whether or not epistemic norms are teleological/consequentialist or non-teleological/deontological.³⁰ In this paper, I want to look in detail at a few views that claim that an otherwise-appealing teleological understanding of our epistemic norms must be incorrect. Though we're attracted by a picture that justifies our epistemic norms in terms of the valuable end states they produce, some of our most central epistemic intuitions, these views suggest, conflict with this picture. My main interest lies in examining the repercussions of such a view for our understanding of the normativity of the epistemic domain. Though the debate about teleological vs. non-teleological epistemic norms has begun to garner quite a bit of attention, we've yet to explore the ramifications of this debate for our understanding of the nature of epistemic value. I'm going to argue that a close look at our epistemic judgments, and in particular, how these judgments compare to judgments in the moral domain, sheds light on the nature of epistemic value, and, ultimately, on how we should think about the fundamental contours that shape the epistemic domain.

Philosophers often introduce views in epistemology by articulating broad underlying assumptions about the domain. A quick search for “epistemic

³⁰ See, for instance, Jenkins 2007; Berker 2013a; Berker 2013b; Greaves 2013; Ahlstrom-Vij and Dunn 2014; Goldman 2015; Wrenn 2016; Friedman 2018.

normativity” pulls up multiple papers that open with proclamations on the nature of epistemology: “Epistemology, many think, is about what we should or should not believe. It is, in other words, about *epistemic norms*, i.e. the norms that specify the...truth-related conditions under which we are required or permitted to believe things” (Côté-Bouchard 2016, 3181). Brian Talbot writes

When most epistemologists consider the question ‘What ought one believe?’ or ‘Ought so and so believe such and such?’ they take the answers to these questions to be binding, take them as arising from reasons and rules which are in turn generated by an epistemic end, and take this end as having true belief as a necessary ingredient. (2014, 616)

One especially common way to begin discussion of epistemology’s normativity is to compare it to a more familiar domain of normativity – the ethical domain. Jaegwon Kim begins his “What is Naturalized Epistemology?” by claiming that “epistemology is a normative discipline as much as, and in the same sense as, normative ethics” (1988, 383). Similar claims that draw out the parallels between our talk about epistemic matters and ethical matters can be found in, to name a few, Feldman (2000), Alston (1989), Pollock (1987), Rinard (2017), and Schroeder (2014).

Selim Berker’s 2013 paper “Epistemic Teleology and the Separateness of Propositions” opens with just such a comparison. He begins: “As I see it, the most fundamental question in ethics is ‘What should I do?’” (2013a, 337). He notes that the question can take on different forms depending on one’s preferred way of putting it. We can ask what one has most reason to do, what’s rational for one to do, what one is justified in doing, etc., and we can ask related questions about what one ought to do or how one ought to live narrowly – Should I lie to my friend about where I was last Saturday? or broadly – What sort of life should one live? Meta-ethicists ask the fundamental questions: What does the question mean? What

would constitute an answer to the question and what methodology should we deploy in seeking an answer? Normative ethics, on the other hand, aims to supply answers to our substantive questions about what to do. Berker continues: “I think it is useful to view epistemology as having a parallel structure. On this way of conceiving of the discipline, the fundamental guiding question in epistemology is ‘What should I believe?’” (Ibid., 338). Like in ethics, we can ask the question in terms of rationality, warrant, justification, or what we have most reason to believe.

So far I’ve tried to point out a few basic assumptions that ground contemporary epistemology: epistemology is a normative domain, it’s often thought to be guided by the normative question about what one ought to believe,³¹ and its normativity is analogous to other familiar kinds of normativity like moral normativity. Berker’s pair of 2013 papers³² makes all of these assumptions and in addition is a good jumping-off point to think about the recent interest in examining whether or not the structure of epistemic norms is teleological. In the next section, we’ll take a look at the analogy between the fundamental projects of ethics and epistemology and the import this analogy has for understanding epistemic normativity.

1 The Ethics and Epistemology Analogy

After his initial sketch of the analogous guiding questions of ethics and epistemology, Berker writes:

In what follows, I want to use this analogy between normative evaluations of actions and normative evaluations of beliefs to tease out, and then argue against, a certain strain of thought that seems to have become an article of faith in much recent epistemological

³¹ “What should I believe?” is not taken to be the only question that guides contemporary epistemology. Virtue epistemologists, for instance, give more weight to the question “What sort of agent should one be?”

³² Berker 2013a, and 2013b “The Rejection of Epistemic Consequentialism.”

theorizing. According to this strain of thought, what distinguishes epistemic norms from other sorts of norms (prudential, moral, and so on,) is that epistemic norms are guided by a distinctive set of goals. Most often these epistemic goals are taken to be the twin goals of acquiring true beliefs and avoiding false ones. Sometimes, though, the list of epistemic goals is broadened to include other items, such as the acquisition of knowledge, or coherent belief-systems, or understanding, or wisdom. But regardless of what the list of epistemic goals looks like, the guiding idea behind this strain of thought is that all other normative notions in epistemology are ultimately explicable in terms of how well the objects of assessment conduce toward, promote, or otherwise subserve these epistemic goals. (2013a, 339)

Berker goes on to attempt to convince us that “this consequentialist/teleological approach to normative epistemology is positively misguided” (Ibid., 340).³³

Consequentialism, in ethics, is a family of theories united – at least roughly - by a certain structural feature: that the *evaluative* determines the *deontic*. In other words, the good – however it is identified – determines in some way what ought to be done. A simple version of a consequentialist theory in ethics is hedonistic utilitarianism. Pleasure is good, pain is bad, and nothing else has value or disvalue as an end. A (maximizing) hedonistic utilitarian would hold that some action is right insofar as it produces on balance the most pleasure and the least amount of pain. Consequentialism, in epistemology, is likewise a family of theories united by the same structural idea. An epistemic consequentialist begins with an evaluative

³³ Berker here is using ‘teleological’ and ‘consequentialist’ interchangeably. Later in the paper I show that there is some ambiguity in the deployment of these terms in recent work in epistemology. I point to work where ‘teleological’ is used to pick out theories whose norms are grounded in or explained by value. This is different from consequentialist theories, which typically assume the rightness of a particular response to value – usually maximizing, sometimes satisficing, but always *promoting* relevant value in some way. On this classification, consequentialism would be a breed of teleology. But this use of teleology doesn’t allow for a sufficient contrast with deontological views, so I will follow Berker in focusing on the *conducing* or *promoting* aspect of teleology and consequentialism, thus using the terms interchangeably.

claim, say, that believing important truths is good and believing falsely is bad.³⁴ The deontic claim follows: we ought to conduct our epistemic lives so that we believe many important truths and refrain from believing falsehoods.³⁵

What's important to notice about the structure of teleological theories is that in identifying the various goods that have final value and constructing a theory of overall value out of the theory of final value there can be conflicts or competition amongst final values. These conflicts can arise due to competition between different species of final value, say between keeping a promise and not causing harm, or competition between the same species of final value, say between keeping one of two mutually exclusive promises. The epistemological equivalent of such tradeoffs might (for an epistemic value pluralist) be a conflict that arises between adopting a belief that is highly explanatory versus one that is more testable, or a conflict that arises due to the risk of believing falsehoods when one sets out to believe significant truths.

Thus there arises a question about how to determine value overall. Teleological theories need a way of "trading off" these values against one another in order to determine a final verdict.³⁶ Say that one accepts a tradeoff when one sacrifices something of final value for something of greater final value. Because teleological theories in ethics can have conflicts between final values and must settle on an overall value verdict, they countenance tradeoffs between these values in determining what action should be done. Berker thinks that if the right

³⁴ Though epistemologists nearly uniformly take epistemic good to involve true belief, some are interested in knowledge, understanding, etc. So long as the picture involves a specification of the good and a promotional relationship towards the good, the formulation of the epistemic good won't matter for the purposes of this paper.

³⁵ In what follows, I will point out some important differences between the behavior of consequentialist norms in ethics and epistemology. There are other differences in consequentialist ethics and epistemology theories also worth noting that I will not discuss – for example, it's common for ethical consequentialists to hold that agents should be impartially concerned with producing goodness for all affected beings. Epistemic consequentialists, on the other hand, focus on the *personal* value of believing (significant) truths. So unlike the ethical consequentialist, epistemic consequentialists apparently do not conceive of epistemic goals as concerning the impartial production of epistemic value.

³⁶ The most common way of determining overall value, familiar in consequentialist theories, is aggregation – the value of available options are summed and then compared.

approach to epistemology were teleological, then we should also expect the correct theory of what to believe to countenance tradeoffs between epistemic goods. But, he claims, no epistemologists – even ones who explicitly embrace teleological frameworks, are willing to countenance such tradeoffs in the epistemic domain, so the teleological approach to epistemic normativity must be incorrect.

Ethical consequentialists often face the accusation that their theories countenance morally problematic tradeoffs. Consider familiar examples in which an unruly mob could be pacified by the scapegoating of an innocent, or a patient at the hospital for a well checkup could be used to save the lives of a handful of people in need of organ donations. At least on the face of it, these seem like cases in which the value to be gained by sacrificing the interests of the one for the many would be greater if we were to act in a way that appears to many of us intuitively morally problematic. Thus non-consequentialists argue that it's not always true that the evaluative determines the deontic. Berker refers to this countenancing of tradeoffs as one of teleology's greatest vices. He notes that Rawls identifies the problematic feature in the above examples as a failure to respect the "separateness of persons" (1999, 16). The problem, as he sees it, is that the wellbeing of the many is aggregated. The people whose wellbeing is lumped together fail to be treated as individual persons with individual perspectives. We frequently think that it's impermissible to sum the many small burdens that a large group of people might experience in order to outweigh a giant burden that we could impose on a single person. If the many people were one person, the larger burden could legislate in favor of imposing the smaller burden.³⁷ But in situations similar to the one above, there is no single person who would carry the weight of the entire burden, and so burdening one for the benefit of the many seems impermissible.

Similar tradeoffs, Berker notes, can occur in epistemology. Take Jane, who has no evidence that God exists. But if she forms the belief that God exists despite

³⁷ Interesting to note, however, is that Rawls in fact endorses a few different kinds of tradeoffs with respect to wellbeing. His Second Principle of Justice aggregates wellbeing over social classes and endorses a tradeoff between limits on gains in wellbeing for the well-off for the sake of benefits to the less well-off. And his General Conception of justice permits tradeoffs between well-being and rights.

her lack of evidence, she'll be more likely to receive a large grant from a religiously affiliated funding agency. The grant will ensure that she goes on to believe many more interesting true propositions than she would believe otherwise (Fumerton 1995, 12). Or take John, a mathematician suffering from an illness that doctors say would ordinarily kill him in two months' time. He's on the cusp of proving the Continuum Hypothesis, but he needs six more months. If John stubbornly clings to his belief that he will recover, it will in fact significantly raise his chances of surviving long enough to complete his proof and derive many other interesting truths about set theory (Berker 2013b, 369)³⁸.

If one were to hold a simple epistemic consequentialist view, say, on which interesting true beliefs were valuable and false beliefs were disvaluable, it would seem – even putting aside possible prudential benefits to exclusively consider the epistemic benefits - that Jane should believe that God exists, and John should believe that he will recover. After all, Jane and John would go on to believe more interesting true things and thus be in states that have much more epistemic value. But the problem is that even if we stipulate that believing that God exists or believing that you'd recover would allow you to accrue more of what we consider epistemically valuable, we don't thereby judge that those beliefs are justified, warranted, rational, or the likes. Even if a particular false belief opened the door to many other important truths, it doesn't seem to make that view any more reasonable. Epistemologists widely agree: Klausen claims that “on reflection,” thinking of epistemic norms as consequentialist in this way, is “preposterous” (2009, 165). After all, how could anything having to do with the epistemic standing of *other beliefs* have any bearing on the epistemic standing of a particular belief – save by their being epistemically or logically connected, which by hypothesis, is not the case? Feldman writes that in considering whether to believe, disbelieve or suspend judgment on proposition p , one is to consider only these options and only the end of getting at the truth about p

³⁸ Berker's choice of the Continuum Hypothesis for this example is not ideal, since its status is known for familiar set theories. Sub in your preferred not-yet-proven hypothesis, such as the Riemann Hypothesis.

...and irrelevant to this judgment are the long-term consequences of adopting the belief. It is the truth about p now that matters. Thus, if believing something now would somehow lead me to believe lots of truths later, that long-term epistemic benefit is also irrelevant to it. (1988, 249)

Berker thinks that given the veritistic teleologist's twin goals of believing true propositions and not believing false propositions, that we should expect tradeoffs to occur ordinarily between these fundamental aims. But what he thinks has not been sufficiently appreciated is that veritistic theories must allow for tradeoffs with respect to different propositions: the theories must balance the value of believing one proposition if it is true/disvalue of believing it if it is false with the value and disvalue associated with all other propositions. "But such tradeoffs", Berker writes, "are beyond the pale" (2013a, 363). He continues:

When determining the epistemic status of a belief in a given proposition, it is epistemically irrelevant whether or not that belief conduces (either directly or indirectly) toward the promotion of true belief and the avoidance of false belief in *other* propositions beyond the one in question...When it comes to the evaluation of individual beliefs, it is never epistemically defensible to sacrifice the furtherance of our epistemic aims with regard to one proposition in order to benefit our epistemic aims with regard to other propositions... (Ibid., 365)

Just as the ethical consequentialist ignores the separateness of persons if her theory endorses chopping up one to save five or throwing one in jail to pacify an unruly mob, Berker thinks the epistemic consequentialist goes wrong when she ignores the epistemic separateness of propositions. Berker is not the first to note

the requirement to consider the justificatory status of propositions separately. In an influential formulation of epistemic normativity, Chisholm writes “...every person is subject to a purely intellectual requirement – that of trying his best to bring about that for every proposition *h* he considers, he accepts *h* if and only if it is true” (1977, 14). Berker locates the teleologist’s problem in her failure to respect the epistemic separateness of propositions in his 2013a. In “The Rejection of Epistemic Consequentialism,” though he continues to find ignoring the separateness of propositions problematic, he expands his view to encompass the more general issue of teleology’s *conducting* or *promoting* relation. In what follows, I will talk more about how conducting or promotion relates to teleology and to epistemic evaluation more generally. But first I will look further into the comparison between tradeoffs in the ethical and epistemic domains.

Is Berker right that ethical non-consequentialism has something to teach us about epistemology? I think he is, but what I think it can help elucidate is not the same as what he thinks it shows us. Curiously, he concludes that epistemic norms must be non-teleological, but in none of his three papers³⁹ on the topic does he offer a satisfying *explanation* of the objectionable feature of teleology. What *makes* countenancing tradeoffs bad in the epistemic domain? *Why* is conducting or promoting value the wrong way to think about epistemic normativity? In fact, I think the specific conclusion Berker draws – that an analogy with ethics can show us that epistemic norms are non-teleological – is not very helpful. I think this for two main reasons. The first is that the analogy between separateness of persons and separateness of propositions elides important differences in the permissibility of tradeoffs in the ethical and epistemic domains. I will argue that ignoring these differences masks a further important difference in how we should understand the *value* and the thus the *normativity* in play in the respective domains. The second is that the purported target of Berker’s project is epistemic normativity – all of it – and by identifying the question that underlies the epistemic domain as the

³⁹ See also Berker’s “Reply to Goldman: Cutting up the One to Save the Five in Epistemology” (2015).

question “What should I believe?” he casts its purview maximally widely. But I think that the epistemic intuitions Berker uncovers don’t in fact spell out answers for what characterizes epistemic normativity, especially if that normativity is taken to be the normativity of *belief* - as he suggests when he casts epistemology’s target question as “What should I believe?”. I think his results actually characterize a narrow subsection of what can plausibly be taken to count as the normative domain of the epistemic -- a domain which is itself only a cross-section of what we might think of as the normative domain of the entire state of belief. In this section I’ll draw attention to a couple of important differences in how tradeoffs are treated in ethics and epistemology, and in the sections that follow, I’ll use this observation alongside Berker’s conclusion to draw out some import for our wider understanding of epistemic normativity.

2 Tradeoffs: Interpersonal and Inter-Propositional

Berker is aware that countenancing tradeoffs is not always a vice in ethics. After all, not all interpersonal tradeoffs ignite the separateness of persons objection. However he thinks it’s the possibility of cross-propositional value tradeoffs that spells doom for the epistemic consequentialist. How, then, do interpersonal and inter-propositional tradeoffs relate? There are different tradeoff structures that could make this point equally well, so I’m simply going to select two. We’ll look first at a kind of tradeoff that does not run afoul of the separateness of persons charge. Think of a very ordinary kind of tradeoff in which some person takes on some small disvalue in order to help another person. Call this an *ordinary cross-person tradeoff*:

OCP tradeoff: A incurs some small harm in order to promote a greater benefit (/prevent a greater harm) to B.

OCP tradeoffs are routine. You and I are participating in an amateur cycling race, cycling together on a secluded portion of the racecourse. If you crash and I'm the only one around, I'm presumably obligated to stop and help you. The benefit of your receiving appropriate medical help outweighs the disvalue I incur by losing time in the race. Ethical consequentialists and deontologists agree that OCPs are often permissible or required. There are many ways OCP tradeoffs can come about and can include (infinitely) many configurations in terms of number of people and benefit/harm allocation. By stipulation, OCP tradeoffs don't raise the separateness of person worry because the relevant sort of cross-person aggregation of consequences isn't taking place – the assessment of the total harm or benefit doesn't involve the questionable lumping of multiple peoples' individual perspectives into one.

More interesting to think about, perhaps, is a different kind of cross-personal tradeoff – one that more closely mirrors the structure of the tradeoffs Berker describes. Think of a far less ordinary kind of tradeoff in which some person incurs great (maybe maximal) disvalue in order to benefit a great number of people. Call this kind of case a *threshold tradeoff*:

Threshold tradeoff: great disvalue to A, very large benefit (/prevention of a very great harm) to a very large group of people

Suppose terrorists have hijacked a large plane and we have extremely strong reason to think that they will blow it up, killing 380 people, unless we kill one person. Interestingly, in such a case most deontologists agree that we ought to kill one person to save the 380 people. Nagel writes “[D]eliberately killing an innocent is impermissible unless it is the only way to prevent some very large evil (say the deaths of fifty innocent people). Call this the threshold at which the prohibition against murder is overridden” (2012, 62).

Make a distinction between absolute deontology and moderate deontology. Whereas absolute deontologists find the violation of deontological constraints

always and everywhere impermissible, thus allowing for the occurrence of morally catastrophic events that could be prevented by a single instance of such a violation, moderate deontologists permit violations of constraints or require promotions of the greater good under extreme circumstances. Though deontology is often characterized by its contrast with consequentialism and that contrast is clearest in its absolutist form, it's exceedingly difficult to find examples of philosophers who actually hold such a view. The best example seems to be common renderings of Kant, who famously claimed "Better the whole world should perish," than that injustice be done (Ladd 1965, 100). As Alexander and Moore point out, some deontologists' treatments of threshold cases seem to simply evade the question – Anscombe says that true moral agents can't even think of such real or imagined cases because the mere thinking of them is immoral, and Williams thinks that such cases are "beyond morality" and beyond reason.⁴⁰ Notice, though, that an absolutist view like Kant's in which some deontological constraint can't be violated or greater good promoted even in order to prevent thousands of occurrences of the very act the constraint prohibits is often taken to devolve into a *reductio*. And failures to treat head-on threshold cases, like Anscombe's and Williams', seem to share the weakness of viewing our moral reasoning as ill-equipped to deal with sufficiently difficult cases – even theoretically – and to resign ourselves to the belief that our moral reasons simply run out.

Moderate deontology, on the other hand, is a quite well-represented view. Thresholds beyond which deontological constraints may or should be violated or greater good ought to be promoted are positively represented or endorsed in – to name a few – Kagan, Nagel, Ross, Moore, Alexander, Scheffler and Kamm.⁴¹ Michael Moore (2010) likens the moderate deontologist's use of thresholds to a dam. Water rises behind a dam until it reaches the threshold height of the dam, after which the water spills over. Moderate deontology's thresholds mark the

⁴⁰ Anscombe 1958, Williams 1973, and Alexander and Moore 2016.

⁴¹ Ross 1930; Scheffler 1982; Kagan 1998; Alexander 2000; Kamm 2007; Moore 2010; Nagel 2012.

point at which constraints can be violated because of the spillover of negative consequences.

By this point the differences in ethical deontology's treatment of tradeoffs and epistemology's purported treatment of tradeoffs should begin to become apparent. Remember that Berker likened failure to respect the epistemic separateness of propositions to failure to respect the ethical separateness of persons. Besides more obvious features which lessen the aptness of the ethics/epistemology tradeoff analogy - such as the fact that unlike in the epistemic case of propositions, in ethics the separateness of persons worry arises from the normative status of persons who we often take to make reasonable claims on us - there are two important things to notice as we consider the fruitfulness of this analogy. As we saw in the case of OCP tradeoffs, there are cross-personal tradeoffs that ethicists of all stripes endorse that don't raise the separateness of persons worry. And for cases that do raise the separateness of persons worry, like the threshold cases, the number of ethicists that fail to endorse the tradeoff appears to be vanishingly small. But the striking thing about the epistemic domain is that it appears as if it licenses *no tradeoffs whatsoever*. Not only do our epistemic intuitions seem to mark epistemically unjustified the adoption of a particular belief for which we have insufficient evidence in order to secure great future epistemic gains, but they (perhaps to an even greater extent) legislate against the epistemic version of the everyday OCP tradeoffs - say, the adoption of one evidentially unjustified belief in order to believe another more important true proposition. As Berker writes

[...] if the consequentialist/teleological approach to normative epistemology were the correct one, we'd expect that the correct theory of what we should believe would also countenance trade-offs between goods - in this case, trade-offs between epistemic goods. However no one - not even those epistemologists who most willingly

embrace the consequentialist/teleological framework – is willing to countenance all such tradeoffs in the epistemic case. (2013a, 340)

So, Berker characterizes epistemic consequentialism as falling prey to a problem analogous to one that plagues ethical consequentialism. Yet the kind of non-teleological theory that would correctly capture the received fundamental epistemic norm of believing the truth about a particular proposition is so extreme/absolute that it's doubtful whether any contemporary deontologist actually endorses such a view. I think it's worth considering why this is true.

3 Teleology and its Relation to Value

Though my ultimate goal for this section is to argue that the extreme deontological picture that Berker supports in epistemology yields some interesting conclusions about epistemic value, I want to first try to gain a clearer understanding of teleology's background relation to value in a given evaluative domain.

In justifying the thresholds at which deontological constraints should be infringed, moderate deontologists typically appeal to the massive amount of *value* at stake. Describing moderate deontology, Heuer writes: "...consequences of a certain kind—catastrophic ones—must be avoided even at the price of violating restrictions, or the restriction can be defeated by consequences whose disvalue is out of proportion with the particular value that the restrictions protect" (2011, 236). This can seem to some like a feature that threatens the deontological nature of the theory. However we shouldn't think that deontologists don't acknowledge the value of consequences that can be promoted, but rather that they recognize *more* features of situations that are normatively valuable than only the value of the consequences. Moderate deontologists typically recognize a constraint against doing harm, but in extreme cases in which enormous amounts of harm can be

prevented or value promoted by doing harm, they believe that the value at stake must be respected (Kagan 1997).⁴²

I think this point is actually related to a confusion epistemologists have in terms of how to think about teleological views. I noted in an earlier footnote that I would make a clarification about how I was deploying the term. Berker uses the terms ‘teleological’ and ‘consequentialist’ interchangeably and means to focus on the *promoting* or *conducing* aspect of the theories. But it’s common for epistemologists writing on this issue to make other kinds of unhelpful terminological choices or characterizations. One such example is Friedman’s “Teleological Epistemology”. She begins by noting that it’s typically thought that epistemic states like knowing, understanding, truly believing, etc. are valuable. They’re states that are good to be in and we often aim or want to be in them. That these states are thought to be good or valuable is what in turn grounds and explains our epistemic norms. Like Berker, Friedman is arguing that this value-based picture, which she calls a teleological conception of epistemic normativity, cannot be correct. She writes “Let’s call norms that are grounded in or explained by values – either valuable states or states of affairs or the things that actual epistemic subjects value – teleological.” She sees the teleologist as making two main claims:

- (1) A value claim: knowledge [or true belief, etc.] is valuable/valued.
- (2) A grounding claim: that knowledge is valuable/valued grounds or explains our central epistemic norms; these norms are teleologically grounded (2018, 3).

⁴² This feature has caused some consequentialists to argue that moderate deontology is actually just a form of consequentialism. I imagine an absolute deontologist might argue the same thing. Though I think moderate deontology is in fact perfectly coherent - defining consequentialism as a theory that articulates what we should do in terms of what promotes desired (satisficing, maximizing, etc.) consequences, and deontology in terms of invoking constraints on that promotional relation, is, I think, a good way to go - the terminological dispute doesn’t concern me here. My aim is merely to draw attention to the fact that deontological views of the kind analogous to what Berker diagnoses in epistemology hardly exist.

And later: “My thought is that if a norm is teleologically grounded, then that explanation will go by way of some valuable or valued state of affairs” (Ibid., 4). Of course “go by way of” leaves room for a number of ways this grounding could work. But to my mind, this characterization of teleological views is far too general - it doesn’t actually suffice either to distinguish teleological views from non-teleological/deontological views or to lend focus to the features that she and Berker actually find problematic. My main complaint is that one can endorse a view that seems robustly *value-based* – one in which valuable states are identified and thought to explain the norms of that domain, thus satisfying her two descriptors of teleology, yet behaves just like familiar non-teleological views, i.e. by in particular cases issuing recommendations that recognize constraints on the promotion of the good. We might even think of the most stereotypically deontological view, Kantian ethics (at least as articulated in the beginning of the *Groundwork*), as beginning with the idea that only the good will has final value and proceeding to treat all norms as born from and explained by that fundamental value.

It seems that there are ways to respond to fundamental value that do not involve maximization or any other kind of direct promotion, but Friedman appears to disagree with this thought. She writes that once we see subjects as valuing some end or other or as in pursuit of some valuable 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. If one is trying to acquire money, and one is faced with a choice between acquiring one dollar, ten dollars, and one hundred dollars, one opts for the latter on pain of irrationality...This is to say that there is pressure to say that there is a norm in [teleology] that demands we maximize or optimize with

respect to our knowledge acquisition. This sort of norm will say that subjects should always come to know as much as they can. (2018, 9)

Friedman writes in a footnote on this passage “I hope that it’s clear why the teleologist is pushed towards demanding the relevant kind of value maximization. In general once we think of subjects as in pursuit of something of value or that they value then it’s natural to think that they should get as much of it as they can (other things equal).” And in one more passage, “If we think that epistemic subjects are or ought to be trying to know more, then plausibly they will be expected to act so as to maximize with every judgment” (Ibid. 12). Friedman doesn’t say anything about what “other things equal” means, but it seems to me like other things would very rarely or never be equal. First off, the epistemic domain is but one evaluative domain. So long as we recognize the existence of other evaluative domains, (morality, prudence, aesthetics, etc.,) we will acknowledge that the values in these domains can conflict with each other. My desire to acquire more knowledge can be trumped by my need to run to catch a bus. Our own cognitive architecture can also impose some limits on what shape these norms take – given storage and access constraints, it simply isn’t the case that more knowledge is always better. Given these limits, even if we understood the fundamental value of the epistemic domain to (more plausibly) consist in *interesting* or *important* true belief or knowledge, it just isn’t the case that we would be cognitively best served by maximization with respect to that goal.

The relationship between the fundamental values of an evaluative domain and the relations that subjects bear to those values is an interesting one. As Friedman, Berker, and many other epistemologists point out, veritism – the thought that the value of true belief (understood in some way) is fundamental to the epistemic domain – is nearly ubiquitous. But just because some value is fundamental to an evaluative domain and explains the value of all other values of that domain, does not itself dictate that one must bear an instrumentalist or *promoting* attitude toward that value. Sylvan makes this point nicely in his 2018

paper, “Veritism Unswamped”. He starts with the thought that “in any evaluative domain, some values will be more fundamental than others, in the sense that the value of everything else in that domain is explained by their value” (382). He acknowledges that there are many values in the epistemic domain – rationality, knowledge, coherence, accuracy, justification – but that not all of them can be taken to be equally fundamental. “We admire some of them from the epistemic point of view because we admire others from the epistemic point of view” (Ibid.). Like other veritists, he thinks that all epistemic values are explained by the fact that we value accurate belief. But he goes on to claim that both opponents and proponents of this view understand it in an unreasonably narrow way and are widely mistaken about what follows from it.

Epistemologists typically assume what Sylvan calls “instrumentalism about derivative value”. This is the thought that there is only one kind of way in which we can explain one epistemic value in terms of a more fundamental epistemic value. They assume that the explanation must “proceed by invoking *instrumental relations*, so that for any X, X is derivatively epistemically good only by i) tending to produce fundamental epistemic goods, or ii) being the product of something with feature i)” (2018, 382). On this view, all derivative epistemic value is either i) production value, or ii) mere product value. Teleologists, and Friedman in addition, think that the proper way to respond to value in the epistemic domain is to accept what Sylvan calls the “teleological answer”: “For any fundamental value V, the only basically proper way to value V is to instrumentally promote V” (Ibid., 394). But this view seems incorrect. After all, as Nozick (1984) points out, there are lots of ways to value things – we can respect them, as we respect the good will in Kantian ethics, we can be loyal or dedicated to them, love them, etc.. The teleologist seems to believe there’s only one fitting response to fundamental value: to promote it. As Sylvan notes, if we were to take that view towards the value of, say, friendship, we would end up spending all of our time trying to make new friends and produce more instances of friendship, by say, procuring friends for

other people, than showing the proper regard for the friends we already have. *Dedication* to our friends is a prime way to express the value of friendship.

In beginning this foray into the question about what makes a view teleological, I said that I thought that Friedman's characterization of teleology "didn't suffice to distinguish teleological views from non-teleological/deontological views or to lend focus to the features that she and Berker actually find problematic." Now we're in a position to see why. Both teleological and non-teleological views can be normatively structured around fundamental values. Though non-teleological views recognize other norms beyond the consequentialist's value promotional norms, they too almost always include norms that sometimes require the promoting of fundamental value, such as when value thresholds have been surpassed. So rather than attempt to distinguish teleological from non-teleological theories by way of a value and a grounding claim, I think it's better to focus on the particular *conducting* relationship that subjects bear to value in the domain.

For the purposes of this paper, then, it will be best to understand teleology in the same sense as consequentialism - as the view that not only is value fundamental, but that the proper response to value is to bring it about - for the teleologist all value is "to be promoted" and all disvalue is "to be prevented."⁴³ Thus on this view teleology includes what Sylvan calls the narrow "instrumentalism about derivative value assumption", whereas a non-teleological veritist might believe that there are other ways to respect fundamental epistemic value other than to promote it.

⁴³ Berker 2013a, 343, citing Pettit 1997. Also see Korsgaard 1993, and Scanlon 1998. Note that this understanding of teleology is different from another conception of teleology which focuses on *telos*, or proper ends or functions. In this sense, teleological views include views like Aristotle's that looks to the proper ends of humans but does not embrace impartial maximization. Though ultimately it seems like it would be best to have a distinction in place between value-grounded views in general and views that entail aggregative responses, for continuity I'm opting for Berker's synonymous usage of "teleological" and "consequentialist" while pointing out the drawbacks of failing to recognize the distinction.

4 Epistemic Value and its Relation to Normativity

We're now in a position to try to draw out some implications of Berker's argument for our understanding of the nature of epistemic value. Berker thinks that teleological theories countenance tradeoffs. Since epistemic norms appear not to countenance tradeoffs, the correct view of epistemic normativity cannot be teleological. What I want to suggest is that the fact that epistemic normativity does not countenance tradeoffs should have direct implications for our conception of epistemic value. My thought is that there is a general relationship between countenancing tradeoffs and value. It's a simple one: tradeoffs are licensed in light of reasonable responses to value. The idea in broad form is that there is a range of reasonable responses to value in a normative domain. We might think (in keeping with moderate deontology) that norms can be directed toward the promotion of value in some cases (and respect constraints on value promotion in other cases) or, in keeping with maximizing consequentialism, that they can be directed toward the promotion of value in all cases, but we can't take them to be directed toward the promotion of value *in no cases*. For if that were the case, I claim, we would be mistaken to think of what explains the norms as value to begin with. And that's what I think is in fact happening in a normative domain in which final values can conflict, yet which countenances no tradeoffs.

What we've seen so far is that almost every plausible non-teleological ethical theory countenances some kinds of cross-personal tradeoffs. Epistemic norms, on the other hand, appear to countenance no tradeoffs whatsoever. As we saw earlier, when non-teleological views in ethics recommend tradeoffs, it's generally because they are taken to be the proper response to great value promotion or harm prevention. But when we learn that there are no tradeoffs that are deemed acceptable in determining whether one ought to believe some proposition for which one has insufficient evidence, we learn that there's no possible amount of epistemic value the promotion of which would make it that we ought to believe *p*. We have the same absolute obligation towards each and every

proposition, no matter how insignificant it is or how much value is at stake. If in every instance in which it would be possible to secure a great amount of value in some domain through a single norm infringement, the norms of that domain instruct us to proceed so as not to secure that value, I claim we have reason to conclude that the relevant value isn't behaving like value at all – most likely we were wrong to think that that domain was normatively governed by value to begin with. After all, if one conceives of something as valuable, it seems that what *makes it a value* is that it is desired or desirable. But what we see in the epistemic domain is that no amount of value of the thing that's supposed to ground the value could dictate that we should secure that value.

How are we left to think about the normativity of the epistemic domain? As is common in ethics and value theory, I take there to be a necessary connection between normativity and value.⁴⁴ At first glance, we might think this view prejudices us toward the family of teleological or consequentialist views. But as I tried to show earlier, we should take those views to assume a particular instrumentalist promotional conception of an agent's permissible response to value. On my view, the normativity of deontological/non-teleological and teleological norms alike comes about in relation to the value that undergirds the domain. In the ethical domain we might take this to be the value (broadly construed) of pleasure, or a more nuanced notion of wellbeing, or the good will. In the epistemic domain, the veritist tells us to conceive of true belief as the fundamental epistemic value that is supposed to give rise to and explain the content of our epistemic norms.

But now consider the extreme deontological conception of epistemic normativity that Berker claims fits our pre-theoretic intuitions. The structure of Berker's argument is that epistemic norms don't countenance tradeoffs. Teleological theories countenance tradeoffs, so epistemic normativity is non-teleological. My response is as follows: Value is such that we might reasonably

⁴⁴ See for discussion Zimmerman 2015.

respond to it by promoting it some of the time or all of the time, but not none of the time. If a theory claims that the norms of a given domain, in which there are possible conflicts between final values, do not countenance any tradeoffs – in other words, in “trading off” final values against one another in order to determine overall value, the norms are structured so that it is impermissible to violate a particular norm in any circumstance, no matter the value that could be gained, then we should infer from that that value is not in fact what gives rise to the norms of that domain. Since there’s a necessary connection between normativity and value, this doesn’t bode well for the normativity of epistemic norms. In what manner should we understand this necessary connection?

Since we’re exploring the contents of epistemic norms, we might have taken it as a given that epistemic norms are *normative* – in other words that they give rise to reasons or oughts that govern what we should do or how we should be. But we must make a distinction between “genuine” or necessary normativity and mere norm-relativity. Broome appeals to this distinction:

[I]n one sense ‘normative’ simply means to do with norms, rules, or correctness. Any source of requirements is normative in this sense. For example, Catholicism is. Catholicism requires you to abstain from meat on Fridays. It is a rule and it is incorrect according to Catholicism to eat meat on Fridays. So Catholicism is normative in this sense. But I do not use ‘normative’ in that sense. In my sense, it means to do with ought or reasons. Given a rule or a requirement we can ask whether you ought to follow it, or whether you have reason to do so. (2007 162)

As Côté-Bouchard puts it,

If a norm *N* forbids ϕ -ing under conditions *C*, then trivially, ϕ -ing under *C* is forbidden, incorrect, wrong, or bad *relative to* the standard

set by N. However, the descriptive *norm-relative* question of what is permitted or required according to N is distinct from the *normative* question of what there is *good* or *genuinely normative* reasons to do. (2016, 3182)

Earlier I wrote that there is a necessary connection between normativity and value. We can think of normativity as encompassing two families of concepts: evaluative concepts like ‘good’ and ‘bad’, and deontic concepts like ‘right’, ‘wrong’, ‘permissible’/ ‘impermissible’, etc. There are two main positions available in the literature that articulate how these concepts relate to each other in giving rise to the oughts that govern what we should do and how we should be. We could think that value facts explain reason facts: one might hold that the epistemic value of true belief is what explains the fact that I have epistemic reason to believe that *p*.⁴⁵ Second, we could think that the reasons we have explain facts about value: one might think that the fact that it’s good for me to believe *p* when I have decisive evidence that *p* is explained by the fact that in that situation, I have reason to believe that *p*.⁴⁶ In some sense, I think that the difference between these two views is a terminological dispute: we could aptly characterize normative structure by treating reasons as fundamental building blocks out of which we explain value or vice versa. That said, I think the picture on which we take value to explain reasons fares better in the literature, is naturally more intuitive, and enjoys a number of distinctive theoretical benefits over the opposite view.⁴⁷ So though I will proceed by adopting this underlying picture, all that’s required to secure my conclusion that a view of epistemic normativity that countenances no tradeoffs does not give

⁴⁵ See for example Frankena 1968; Stich 1990; Moore 1903 (Baldwin 1993); Reisner 2009; Steglich-Petersen 2011.

⁴⁶ See Scanlon 1998, 78-107.

⁴⁷ See for example Andrew Reisner’s 2009, especially his take on the “inaccuracy objection” – the thought that we lose a lot in the reduction from value to reasons since our value concepts are much finer-grained than our relatively limited set of pro-attitudes – and his thought that the most promising solutions to the Wrong Kinds of Reasons problem get their currency from appeals to independent and more fundamental intuitions about value.

rise to necessary normativity is a necessary connection between either reasons or value and normativity – no matter which family of concepts takes explanatory priority.

On the picture we'll henceforth adopt, reasons come about in accordance with value. It's the value of ϕ -ing that gives rise to our reason to ϕ .⁴⁸ In cases in which there is no value (to promote, to respect, to honor, etc.,) in ϕ -ing, we have no reason to ϕ . In "Evidential Exclusivity, Correctness, and the Nature of Belief" I endorse this same structure and make use of the distinction between "rule-implying norms" and "reason-implying norms". It's a distinction familiar in value theory between norms that act as general descriptive principles or serve to categorize or classify actions into distinct types, and norms that act as prescriptions or governance principles and give us aims to follow. Parfit calls the first type, which are generally individuated by a rule or practice and don't give rise to necessary reasons to follow them, "rule-implying". An example might be the collection of permissible moves in chess. The second kind, which do necessarily give us some reason to follow them, he calls "reason-implying" (2011, 309). An example of a reason-implying norm might be a moral norm, like "one should keep one's promises."

When we consider the picture of non-tradeoff-countenancing, absolutist non-teleology that Berker thinks characterizes the epistemic domain, we see a picture on which there's a rule to believe the truth for every proposition we consider – regardless of whether it's an utterly insignificant proposition that has no value for us and which could be sacrificed for a great number of truths that have great significance for us. Now we might suggest that just because a norm doesn't countenance tradeoffs doesn't mean it's not value-based. Might it not be the case that it's the global value of believing the truth for every proposition that explains the value of the norm and thus gives rise to its normative force? What I've suggested so far is that a norm that doesn't recommend value promotion even in

⁴⁸ See Maguire 2016 for a comprehensive overview of the value-based theory of reasons. He treats this conception as an *analysis* of what it is to be a reason.

instances in which the greatest amounts (infinite!) of value of that domain are at stake isn't behaving like norms grounded in value ordinarily behave and thus doesn't seem to be a norm that is governed by value at all.

Further, contrast cases in ethics: think of a scenario in which a maximizing consequentialist disagrees with a moderate deontologist because she believes in some instance that causing someone harm would generate more value overall, and the moderate deontologist believes that one ought not cause harm. Though we often describe such cases as disagreements about whether to maximize value, it seems like this description is unfair to the deontologist. Presumably this is a terminological dispute, but one that is better captured by locating the disagreement in the *way* in which value is best responded to.

We might think that part of what we mean to do by characterizing consequentialism with the slogan "the good determines the right" is to pick out a difference in *methodology*- an act consequentialist first thinks of the possible outcomes of available acts, compares the goodness that might come from each outcome, and then determines the right act. The deontologist, on the other hand, presumably thinks that the nature and overall disvalue of harming others dictates that we ought to respond to possible instances of harm-causing in a certain sort of way - not by considering in each case whether it could promote the best outcome, but by assuming that we should not cause harm, because we already know that causing harm is a thing that ought not be done. Now in cases in which a great amount of harm could be prevented by causing some amount of harm, moderate deontologists will recommend infringing the norm that prohibits causing harm. But otherwise, the deontologist's mindset differs from the consequentialist's - it's typically not centered on assessing and aggregating the comparative goodness that comes about from particular acts.⁴⁹

⁴⁹ This description of the moderate deontologist makes the view look very similar to rule consequentialism. Whereas it's easy to articulate the difference between absolute deontology and consequentialism, I think it's actually not so easy to articulate the difference between moderate deontology and rule consequentialism.

The way I've represented consequentialists and moderate deontologists has them both responding, albeit in different ways, to the value at stake. Now contrast that to the epistemic deontologist's response to a scenario in which someone knows that by believing some very insignificant falsehood, she could come to believe many significant truths. It doesn't seem like there's any recognition of value at all that goes on in heeding the corresponding epistemic norm. Moral norms, like other genuine norms, are grounded in necessary value. Presumably the value of the norm that prevents causing harm to others consists in the proper way to respect the wellbeing of other humans. This is a kind of value that exists no matter if it turns out that it might be better in some instances to forgo this value in favor of promoting greater value. There is no such necessary value, or so I've argued, in believing truly. The apparent value of true belief is not explained by the intrinsic value of each and every true belief, but rather by the value of getting it right about things that matter to us. Unlike preventing harm, which we might reasonably think always has at least some value, believing the truth about propositions that are entirely insignificant need not have any value at all. So though we might take both the moderate deontologist and the consequentialist to demonstrate reasonable ways to respond to the value of preventing harm, I think the epistemic deontologist's response, in endorsing a norm that recommends sacrificing great amounts of significant truths for an insignificant truth, cannot be taken to be a reasonable response to value.

I've argued that the fact that non-teleological epistemic norms don't countenance any tradeoffs is an indication that non-teleological epistemic norms cannot be taken to arise from or be explained by epistemic value. And I've claimed that because all necessary normative force arises from value, we cannot take the kinds of norms Berker focuses on to give rise to necessary normativity. They are better taken to be rule-implying as opposed to reason-implying norms. Though this rule-implying norm seems to dictate a response to a very particular question – what belief should I now adopt about this proposition? – we shouldn't think of this norm as necessarily conferring normative force. In other words, we need not see it

as supplying us with reason to follow it. So where does this leave us in our quest to understand epistemic normativity?

5 Rethinking Epistemic Normativity

Berker, like many other epistemologists, tells us that epistemology is about answering the question “What should I believe?”. The argument he goes on to make, I’ve tried to show, suggests that we understand epistemic normativity in a way that does not give us a satisfying answer to that question. It’s unsatisfying in two ways: the first is that I don’t think the kinds of epistemic norms he targets really give us an answer to the question of what we ought to believe without further specification.⁵⁰ The second is that it leaves us with a conception of the domain of the *epistemic* as especially narrow and toothless. To put it more exactly: Berker incorrectly treats epistemic norms as the norms of the entire belief state, and his characterization of the epistemic domain as necessarily non-teleological commits us to thinking of the domain as one that isn’t grounded in necessary value or normativity. I think the reason for this is that Berker actually frames the question he asks unnecessarily narrowly. He confines himself, I’ll try to show, to a view that treats the question “What should I believe now about this proposition?” as capable of informing us about the normativity of belief and the contours of the epistemic domain. But this is merely one question that epistemologists might be interested in answering – indeed, it neglects to consider many valuable projects that have long engaged epistemologists - and we might find that different ways of framing the question deliver different suggestions about the norms of belief and how we ought to understand the projects of epistemology. In what follows, I want to try to reshape our view of the focus of epistemology. In doing so, we will come

⁵⁰ I think he unduly privileges one narrow interpretation of that question: “What should I believe about p (now, as I actively deliberate about whether or not p is true)?” As I argued in “Evidential Exclusivity, Correctness and the Nature of Belief” this leaves out other equally relevant interpretations of the question that will sometimes receive different answers.

to see that there's use for epistemic norms other than the non-teleological proposition-centered norm Berker advocates.

Berker writes of the teleological mindset that it is “widely endorsed (and almost always without being explicitly argued for)” and that this “entire approach to normative epistemology is misguided” (2013b, 340). Though he shows the centrality of the deliberative proposition-centered approach to understanding epistemic normativity he does consider norms that have as an evaluative focal point belief-forming processes as opposed to beliefs about particular propositions. Ultimately he rejects such a view as insufficiently aimed at the proper goals of epistemology. In commenting on the scope of his argument, he says that his recipe of concocting problems for views that license tradeoffs applies to “all varieties of epistemic consequentialism that do not go so far as to restrict the conducting relation to the instantiation relation, or (*even more radically*) to give up on the *epistemic evaluation of individual beliefs all together*” (2013b, 379). Berker comes from a long history of interpreting the fundamental epistemic viewpoint similarly narrowly. Chisholm's formulation of the overarching epistemic norm is that it's a purely intellectual requirement to bring it about that for every proposition we consider, we accept it if and only if it's true (1997, 14). Foley restricts the norm to “*now* believing what is true and now avoiding what is false” (1987, 8 my emphasis), and Fumerton tries to show that even this norm isn't restricted enough – that the goal must be articulated as now believing what is true with respect to a *single proposition* (1995, 12).

I think the history of conceiving of this narrowly single-subject (usually first-person), time-now, proposition-centered perspective as the focal perspective of epistemology is in large part due to a background assumption that epistemic norms are simply norms about belief, which is itself normatively or constitutively truth-aimed. I've argued elsewhere that this isn't true.⁵¹ Now I want to argue

⁵¹ See “The Doxastic Account of the Epistemic” and “Evidential Exclusivity, Correctness and the Nature of Belief”. I argue there that, though the phenomenon of “transparency” can explain why we might have this conception of belief, it is merely a feature of how we're apparently constructed to treat

that, since epistemic normativity is not the same as the normativity of the state of belief, we need not be confined to such a narrow perspective in our understanding of the appropriate purview of the epistemic domain. Instead, I think we should conceive of the contours of epistemic normativity and the resulting norms as those that furnish recommendations for how we ought to conduct our epistemic projects – namely those projects that are aimed at inquiring about the truth of some question or collection of questions that matter to us.

Whereas the focal question has largely been taken to be “What should I believe now about this proposition?” we can think about a number of related questions that the epistemic domain can also rightfully consider.

- What should *I / some subject / some group of subjects* believe about
- *this proposition / this collection of propositions*
- *now / at some time / over some range of time in the future*

And though we might think that belief is a mental attitude that is central to our investigation of the proper focal questions of inquiry, we must also ask more broadly what mental and non-mental actions we should take to best succeed in our inquiries. What (sometimes evidentially insufficiently-justified) propositions should we accept for further investigation? Would endorsing some particular evidentially insufficiently-justified proposition raise the likelihood that we will go on to acquire important true beliefs about some topic of inquiry? What habits should we form to be better inquirers?

At this juncture, a common line of objection will be raised. Epistemologists will argue that the norms associated with the questions I’ve raised will be *practical*, not *epistemic* norms. But how should we interpret the force of this apparent objection? Consider again the case in which Jane is determining whether

deliberation, not a feature that gives rise to normative facts about belief, nor one that adequately applies to the full functional profile of belief.

to believe in God against her evidence. We can contrast the epistemic with the non-epistemic benefits of adopting this belief. We took the epistemic benefits to include the many significant true beliefs she would go on to have if she secured funding from the religious funding agency. A non-epistemic benefit might include that a belief in God would lessen her fear of death. So in this sense there's a straightforward distinction in play – we can characterize Jane as believing for epistemic reasons (in at least one sense) because her belief is directed toward the pursuit of truth. It has as its goal the end that characterizes the epistemic domain.

But some will nonetheless say that a norm that licenses adopting a belief for its benefits is a *practical norm* aimed at an *epistemic end*. I take the root of such an objection to be that the norm is about *action* and not *belief*. In “The Doxastic Account of the Epistemic” I wrote about the problems with understanding “epistemic” as synonymous with “belief”. But I think there are further positive reasons to shift our understanding of the epistemic domain to one that is characterized by the pursuit of significant truth (or knowledge or understanding). This is, after all, how other familiar normative domains are oriented – we can think of aesthetic norms as oriented toward beauty, moral norms as oriented toward moral goodness, etc. And once we give up on the idea that epistemic norms *just are* the norms of the belief state in favor of a conception of epistemic norms as norms that are directed toward the pursuit of truth, then we can see that we will need to recognize more means to achieving the epistemic end than just belief adoption.

One observation we will make as we expand the purview of the epistemic domain is that so far we've talked about norms of belief adoption in an unabashedly instrumentalist way. We consider whether we should or shouldn't believe some proposition in order to secure some further benefit. Both teleologists and non-teleologists about epistemic normativity often proceed as though such consideration is appropriate – that it makes sense to ask ourselves whether we should believe or refrain from believing and accept the answer as if it gives us a directive for what we go on to *do*. But many will have already objected to a much

more basic aspect of this picture, which is that it doesn't make sense to treat belief adoption in this instrumentalist way because belief adoption is not something we *do* so much as something that *happens to us*. Resistance to this instrumentalist conception can come from two different but related worries. One is that as a descriptive fact it seems that we simply don't have the kind of voluntary control over belief that makes instrumentalist norms fitting. The other is that norms are prescriptive – they tell us what to do or how to be. But it's commonly held that it's not the case that we *ought* to do something or that we *ought* to be a certain way if it's not within our power to do that thing or be that way.

Though I in fact think that both of these worries about doxastic voluntarism are overstated,⁵² I also think that epistemologists' insistence on confining the evaluative focal point to the very question that it makes least sense to ask given our desire to deploy instrumentalist norms – in other words confining our focal point to current active deliberation about a proposition, the scenario in which we are most forced to take into account exclusively evidential considerations - is perplexing.

Whereas it has been common to respond to prescriptions for instrumentalist belief adoptions with rejoinders like “But that's not really *belief*, that's *acceptance*!” where the critic means to imply that we've exited the domain of prescriptive *epistemic* norms because we are no longer talking about the state of belief, we can now happily affirm that it is indeed acceptance we're speaking of, and that it's no less epistemic for that fact. And further, that it's much more fitting to talk about acceptance, endorsement (or other more suitably instrumentalist attitudes/actions) than belief, because we're more readily able to accept or

⁵² First, I think that the degree to which classically voluntary actions are under our direct control is overstated and there are many actions that we talk about as voluntary over which we have at best indirect control (See Rinard's “No Exception for Belief” for great examples, and second, I think it makes perfectly good sense to issue normative judgments about departures from ideal states because these are not necessarily condemnations or declarations of blameworthiness that carry implications about responsibility.

endorse insufficiently evidentially justified propositions than we are to believe them in cases of conscious deliberation.

Though we have ample examples that point to the dominance in epistemology of the proposition-centered veritistic notion that governs belief and credence, philosophers have done valuable work on other projects that will help fill out our broadened conception of epistemic normativity. Norms of scientific investigation and inquiry have long been of interest to philosophers of science and have enjoyed renewed interest amongst contemporary epistemologists. Van Fraassen, for instance has argued that it is appropriate for scientists to *accept* scientific theories, and that *informativeness* can justify acceptance though it does not justify belief. He characterizes acceptance as involving a certain kind of commitment, which for scientists entails the adoption of a certain research program (van Fraassen 1980).

Kitcher has famously talked about the division of cognitive labor in scientific communities as a key component of epistemic success. The fact that scientists are not individual agents working alone in laboratories developing theories solely on the basis of their own findings, but are rather integrated into communities where there are external factors like funding and prestige can lead them to divide resources and coordinate research efforts in an investigative domain (Kitcher 1992). It's interesting to note that van Fraassen's work on acceptance and Kitcher's work on the division of cognitive labor were both met with responses that targeted whether or not such observations had an impact on what we saw as *genuinely epistemic* inquiry - a large response to van Fraassen's work has consisted in how we should understand the relationship between the attitudes of acceptance and belief, and it's been highlighted that Kitcher's work on the division of cognitive labor actually focuses on the *practical* reasons for researching views that are less likely to be true.⁵³

⁵³ See for examples Maher 1990 and Horwich 1991.

Contemporary projects include Fleisher’s “Rational Endorsement”, in which he pushes for us to recognize a distinct attitude from both belief and acceptance which he calls “endorsement.”⁵⁴ This is the attitude he thinks is rational for researchers to take during the course of inquiry. He points out that the recent work on tradeoffs that has been the subject of this paper, and in addition recent successes in accuracy-based utility theory (such as vindications of probabilism, conditionalization, and other coherence norms by appeal to accuracy alone) have lent support to the view that the rationality of belief is determined solely by considerations of whether it itself is likely to be true (Fleisher 2017).⁵⁵ But in addition to reasons to believe a particular proposition is true – reasons he calls “intrinsic reasons” we should also make a place for reasons that are *about* promoting or attaining the truth, but aren’t *themselves* indicators of truth. Reasons for endorsing a theory that stem from the fact that doing so would promote healthy inquiry are “extrinsic epistemic reasons” (Fleisher 2017, 15, citing a distinction made in Steel, 2010). Unlike Kitcher’s reasons for the distribution of cognitive labor which cite funding and prestige, these reasons seem undoubtedly epistemic – they’re reasons that target how to pursue truth in inquiry, even if this pursuit is global and less direct.

Information gathering and inquiry are central to our human existence – equally central, we might think, as questions about whether or not to believe the truth about a particular proposition, if only because such deliberation usually happens below the level of conscious awareness of us, and thus is not even typically subject to instrumentalist considerations. Though sometimes we’re in situations in which we have no conflicting goals and we’re deliberating about what to believe, more often our pursuit of truth puts us in other kinds of scenarios with differently configured aims. Thinking about which methods we ought to employ to better achieve those aims is thus a worthy enterprise. When we see the epistemic

⁵⁴ See in addition Friedman’s “Inquiry and Belief” (forthcoming) for another recent example of work that considers the need for a wider epistemic focal point.

⁵⁵ See for examples Joyce 1998 and Pettigrew 2016.

domain as united by its orientation toward the fundamental value of significant truth or knowledge and as reflecting this broader array of perspectives, scenarios, and aims, we see it as giving rise not just to a rule-implying norm that provides an answer to one particular epistemic question, but in addition to norms that appear to confer necessary normative force.

References

- Ahlstrom-Vij, Kristoffer, and Jeffrey Dunn. 2014. "A Defence of Epistemic Consequentialism." *Philosophical Quarterly* 64 (257): 541-51.
- Alexander, Larry. 2000. "Deontology at the Threshold." *San Diego Law Review* 37: 893-912.
- Alexander, Larry, and Michael Moore. 2016. "Deontological Ethics." *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta. <https://plato.stanford.edu/archives/win2016/entries/ethics-deontological/>.
- Alston, William. 1988. "The Deontological Conception of Epistemic Justification." In *Epistemic Justification: Essays in the Theory of Knowledge*, 115-152. Ithaca, NY: Cornell University Press.
- Anscombe, G. E. M. 1958. "Modern Moral Philosophy." *Philosophy* 33 (124): 1-19.
- Berker, Selim. 2015. "Reply to Goldman: Cutting Up the One to Save the Five in Epistemology." *Episteme* 12 (2): 145-53.
- . 2013a. "Epistemic Teleology and the Separateness of Propositions." *Philosophical Review* 122 (3): 337-93.
- . 2013b. "The Rejection of Epistemic Consequentialism." *Philosophical Issues* 23 (1): 363-87.
- Broome, John. 2007. "Is Rationality Normative?" *Disputatio* 2 (23): 161-78.
- Côté-Bouchard, Charles. 2016. "Can the Aim of Belief Ground Epistemic Normativity?" *Philosophical Studies*, 173 (12): 3181-98.
- Chisholm, Roderick M. 1977. *Theory of Knowledge*. 2d ed. Englewood Cliffs, N.J.: Prentice-Hall.
- Feldman, Richard. 2000. "The Ethics of Belief." *Philosophy and Phenomenological Research* 60 (3): 667-95.
- . 1988. "Epistemic obligations." *Philosophical Perspectives* 2: 235-56.

- Fleisher, Will. 2017. "Rational Endorsement." *Philosophical Studies* (September): 1-27. <https://doi.org/10.1007/s11098-017-0976-4>
- Foley, Richard. 1987. *The Theory of Epistemic Rationality*. Cambridge, MA: Harvard University Press.
- Frankena, William K. 1968. "Obligation and Value in the Philosophy of G. E. Moore." In *The Philosophy of G. E. Moore*, edited by Paul Arthur Schilpp. 3rd ed., 93-110. LaSalle, IL: Open Court.
- Friedman, Jane. 2018. "Teleological Epistemology." *Philosophical Studies*: 1-19. <http://doi.org/10.1007/s11098-018-1033-7>
- forthcoming. "Inquiry and Belief." *Noûs*. <https://doi.org/10.1111/nous.12222>
- Fumerton, Richard A. 1995. *Metaepistemology and Skepticism*. London: Rowman & Littlefield.
- Goldman, Alvin. 2015. "Reliabilism, Veritism, and Epistemic Consequentialism." *Episteme* 12 (2): 131-43.
- Greaves, Hilary. 2013. "Epistemic Decision Theory." *Mind* 122 (488): 915-52.
- Heuer, Ulrike. 2011. "The Paradox of Deontology, Revisited." In *Oxford Studies in Normative Ethics*, Volume 1, edited by Mark Timmons, 236-67. Oxford: Oxford University Press.
- Horwich, Paul. 1991. "On the Nature and Norms of Theoretical Commitment." *Philosophy of Science* 58 (1): 1-14.
- Jenkins, C. S. 2007. "Entitlement and Rationality." *Synthese* 157 (1): 25-45.
- Joyce, James M. 1998. "A Nonpragmatic Vindication of Probabilism." *Philosophy of Science* 65 (4): 575-603.
- Kagan, Shelly. 1998. *Normative Ethics*. Boulder, CO: Westview Press.
- Kamm, F. M. 2007. *Intricate Ethics: Rights, Responsibilities, and Permissible Harm*. Oxford: Oxford University Press.
- Kant, Immanuel. (1780) 1965. *The Metaphysical Elements of Justice: Part 1 of The Metaphysics of Morals*. Translated by John Ladd. Indianapolis: Hackett Publishing.

- Kim, Jaegwon. 1988. "What is Naturalized Epistemology?" *Philosophical Perspectives* 2: 381-405.
- Kitcher, Philip. 1990. "The Division of Cognitive Labor." *Journal of Philosophy* 87 (1): 5-22.
- Klausen, Søren Harnow. 2009. "Two Notions of Epistemic Normativity." *Theoria* 75 (3): 161-78.
- Korsgaard, Christine M. 1993. "The Reasons We Can Share: An Attack on the Distinction between Agent-relative and Agent-Neutral Values." *Social Philosophy and Policy* 10 (1): 24-51.
- Lowe, E. J., and Stephen P. Stich. 1990. "The Fragmentation of Reason: Preface to a Pragmatic Theory of Cognitive Evaluation." *Philosophical Quarterly* 42 (166): 98.
- Maguire, Barry. 2016. "The Value-Based Theory of Reasons." *Ergo*, 3 (9): 233-262.
- Maher, Patrick. 1990. "Acceptance without Belief." *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association* 1990 (1): 381-92.
- Moore, G. E. (1903) 1993. *Principia Ethica*. Edited by Thomas Baldwin. Cambridge: Cambridge University Press.
- Moore, Michael S. 2010. *Placing Blame: A Theory of the Criminal Law*. Oxford: Oxford University Press.
- Nagel, Thomas. 2012. *Mortal Questions*. Cambridge: Cambridge University Press.
- Nozick, Robert. 1984. "Philosophical Explanations." *Ethics* 94 (2): 326-7.
- Parfit, Derek. 2011. *On What Matters*. Oxford: Oxford University Press.
- Pettigrew, Richard. 2016. *Accuracy and the Laws of Credence*. Oxford: Oxford University Press.
- Pettit, Philip. 1997. "The Consequentialist Perspective." In *Three Methods of Ethics*, edited by M. Baron, P. Pettit and M. Slote, 92-150. Blackwell Publishers.
- Pollock, John. 1987. "Epistemic Norms." *Synthese* 71 (1): 61-95.
- Rawls, John. 1999. *A Theory of Justice*. Revised edition. Cambridge, MA: Belknap Press of Harvard University Press.

- Reisner, Andrew. 2009. "Abandoning the Buck Passing Analysis of Final Value." *Ethical Theory and Moral Practice* 12 (4): 379-95.
- Rinard, Susanna. 2017. "No Exception for Belief." *Philosophy and Phenomenological Research* 94 (1): 121-43.
- Ross, W. D. 1930. *The Right and the Good*. Oxford: Oxford University Press.
- Scanlon, Thomas. 1998. *What We Owe to Each Other*. Cambridge, MA: Belknap Press of Harvard University Press.
- Scheffler, Samuel. 1982. *The Rejection of Consequentialism: A Philosophical Investigation of the Considerations Underlying Rival Moral Conceptions*. Oxford: Oxford University Press.
- Schroeder, Mark. 2014. *Explaining the Reasons We Share, Explanation and Expression in Ethics, Volume 1*. Oxford University Press.
- Steel, Daniel. 2010. "Epistemic Values and the Argument from Inductive Risk." *Philosophy of Science* 77(1): 14-34.
- Steglich-Petersen, Asbjørn. 2011. "How to be a Teleologist about Epistemic Reasons." In *Reasons for Belief*, edited by A. Steglich-Petersen and A. Reisner, 13-33. Cambridge University Press.
- Stich, Stephen P. 1990. *The Fragmentation of Reason: Preface to a Pragmatic Theory of Cognitive Evaluation*. Cambridge, MA: MIT Press.
- Sylvan, Kurt. 2018. "Veritism Unswamped." *Mind* 127 (506): 381-435.
- Talbot, Brian. 2014. "Truth Promoting Non-Evidential Reasons for Belief." *Philosophical Studies* 168 (3): 599-618.
- van Fraassen, Bas C. 1980. *The Scientific Image*. Oxford: Oxford University Press.
- Williams, Bernard. 1973. "A Critique of Utilitarianism." In *Utilitarianism: For and Against*, edited by J. J. C. Smart and Bernard Williams, 77-150. Cambridge: Cambridge University Press.
- Wrenn, Chase. 2016. "Tradeoffs, Self-Promotion, and Epistemic Teleology." In *Epistemic Reasons, Norms and Goals*, edited by P. Schmechtig and M. Grajner, 249-276. De Gruyter.

Zimmerman, Michael J. 2015. "Value and Normativity." In *The Oxford Handbook of Value Theory*, edited by Iwao Hirose and Jonas Olson, 13-28. New York: Oxford University Press.