

An Epistemic Non-Consequentialism

Abstract. Despite the recent backlash against epistemic consequentialism, an explicit systematic alternative has yet to emerge. This paper articulates and defends a novel alternative, *Epistemic Kantianism*, which rests on a requirement of *respect for the truth*. §1 tackles some preliminaries concerning the proper formulation of the epistemic consequentialism / non-consequentialism divide, explains where Epistemic Kantianism falls in the dialectical landscape, and shows how it can capture what seems attractive about epistemic consequentialism while yielding predictions that are harder for the latter to secure in a principled way. §2 presents Epistemic Kantianism. §3 argues that it is uniquely posed to satisfy the desiderata set out in §1 on an ideal theory of epistemic justification. §4 gives three further arguments, suggesting that it (i) best explains the normative significance of the subject's perspective in epistemology, (ii) follows from the kind of axiology needed to solve the swamping problem together with modest assumptions about the relation between the evaluative and the deontic, and (iii) illuminates certain asymmetries in epistemic value and obligation. §5 takes stock and reassesses the score in the debate.

1 Introduction

While there has been a recent backlash against epistemic consequentialism, an explicit systematic alternative has yet to emerge.¹ Indeed, the positive activity has all been on the consequentialist side, with epistemic utility theorists enjoying a productive and enviably funded heyday. Compared to these peers, epistemic non-consequentialists risk looking like mere reactionaries shaking their fists at a glossy establishment, and perhaps one bound to prevail because the battleground is not politics but academic philosophy.

In this paper, I aim to fill the gap in the literature by offering a detailed positive version of epistemic non-consequentialism, showing that it is at least as attractive as any consequentialist epistemology, and more attractive on several counts. The plan is as follows. In the remainder of this section, I cover some preliminaries, explaining how the epistemic consequentialism/non-consequentialism debate should be understood, why the shining successes of epistemic utility theory are red herrings, and why it is unclear that epistemic consequentialism ever was the dominant position before epistemic utility theorists started treating

¹I say 'explicit' because while there are arguably some long-standing non-consequentialist views around (e.g., responsibilist virtue epistemology and perhaps some versions of coherentism and evidentialism), their proponents have not spoken up in the new literature explicitly devoted to the issue.

it as the default approach. I then develop a positive non-consequentialist view I call *Epistemic Kantianism* in §2. In §3, I argue that Epistemic Kantianism is uniquely poised to secure a number of desiderata on an ideal theory of epistemic justification. In §4, I give three further arguments for the view. After taking stock in §5, I close by reassessing the score in the epistemic consequentialism/non-consequentialism debate.

1.1 Understanding the Divide

While this paper offers a positive theory, it is a special case of a broader negative position—viz., epistemic non-consequentialism. To understand what that broader position maintains, we must first understand the view it opposes.

Following Berker (2013a), I take it that epistemic consequentialism answers the question ‘What should I believe?’ in a way that parallels the answer given by ethical consequentialists to the question ‘What should I do?’ Because consequentialism in ethics came first, a responsible discussion of epistemic consequentialism should draw on relevant themes and lessons from the more extensive literature in ethics.

Consequentialism in ethics is a first-order theory of right action, and consequentialism in epistemology is best understood as a theory of justified belief of the same sort.² Following Kagan (1992, 1997), I divide first-order theories into two classes:

Factoral theories of a normative status S seek to determine what factors bear on whether something instantiates S—i.e., *what are the S-relevant factors*.

Foundational first-order theories go deeper, seeking to provide fundamental normative explanations of *why these factors are the S-relevant factors*.

Despite being foundational in one sense, the latter are first-order rather than meta-normative, at least on one natural way of drawing that distinction. These theories do not seek *analyses*

²While I will most often discuss *justified belief* because of its central place in epistemology, epistemic consequentialism and non-consequentialism really are broader first-order theories about not just justification but other deontic categories like *requirement*, *obligation*, and *permissibility*.

At some points in what follows, it will be more natural to use these other deontic notions. I will assume for simplicity that they are related as follows. Permissibility is the dual of requirement (a strong necessity modal (‘must’)), and justification is the dual of obligation (a weak necessity modal (‘ought’)); for a defense of this way of classifying justification, see Beddor (2017). Like many epistemologists, I will take the theoretically central notions of justification, obligation, requirement, and permission to be perspective-dependent notions, though crucially—as we will see in §4—living up to perspective-dependent demands seems to have perspective-transcendent significance that needs explaining.

of normative properties, but rather the deepest normative explanations of why the factors relevant to the instantiation of these properties are relevant.³ Shallower explanations are possible and compatible. These are given by factoral theories.

Following Kagan, I take consequentialism in its most interesting form to be a foundational first-order theory. This classification matters. For there might be theories that look consequentialist at the factoral level but are non-consequentialist at the foundational level. To illustrate this point vividly, suppose for argument's sake that Parfit (2011) is right that rule consequentialism, Kantianism, and contractualism make the same predictions in their best formulations. Even if this claim were true, ethics wouldn't be over. For even if these theories make the same predictions, there remains a question about which provides the *best explanation* of why those predictions are correct. Perhaps the best explanation of why we should perform the actions rule consequentialists recommend is Kantian. If so, that theory would at bottom be Kantian.⁴

Of course, the underdetermination of foundational theory by factoral evidence might not be as extreme as Parfit (2011) suggests. But it would be surprising if there weren't some plausible extensionally equivalent pairs of theories at the foundational level, one consequentialist and another not. Some believe there is a procedure for generating such pairs: take

³Although I use the phrase 'normative explanation' for convenience, one needn't assume a pluralist view about metaphysical explanation or grounding of the sort attacked by Berker (forthcoming) to draw the needed distinction. There is surely a distinction between basic right-making features and *constituents* of the *property* of rightness. Rightness might be a simple property, *à la* some non-naturalisms, despite there being basic natural right-making characteristics. Perhaps the constitution relation isn't a grounding relation. If so, we can distinguish one kind of meta-normative theory and a foundational first-order theory without different grounding relations, but just a distinction between *constituents* and *explainers*. I don't see why Berker's arguments conflict with this distinction, which isn't a distinction in kinds of grounding. But if they do, then I happily commit to answering them elsewhere.

⁴I don't deny that it could be useful to use 'consequentialism' less stringently, so that views like (e.g.) Hare (1981)'s qualify as consequentialist in virtue of their factoral predictions, despite their Kantian foundations. But I think the most important disagreements aren't factoral. This is especially clear in epistemology, and this is why I use the term more stringently throughout.

In ethics, there is often disagreement between consequentialists and non-consequentialists on the significance of commonsense factors like the doing/allowing distinction, the intending/foreseeing distinction, negative and positive rights, and so on. Consequentialists are often willing to be revisionary about factors. By contrast, the epistemologists that seem to be the best candidates for being dubbed 'consequentialists' often argue from factoral intuitions (consider Goldman (1979, 1986)), and insist that it is their opponents who violate commonsense intuition. The opponents rarely respond by agreeing that their views require rejecting much of common sense (though see Fumerton (1995) for a notable exception). Given that most epistemologists in both camps put a premium on consistency with commonsense intuition, the key disagreements seem to be about foundations, not factors. For example, internalists insist that there is no *principled reliabilist explanation* of intuitions about demon victims (see, e.g., Foley (2004)), while reliabilists insist that internalist factors like sensory experience are *epiphenomenal* (see, e.g., Lyons (2009)).

whatever you find a plausible non-consequentialist theory and apply Dreier (1993)'s 'consequentializing' recipe to get an extensionally equivalent consequentialism.

There are many consequentialisms. What they have in common is a certain kind of explanation of rightness in terms of value. The qualifier 'a certain kind' matters. Many virtue ethicists think value is prior to rightness, but they are not consequentialists. Some Kant scholars and Kantians think he took the value of personhood, autonomy, or the good will to be normatively fundamental, but their interpretations needn't make him a consequentialist.⁵ Accordingly, the Rawlsian gloss on consequentialism as the view that value is prior to rightness is too simple. Consequentialists think value explains rightness in an *instrumental* or *means-end* way, where 'means' is understood broadly to include constitutive means, not just causal means. More officially, the key schematic thought is:

C: The fundamental normative explanation of why right actions are right is that they stand in some to-be-specified link L to certain means to final value.

Direct consequentialists take L to be the identity relation: right actions are themselves the means. *Indirect* consequentialists, such as rule consequentialists, take L to be some other relation (e.g., being permitted by an optimific rule).

Paralleling (C), epistemic consequentialism in its most general form makes the following schematic claim about justified belief:

EC: The fundamental normative explanation of why justified beliefs are justified is that they stand in a to-be-specified link L to certain means to final epistemic value.

Like consequentialism in ethics, epistemic consequentialism then has direct and indirect versions. The direct version takes the following form:

EC_d: The fundamental normative explanation of why justified beliefs are justified is that they are means to final epistemic value.

The more plausible, indirect version takes the following form:

EC_i: The fundamental normative explanation of why justified beliefs are justified is that they are indirectly linked to producers of final epistemic value.

⁵See, e.g., Dean (2006), Guyer (2000), Kamm (1992), and Wood (1999). See also Herman (1993: Ch. 10) for the broader point that Kantian ethics can be seen as value-based and not brutally deontological.

EC_i's many implementations are distinguished by how the link and the producers are specified. The link might, for example, be the *directly-caused-by* relation, and the producers might be *reliable processes*, yielding an explicitly consequentialist process reliabilism.

Although EC_i is not implausible, it is less clear to me than it is to Berker (2013a,b) that it has ever been the standard view in traditional epistemology, or one in the foreground of many classic discussions. For importantly, EC_i isn't clearly implied by reliabilism. The earliest versions of reliabilism weren't normative theories at all, and *a fortiori* weren't foundational first-order normative theories. They were rather attempts to unpack the non-accidentality condition on knowledge. Armstrong (1973), the early Goldman (1967, 1975, 1976), and Dretske (1981) didn't take knowledge to be partly constituted by justification, but just saw it as a straightforwardly natural, conceivably non-normative relation. Even if it were normative, non-accidentality wouldn't be the property of consequentialism; it is also central to Kantian conceptions of moral worth.⁶

While reliabilism developed into a first-order normative view about justification in Goldman (1979), it was not always linked to a consequentialist project. Goldman (1979) just defended it via standard intuitional methodology: reliabilism is preferable because it provides the best explanation of our considered epistemic intuitions. He only gave reliabilism consequentialist foundations in his (1986), and elsewhere only relied on consequentialist themes in arguing against opponents in his (1980), a paper which anticipated the rule consequentialism of his (1986). One could instead defend reliabilism as a non-normative theory or a factoral first-order theory. In neither form would reliabilism be a version of epistemic consequentialism. The mere fact that it could *look* consequentialist from a certain angle is irrelevant. Any theory can look consequentialist from a certain angle, given Dreier (1993)'s point that every plausible normative theory has an extensionally equivalent—if often unhelpful—consequentialist reformulation.

This is not to say that Berker (2013a) is wrong that many epistemologists *would* sign up to epistemic consequentialism in the foundational sense at issue here, or *do* make consequentialist claims in the foundational sense at issue here. Indeed, I think the results of

⁶See, e.g., Herman (1981), Baron (1995), and Stratton-Lake (2000).

the epistemic consequentialism vs. non-consequentialism debate are extremely important for deciding among the standard internalist and externalist theories of justification. But the foundationally consequentialist claims that have been made have not—with the exceptions of Alston (2005), Goldman (1986), Joyce (1998, 2009)—been *foregrounded* in defending first-order intuitions and factors important to the relevant theorists. Indeed, the traditional epistemologists other than Alston and Goldman who most clearly make foundationally consequentialist claims—namely, BonJour (1985: 7-8) and Foley (1993: 19)—would struggle to give a satisfactory rationale for the factor central to their theories, which is *coherence*: it is very unclear how coherence is truth-conducive.⁷

A broader point is worth stressing. Debates between reliabilists and their opponents have tended to center on how best to explain intuitions about cases, like forgotten evidence cases (see Goldman (1999)), clairvoyance cases (see BonJour (1980)), and the new evil demon thought experiment (see Cohen (1984)). I agree that there are serious case-based arguments for a *factoral* reliabilism. But they establish no foundational normative theses. At most they establish that there is some important epistemic status which implies or is necessarily equivalent to having reliable ancestry. This claim is compatible with a more fundamental non-consequentialist view. Just as Parfit gave a Kantian argument for the predictions of rule consequentialism, so one might offer a deontological argument for a factoral reliabilism.

If we understand matters in this way, we should conclude that too much literature on epistemic consequentialism has focused on extensional disagreements. Berker (2013a,b)'s main argument is that epistemic consequentialism has implausible predictions about cases, and most literature responds to this argument. Extensional disputes are not the most important disputes. I agree with Ahlstrom-Vij and Dunn (2014) that extensional arguments against epistemic consequentialism fail: some versions of indirect epistemic consequentialism don't make bad predictions about cases. This fact won't surprise anyone familiar with ethics. Consequentialism's flexibility is well-known. One can gerrymander it to avoid extensional difficulties. The more important question is how the explanations offered by these

⁷This is not to say that consequentialist thinking doesn't play an important role in some of their work, but just that it doesn't really seem to provide a foundation for their central factor. For example, as Berker (2009a) notes, BonJour did use consequentialist ideas in arguing against foundationalism, but eventually abandoned coherentism partly on account of the non-truth-conducivity of coherence.

reformulations stack up against the non-consequentialist's.

Besides the literature surrounding Berker (2013a,b), the other work being done under the banner of epistemic consequentialism is by epistemic utility theorists. While technically impressive, this literature is also hard to read as addressing fundamental normative issues. 'Epistemic utility theory' names a way of formally modelling epistemic normativity. Pettigrew (2015) doesn't even use it to model epistemic rightness. He rather uses it to model epistemic value, making it not even potentially analogous to a central version of consequentialism. Non-consequentialism may even be compatible with EUT's providing a useful formal model of the conditions for epistemic justification. But that model doesn't, non-consequentialists claim, provide a fundamental normative explanation.

EUT-ists do claim to give a non-pragmatic 'vindication' of probabilism. But it is one thing to vindicate a practice and another to give a fundamental normative explanation of its rightness. A vindication provides a clear justifying reason for a practice. One can give such a reason without its being the only or most fundamental reason. If, for example, one thinks accuracy is only worth caring about because it is practically valuable, one could accept EUT's vindication but think the fundamental story remains pragmatic.

1.2 Varieties of Epistemic Non-Consequentialism

Let's turn to the non-consequentialist side of the aisle. Epistemic non-consequentialism minimally claims that something other than a link to the promotion of final epistemic value fundamentally explains why justified beliefs are justified. There is considerable room for disagreement about what this other fundamental factor or set of factors might be. We can first distinguish *brute deontological views* and non-brute views. Brute views hold that the fundamental explanation of why certain responses are justified is that they conform to certain normatively bedrock rules.

This kind of non-consequentialism has become rare in ethics. Consequentialists would attack a strawman in attacking this view. But while it is not my preferred tack, brute deontology might be a more serious position in epistemology. Consider the debate about whether 'the norm of belief' is truth, knowledge, or something else. This debate could be understood

as concerning the sole absolute epistemic injunction. The various proposals that have been offered could then be construed as versions of brute epistemic deontology. One might, for example, pair Sutton (2007)'s knowledge-first angle with a brute deontology, by accepting the following claims:⁸

BD: The fundamental explanation of why justified beliefs are justified is that they are in conformity with the norm of belief.

KN: The norm of belief is knowledge.

Non-consequentialism needn't, however, be brute. One central non-consequentialist position in ethics—Kantianism—is hardly brute. Indeed, on a familiar reading, Kant sought to derive ethics from structural rationality.⁹

A different reading inspires my non-brute non-consequentialism. As I noted earlier, some readings of Kant are consistent with his being a value-firster—e.g., ones on which he takes the value of personhood, the good will, or autonomy to be the ground of moral principles. On such interpretations, Kant differs from consequentialists not in denying that right action is a matter of responding to value, but rather in holding that the response fundamentally demanded by the central value is *respect*.

The disagreement hence needn't be about *whether* values are the touchstones of justification, but about *how* they are. On one Kantian view, the ultimate value by reference to which moral obligations are explained is the value of personhood. This value doesn't call fundamentally for production—as if having babies were a near-fundamental obligation!—but rather respect. Here obligations are explained via a relation to value, but the relation of derivation is non-instrumental.

This way of explaining obligations by values is also familiar in political philosophy. Consider the relational egalitarianism of Anderson (1999). Anderson agrees with telic egalitarians like Temkin (1997) that equality is intrinsically valuable and a touchstone of justification. But she disagrees by holding that equality's value is fundamentally embodied not in the fact that *equal distributions of goods are to be promoted* (though that may be an upshot), but rather in the fact that *we are to relate to each other as equals*. Behavior can hence be justified

⁸For a more sophisticated theory of this sort, see Gibbons (2013).

⁹See Korsgaard (2008) and Markovits (2014).

by appeal to the value of equality, but not *fundamentally* on the grounds that it promotes equal welfare distributions: rather, in the first instance, by the fact that it manifests a way of regarding others as equals. This justification bottoms out in an appeal to value, but it is non-consequentialist.

Mark well that it is not essential for someone who embraces this order of explanation to hold that promoting value doesn't matter at all. Even on a strong version of this order of explanation, respect or something else other than promotion will just be what is *fundamentally* demanded by the value, with room for other responses of *derived* significance that are constrained by demands of respect, so that forms of promotion incompatible with respect are forbidden.

Political philosophy again affords a helpful illustration. Relational egalitarians needn't deny that promoting equal distributions of goods matters. But the strongest kind of relational egalitarian will claim that the importance of promoting equal distributions is derivative, and that a distribution should be considered valuable only if it *manifests* or *embodies* relational equality. Still, one could easily imagine that redistribution might be required for full relational equality. Indeed, a Marxist might insist that nothing short of the elimination of class hierarchy and non-trivial economic inequality is compatible with genuine relational equality. But what equality demands in the first instance is a certain way of relating to each other: one might then derive an obligation to promote distributive equality *assuming* a suitably progressive conception of relating to each other as equals.

One could call *purely* non-consequentialist those theories that explain all obligations solely by appeal to values that fundamentally call for respect, and only at best derivatively for promotion in certain cases. One could also imagine mixed theories according to which values may fundamentally call for both respect and promotion, but where the demand for respect trumps the demand for promotion. Such a view wouldn't derive duties to promote from duties of respect or the like. Instead, promotion would be demanded as such, provided respect is in place. Non-consequentialist views that accept side-constraints and explain them by appeal to a value that is not fundamentally to be promoted could also be examples: such theories would recognize two sorts of value—values that have dignity, and values that

don't—and explain the priority of respect by the lexical ranking of the first kind of value over the second; perhaps this is Nozick (1974)'s view. One could call such theories *impurely* non-consequentialist.

Kant could be understood in either way. I will count versions of either kind of theory as potentially Kantian for this paper's purposes. It is beyond the scope of this paper to defend a view that is Kantian in the strong sense of perfectly mirroring Kant's ethics—this would require more exegetical argument than is fitting for this occasion—but I do hope to establish a close analogy with some interpretation of Kant.

Some brief comments are in order to explain why I think a view that acknowledges derivative duties of promotion could be properly Kantian. Kant recognizes an imperfect or 'wide' duty of beneficence.¹⁰ One might take this duty to be grounded in the value of others' happiness, where that is understood as a value fundamentally to be promoted, provided that duties of respect are not violated. More ambitiously, one might try to argue that respect itself can require beneficence, perhaps by arguing that some other formulation of the Categorical Imperative can require beneficence. A further, more moderate possibility would be to take the significance of beneficence to be derived from a more fundamentally non-consequentialist demand that is distinct from, and operates in tandem with, requirements of respect. Perhaps the value of personhood calls fundamentally for both respect and the *practical love* Kant describes briefly in *Groundwork* (4:399) and more in *Metaphysics of Morals* (see, e.g., 6:449-450), where practical love is essentially *consistent* with respect (how else could it be love?) but not *derived* from it. One might then derive duties of beneficence from the fact that personhood calls for practical love alongside respect. I cannot settle here which of these views is closest to Kant's, but I take views structurally like any of these to merit the name 'Kantian'.¹¹

¹⁰See, e.g., *Groundwork* 4:423 and *Metaphysics of Morals* 6:450-2.

¹¹In the *Metaphysics of Morals* (6:450), Kant describes beneficence as the result of benevolence or practical love, and later (6:452) lists it as a duty of love, where this seems a category of duty *coordinate* with duties of respect (see Wood (1999: 324)). So I have some inclination to take Kant's late view to be the third, and to take humanity to call for both *respect* (which necessitates not treating others as mere means) and practical love (which necessitates at least benevolence, and a wide duty of beneficence). And I follow Wood (1997, 1999) in taking Kant's late view in the *Metaphysics of Morals* seriously, despite the focus of so much literature on the *Groundwork* and *Critique of Practical Reason*.

Herman (1993: Ch. 3) seems to adopt the second position in analyzing Kant's discussion of beneficence when giving examples of the universalizability test in *Groundwork* 4:423. In later work Herman (2007: Ch. 11) also

The view I will develop is a purely non-consequentialist one that regards fundamental epistemic value as calling for respect in the first instance, though I take respect in turn to require minimizing a certain epistemic disvalue in a certain way, and this could be understood as a negative promotional duty. It hence most closely resembles the second of the foregoing three Kantian views, where the fundamental value is taken to fundamentally demand just respect, but where respect is understood robustly, so that promotional duties might be derived from it. My view has the following three elements:

Value-First Epistemic Kantianism

1. **Valuing Thesis:** The fundamental normative explanation of why justified beliefs are justified is that they manifest certain ways of valuing fundamental epistemic value.
2. **Kantian Conception of Valuing:** The fundamental way of valuing epistemic value is respect.
3. **Veritist Conception of Fundamental Epistemic Value:** Accuracy is the fundamental epistemic value.

Epistemic consequentialists could be understood as agreeing with (1) and (3) but holding that the fundamental way of valuing epistemic value is promotion. If they were to agree to (1) and (3), my main argument against them (see §4.2) would be an argument against their instrumentalist alternative to (2):

- 2*. **Instrumentalist Conception of Valuing:** The fundamental way of valuing fundamental epistemic value is to promote it.

Conceivably, however, epistemic consequentialists could resist (1), holding that justification is explained directly by some connection to the promotion of value, not by the two alleged facts that (i) promotion is the fundamental way of valuing epistemic value and (ii) epistemic justification achieves such promotion. §4 hence gives other arguments.

describes a way of deriving duties of beneficence *after* the test imposed by the Categorical Imperative is met, and this looks more like the first position. In a discussion of the second formulation of the Categorical Imperative, Korsgaard (1996: Ch. 4) suggests that the part of this requirement that enjoins us to treat humanity as end-in-itself may have beneficence as an upshot. Whether Korsgaard's treatment is best understood as a relative of the first position that takes humanity to be *preserved* as well as respected (where preservation is teleological in the same way as promotion) or as more similar to the third position is unclear (though she doesn't mention practical love in this chapter).

1.3 Desiderata

To assess the view, it will help to establish some desiderata. Obviously, any acceptable view should have certain stock virtues. It should comport *ceteris paribus* with commonsense intuition. Where it doesn't, it should be paired with an error theory explaining why *cetera* aren't *paria*. It should also be parsimonious in explaining the data, and its explanations should be at least as compelling as its competitors'. Finally, it should be compatible with wider attractive outlooks like naturalism.

There are two more specific desiderata that I want epistemic non-consequentialism to satisfy. The first is the ability to explain why many epistemologists have been wrongly attracted to epistemic consequentialism. As part of this desideratum, I would like epistemic non-consequentialism to be consistent with the importance of reliability and of minimizing expected inaccuracy.

Another desideratum I want epistemic non-consequentialism to satisfy is that it not disconnect epistemic duty from epistemic value. Obviously, I do not want epistemic duties to hinge wholly on instrumentality to final epistemic value. But avoiding that kind of connection doesn't imply avoiding any strong deontic-evaluative connection. Though I don't consider it non-negotiable, I want a view that honors the following thought:

Deontic-Evaluative Alignment (DEA): One's doxastic attitude is epistemically justified iff it is (one of) the epistemically best open to one.

DEA might be confused with direct epistemic consequentialism. But it is weaker and compatible with several non-consequentialist approaches.

Note that the construction 'best doxastic attitude' can be understood as an attributive evaluative construction,¹² like 'best chess move'. Such evaluations are in an important way *insulated*. Most obviously, they are insulated from normativity not constitutively associated with the kind of thing the adjective modifies. They are also insulated from some evaluative facts in the same domain. For illustration, suppose some available move would uniquely checkmate your opponent, but that his resulting loss will make him so determined that he'll win all later games. It remains the best chess move open to one. As Berker (2013a)'s 'sepa-

¹²See Geach (1956) and Thomson (2008) for discussion.

rateness of propositions' counterexamples suggest, epistemic evaluation is similarly insulated.

While it would involve breaking the analogy with ethics, one could imagine epistemic consequentialists agreeing that epistemic evaluation is thus insulated and adopting a proposition-relative view on which

(*) instrumentality to the (attributively) epistemically best attitude open to one vis-à-vis p explains one's epistemic justification to take a certain doxastic attitude vis-à-vis p.

(*) is, however, independent of DEA. One could run the explanation the other way, insisting that certain attitudes are the epistemically best because they are the attitudes one is epistemically justified in forming.

While this is a respectable option, my preferred vindication of DEA is different. I hold that

whether one is epistemically justified in forming some doxastic attitude is determined by whether it would manifest the response(s) demanded by fundamental epistemic value.

Like Wood (1999)'s Kant and Raz (1999, 2012), I think value generates all demands, where these aren't exhausted by promotion. To vary an example from Scanlon (1998), friendship's value demands that I not betray my friend Mike even to cause several new friendships to form. While my being a good friend to Mike is finally valuable, it is so *because* it respects our friendship's value. There can, after all, be final value that is derivative, as is widely recognized in ethics.¹³ Hence, it is backwards to claim that I should thus relate to Mike because it brings about a finally good state of affairs. While its value is non-instrumental, that state of affairs—and states of affairs generally, as Anderson (1993) argues—derives value from the fact that it manifests respect for our friendship's value.

We should similarly explain the alignment between the epistemically justified and epistemically best attitudes: it is because certain attitudes respect accuracy's value that they are best, and also because they respect accuracy's value that they are the attitudes one is epistemically justified in forming. Value is at the bottom—just non-instrumentally.

¹³See Anderson (1993), Herman (1993), Hurka (2001), Kagan (1998), Korsgaard (1983), Moore (1903), Parfit (2011), Rabinowicz and Ronnow-Rasmussen (2000), Rashdall (1907), Scanlon (1998), and Zimmerman (2001).

An epistemic non-consequentialism that satisfies these desiderata couldn't be seen as positing brute duties that there is no clear reason to obey. By being consistent with motivations for epistemic consequentialism, it would emerge as a serious competitor. Given its other advantages, it could stand out as the superior option, not some reactionary view.

2 Veritist Epistemic Kantianism

I turn now to develop such an epistemic non-consequentialism. I'll begin, however, by stating some background assumptions that are not specifically non-consequentialist.

Like many theorists, I take the domain of epistemic value to be structured: some items are *fundamentally* epistemically valuable, the rest *derivatively* so. To say that certain things are derivatively epistemically valuable is not to say that they are valuable merely as means. As I noted earlier, many axiologists reject the claim that all derivative value is instrumental value, including some consequentialists (e.g., Hurka (2001)).

Like many epistemic consequentialists, I take the fundamental axiological level to be sparse—a single thing is fundamentally epistemically valuable: *accuracy*. Unlike some epistemic utility theorists, however, I don't think that embracing an accuracy-first epistemic axiology makes one an epistemic consequentialist.¹⁴ The more general principle needed to secure this thought is clearly false. Being a monist of a certain kind about fundamental *value* in a domain implies nothing about whether one is a consequentialist, since consequentialism is standardly a view about *rightness*.

Let's turn to some more distinctive and less consequentialism-friendly ideas. As I mentioned before, I think that values demand or at least justify certain responses and explain why these responses are demanded or justified. There is often a variety of responses demanded or justified by a value. Personhood demands respect, but it also calls for preservation, and some might even think the nature of personhood justifies or at least permits promotion. Friendship demands loyalty and a kind of love, justifies guarding and protecting one's friends, and permits celebrating them and promoting their welfare. Fidelity demands honoring one's promises, and also permits and perhaps justifies promoting promise-keeping by encouraging

¹⁴Consider Pettigrew (2015) and Joyce (2009)'s way of using 'accuracy-first epistemology'.

people to keep their promises.

Among responses that are demanded or justified we should draw a certain distinction. Some responses are *fundamentally* demanded or justified by a value. Consider personhood. To vary Narveson (1973)'s point, respecting existing people is one thing and causing states of affairs to exist in which people (potentially new ones) are respected is another. While both could be justified, only the former is *fundamentally* demanded by personhood. To the extent that creating new respected people can be justified, it seems *derivatively* so: at best, it is *because* people call for respect that promoting states of affairs in which people are respected can be good.

2.1 Veritist Kantianism about Epistemic Value

The first pillar of my Veritist Epistemic Kantianism assumes these ideas:

- (i) **Kantian Conception of Accuracy's Epistemic Value:** The value of accuracy is fundamentally to be respected.¹⁵

Hence, while promoting true belief may be epistemically permissible or justified, it would be so only derivatively.

As I will further explain in §4, this first pillar has some of the same appeal as Narveson (1973)'s point in population ethics. Just it is attractive to replace:

Strong P-Obligation: We ought to cause people to exist who are treated well.

with

Weak P-Obligation: We ought to treat existing people well.

so it is attractive to replace:

Strong T-Obligation: We ought to produce accurate doxastic attitudes.

with

Weak T-Obligation: We ought to ensure the accuracy of our existing doxastic attitudes (where 'ensure' doesn't mean 'bring about', but rather something like 'certify').

¹⁵While I happen to prefer Veritism, one could imagine a Knowledge-First Epistemic Kantianism that makes a parallel claim about the value of knowledge.

We don't have an unqualified obligation to produce true beliefs, as stock concerns about trivial truths suggest. And it is irrelevant to epistemic evaluation proper to note that we can restrict the Strong T-Obligation by sticking 'about subject matters of interest or importance' after 'doxastic attitudes'. Such qualifications may matter if we are trying to understand obligations *simpliciter* with an epistemic subject matter. But if we focus on epistemic evaluation in the attributive sense, such qualifications have no place.

This fact is stark with respect to negative epistemic evaluation, as Grimm (2009) noted. Many of our basic beliefs—e.g., most perceptual beliefs—are scarcely more interesting than beliefs about blades of grass. If these turn out radically false, they are not less epistemically bad just by being boring. This is one illustration of why such restrictions are irrelevant to epistemic evaluation. But the more important point is that once we distinguish epistemic values proper from values that are epistemic, we should see that the intuitions behind a pragmatically qualified Strong T-Obligation are unreliable.

Granting the necessary distinctions, then, it seems plausible that there is an epistemic asymmetry between producing accurate beliefs and ensuring the accuracy of our existing doxastic inventories, akin to the practical asymmetry between producing states of affairs in which people are respected and respecting existing people. This asymmetry is what one would expect if (i) were true. So there is some intuitive support for (i).

2.2 Respect: General Observations

What is it to respect accuracy? Let's begin with some general points about respect to zero in on the relevant notion.

The language of respect is used in conflicting ways, and only some express what I have in mind. By excluding some irrelevant uses, we can get clearer on the relevant notion. Darwall (1977) helpfully distinguished two uses. On the one hand, there is *appraisal respect*, which involves *esteem*. When I say 'I respect Barack Obama', I mean I hold him in esteem (though I respect him in other senses too). On the other hand, there is a different kind of respect that involves appropriate *deliberative constraint* rather than esteem, which Darwall dubs 'recognition respect':

There is a kind of respect which can have any of a number of different sorts of things as its object and which consists, most generally, in a disposition to weigh appropriately in one's deliberations some feature of the thing in question and to act accordingly. The law, someone's feelings, and social institutions with their positions and roles are all examples of things which can be the object of this sort of respect. Since this kind of respect consists in giving appropriate consideration or recognition to some feature of its object in deliberating about what to do, I shall call it *recognition respect*.¹⁶

This is close to the respect that figures in my view. I hold that we ought to have something like Darwallian recognition respect for accuracy. Having this kind of respect for accuracy is something that we achieve in virtue of giving accuracy appropriate weight in our doxastic deliberation and attitude formation, where that is done by forming beliefs on the basis of the apparent evidence. I do not assume that sufficient apparent evidence is never misleading. So I think respecting accuracy is consistent with forming beliefs that are, in fact, inaccurate.

The fact that respect can demand forming false beliefs might seem to create a tension with an accuracy-first axiology. But it doesn't. What (i) suggests is that while accuracy considered as an ideal is fundamentally epistemically valuable, particular *states* of accurate belief might not be worth producing. It is natural for epistemic Kantians to deny that states are the primary bearers of fundamental value; non-consequentialists in ethics propose similar axiologies.¹⁷ This claim can be rejected while the accuracy-first theme is preserved: we simply deny that it is the state of believing accurately that is fundamentally valuable. On this view, it is coherent to think that particular accurate belief-states might be epistemically bad overall. For what the combination of (i) and an accuracy-first axiology suggest is the following sort of value derivation:

Derivation through Respect: A cognitive state is *pro tanto* derivatively epistemically good if it manifests respect for accuracy, and *pro tanto* derivatively bad if it manifests disrespect for accuracy.

Accurate belief states that manifest disrespect for accuracy—e.g., those formed against the evidence—will count as *pro tanto* derivatively epistemically bad for accuracy-first reasons. Whether these states have any *pro tanto* intrinsic epistemic goodness is not clear. (i) doesn't claim that promotion is never a response justified by the value of accuracy. The analogue of

¹⁶Darwall (1977: 38).

¹⁷See especially Anderson (1993).

Narveson's asymmetry in population ethics does suggest that promotion isn't the response *fundamentally* justified by this value. Nonetheless, it could conceivably be *among* the justified responses. If so, we needn't insist that an unjustified accurate belief has no epistemic value whatsoever. We'll return to this issue.

2.3 Recognition Respect in General: Further Details

Let's now try to characterize recognition respect of the kind that matters for Epistemic Kantianism in more detail. Remember that according to the early Darwall, recognition respect is a relation one can bear to many different sorts of things, and one manifests it when one 'weigh[s] appropriately in one's deliberations some feature of the thing in question and... act[s] accordingly.'¹⁸ Darwall's gloss is not entirely helpful. When does one weigh the feature *appropriately*?

Here it is useful to distinguish more and less subjective glosses, and hence to distinguish weaker and stronger notions of respect. To bring out these glosses, I will begin with the idea that the way in which one appropriately 'weighs' a value in deliberation, action, or thought is by considering and responding to the *reasons* given by that value. Now, it is common to distinguish several kinds of normative reason. On the one hand, there are the normative reasons we have *given our beliefs* about the relevant facts. I will call these *mere subjective reasons*. There are then the normative reasons we have *given the apparent evidence* about the relevant facts.¹⁹ These are *apparent evidence-relative reasons*; to save words, I will just call them *evidence-relative reasons*, though I think that not all evidence is apparent, and also suspect that not all apparent evidence is real evidence. Finally, there are the *perspective-transcendent* reasons given by a value relative to all the facts, which may not be apparent.

With these distinctions in hand, we can begin by distinguishing the following two kinds of recognition respect for a value V:

¹⁸Darwall (1977: 38). In Darwall (2006), recognition respect is taken to be a second-personal notion. But Darwall (1977)'s original idea was broader. I set aside Darwall's more recent use of this idea.

¹⁹Apparent evidence can itself be understood more or less subjectively. On the one hand, to borrow some terminology from Schellenberg (2013), there is the *phenomenal evidence* one shares with an envatted counterpart with the same experiences and other internal states. On the other hand, there is the *factive evidence* that one only has in the good case. We can hence distinguish between phenomenally and factively apparent evidence. We might as a result distinguish between two kinds of evidence-relative reasons.

- A person *S* *weakly respects* *V* in ϕ -ing iff ϕ -ing is favored by *S*'s merely subjective *V*-given reasons, and *S* ϕ s for these subjective reasons.
- A person *S* *strongly respects* *V* in ϕ -ing iff ϕ -ing is favored by *S*'s evidence-relative *V*-given reasons, and *S* ϕ s for these evidence-relative reasons.

There is something good from the epistemic point of view about manifesting each form of respect, but the second is better than the first. A person exhibits a sort of *conscientiousness* in weakly respecting a value, but conscientiousness isn't much of an achievement. It may even get in the way, as I will discuss later.

Now, I take it that it is possible for a person to strongly respect a value *V* even if she fails to do what *all* the perspective-transcendent *V*-relevant reasons recommend. Consider the value of honesty. Suppose the person's evidence overwhelmingly counts in favor of *P*, but remarkably happens to be misleading. She could respect the value of honesty if she stated *P* as her honest opinion under these conditions. Still, *P* is false, and what is ideally honest is to tell the truth as a manifestation of respect for the importance of truth-telling.

One could introduce a further, yet stronger kind of respect to capture what is missing in this case. I don't want to do so, since it is not intuitive to me to say that *respect* is what is missing here. But one might at least introduce a further orientation to value here under a different name. Let's say that

- A person *adheres* to *V* iff ϕ -ing is favored by both the perspective-transcendent and the evidence-relative *V*-relevant reasons, and the person's ϕ on account of this fact.

Being ideally honest perhaps requires adherence to the value of honesty, not just respect. It is, however, not plausible that a person is dishonest in virtue of failing to adhere to honesty in this stipulative sense, nor is it clear to me that it would be a strike against a person's moral worth merely to fail to adhere, if strong respect were given.

The foregoing glosses are not intended as analyses or even as first-order accounts of weak and strong respect; instead, since talk of respect needs regimentation to be theoretically illuminating, I am content to take the indented claims on board as *explications* of ordinary respect concepts, as Carnap (1947) would say. When telling their first-order story, Epistemic Kantians will read such claims *backwards*, insisting that the value of conscientiousness and the normativity of subjective reasons is grounded in a more fundamental value via its demand

for weak respect, and that the normativity of evidence-relative reasons is grounded by its relation to a more fundamental value via its demand for strong respect. Nevertheless, the reasons to which these claims appeal are those independently investigated in the literature on reasons and rationality. Since we have an independent grip on them, we can appeal to them and their necessary link with respect to understand the conditions under which one manifests respect. The Kantian thought is then that if we want to know why heeding reasons *matters*, we should look to a necessary link with respect for more fundamental value.

Which notion of respect is the relevant one for my purposes? Both matter, but only the strong kind of respect matters much, as I'd already hinted. As writers in the literature on moral worth like Nomy Arpaly have suggested, it is not clear that conscientiousness is even required for outright moral worth, though it might sometimes add a little something. Huck Finn provides the familiar illustration. Although Huck thinks he acts wrongly in helping Jim, he is not only acting in conformity with the reasons provided by Jim's humanity, but is unwittingly complying with them. While he fails to heed his conscience, his action is—as Arpaly (2003) says—morally rational. Since moral rationality lines up with heeding evidence-relative moral reasons, we should also conclude that unbeknownst to Huck, he both has and is heeding sufficient evidence-relative moral reasons. If we got inside his psychology, it wouldn't be hard to see has such reasons: at some level, he can presumably see that Jim is a person with dignity, and this is presumably part of what moves him if he is praiseworthy.

On my broader story, the reason why apparent reasons in a domain are real reasons is that respect for some underlying value constitutively requires compliance with them. Epistemic Kantianism's explanation of the perspective-transcendent significance of the perspective-dependent sort of justification central in traditional epistemology is a special case of this story. While it may be possible to have the justified false belief that one is responding to *all* perspective-transcendent demands, there is *a* perspective-transcendent demand one heeds whenever one manifests respect. Since Epistemic Kantianism claims that respect is the response fundamentally demanded by accuracy, it follows that one is excelling in a perspective-transcendent sense in the epistemic domain if one believes justifiedly, even if one believes

falsely.

2.4 Recognition Respect for Accuracy: Further Details

Having explicated these general notions of respect, it is not difficult to grasp the special case of respect for accuracy in more detail. As we've seen, we can understand recognition respect for a value V as constituted by a certain kind of reasons-responsiveness. When the value in question is accuracy, the reasons that matter are accuracy-relevant reasons. I take these simply to be kinds of evidence. We can then distinguish between the merely subjective evidence constituted by (contents of) a person's beliefs and other doxastic attitudes, the fundamentally non-doxastic apparent evidence constituted by (contents of) her non-doxastic mental states (e.g., experiences), and the perspective-transcendent evidence that is simply out there in the world waiting to be discovered (e.g., evidence that might later be uncovered in examining the scene of the crime). We can then say that a thinker S

weakly respects the value of accuracy in holding a doxastic attitude $D(p)$ iff $D(p)$ is favored by S 's merely subjective evidence, and S complies with this subjective evidence in $D(p)$ -ing;

strongly respects the value of accuracy in holding $D(p)$ iff $D(p)$ is favored by S 's apparent evidence, and S complies with this apparent evidence in $D(p)$ -ing,

- where S counts as *complying* with some evidence E iff S forms the attitude favored by E as a manifestation of a disposition to form attitudes given E -like evidence only if they are in fact favored E -like evidence.

Again, the Epistemic Kantian does not read these biconditionals as *analyses* of the relevant notions of respect. On the contrary, complying with the different kinds of evidence derives significance via a link to these forms of respect for accuracy.

It is interesting to consider the notion of adherence when the value in question is accuracy. It turns out to correspond to an important epistemic value, at least when that value is understood as Kant did. A thinker S

adheres to the value of accuracy in holding a doxastic attitude $D(p)$ iff $D(p)$ is both favored by the perspective-transcendent evidence and by S 's apparent evidence, and S complies with this overlapping evidence in $D(p)$ -ing.

What might this involve in particular cases? Suppose that it is a fact that *p*, and hence that there is sufficient perspective-transcendent evidence for *p* consisting in this fact. Suppose furthermore that *S* sees that *p*. Here the doxastic attitude favored is belief, and the person's belief will adhere by being accurate as a manifestation of a disposition to respond to this sort of evidence. In this case, a person achieves something that lines up closely with what Kant himself deemed *knowledge*. As Schroeder (2015) brought to the attention of recent epistemology, Kant (1781/1999: A822/B850) held that knowledge is *belief for sufficient objective and subjective reason*. In the kind of case just considered, adherence consists in just such belief. Of course, one might doubt that all knowledge requires evidence-responsiveness, and I agree that what Sosa (2007) calls *animal knowledge* does not. The point is just that Epistemic Kantianism does seem poised to explain *reflective knowledge*, and it is only this kind of knowledge that is plausibly understood as normatively constituted. The animal knowledge that a person has in virtue of seeing that *p* or remembering that *p*, for example, seems best understood just as a general factive mental state, in itself no more normative than seeing that *p* or remembering that *p*.

Having made the connection between adherence and reflective knowledge, it is easy to see how other familiar epistemic notions might map onto our notions of respect. Doxastic attitudes that weakly respect accuracy will be in a certain way *coherent*, in virtue of fitting with the subject's overall doxastic take on how the evidence stands. Doxastic attitudes that strongly respect accuracy will be *justified* in virtue of complying with the evidence.

A complete epistemology would include a story about graded belief. Here I don't blush to confess that I am happy to agree with some epistemic utility theorists that one's credences should minimize expected inaccuracy relative to the evidence, where this is measured by some scoring rule (e.g., the Brier score). This is not an essentially consequentialist idea: we can take the notion of respect to be more fundamental, and to explain why we should care about minimizing expected inaccuracy (viz., doing so constitutes respect for accuracy in credence-formation).

The point has precedent in ethics. As Scanlon (1998) noted, Scheffler (1982) was wrong to think that the seeming inescapability of decision theory in formally modeling rationality

implies some version of consequentialism. A more fundamental, non-consequentialist rationale can be given for demands like the demand to maximize expected value. Expected value is, after all, by itself just *expected* value. If we want to explain why maximizing expected value is *genuinely* rather than *apparently* good, one must appeal to some different ideal, like respect.

2.5 Veritist Kantianism about Epistemic Justification

So far we have considered the Epistemic Kantian's account of epistemic value and the response fundamentally demanded by it. Consistently with this story, one might deny that the facts about epistemic justification are fully explained by facts about epistemic value and the responses demanded by it, and uphold a brute deontological view. But this approach is not mine. As I hinted before, I accept Deontic-Evaluative Alignment. If I took DEA on board as a background assumption, I could together with (i) derive a Kantian picture of justification as follows:

1. (Assumption) A doxastic attitude is (one of) the attributively best available to one iff it manifests the response fundamentally demanded by fundamental epistemic value.
2. (DEA) A doxastic attitude is epistemically justified iff it is (one of) the attributively best doxastic attitudes available to one.
3. So, a doxastic attitude is epistemically justified iff it manifests the response fundamentally demanded by fundamental epistemic value.
4. Respect for accuracy is the response fundamentally demanded by fundamental epistemic value.
5. So, a doxastic attitude is epistemically justified iff it manifests respect for accuracy.

Call (5) the *Kantian Equivalence for Epistemic Justification*, which I will abbreviate as $J=R$ because of the fact that it makes epistemic justification extensionally equivalent to respect for accuracy. While this reasoning provides one way of reaching $J=R$, I don't accept DEA as a background assumption. Epistemic Kantianism can *explain* why DEA is true provided that one takes on a more explicitly Kantian background assumption:

- (T) Whether one is epistemically justified in forming some doxastic attitude is determined by whether it would manifest the response(s) demanded by fundamental epistemic value.

One can derive $J=R$ from (T) and (i), and can derive DEA from (T), (i) and (Assumption). This latter fact suggests my reason for accepting DEA: it follows from Kantian premises that also support $J=R$.

$J=R$ brings us near the second central component of Epistemic Kantianism. Since $J=R$ is just a necessary biconditional, its truth may demand explanation. The second component of Epistemic Kantianism provides this explanation:

- (ii) **Epistemic Kantianism about Justified Belief (EKJ):** The fundamental normative explanation of why justified beliefs are justified is that they manifest strong respect for accuracy.

EKJ is not intended as a rival to various theories of epistemic justification. It is not a rival to any metaphysical *analysis* of epistemic justification. Attempts to naturalize epistemic justification belong to this category. It is not sufficient to provide a naturalistic analysis of justification that one specifies what natural properties *make for* justification in the sense of a first-order theory of justification-making characteristics. Compare how non-naturalist utilitarians like Sidgwick and Parfit agree that certain natural properties ‘make for’ right action, even though they think the normative property for which these natural properties provide a criterion is non-natural.

EKJ is also no rival to any first-order factoral theory of epistemic justification. Exactly which theories are factoral is unclear; the answer depends on how their proponents would intend them if they were made aware of the factoral/foundational distinction. But numerous views about epistemic justification, like phenomenal conservatism, virtue reliabilism, evidentialism, etc., *could* be read as factoral theories. One could agree with phenomenal conservatives and some evidentialists that a belief is justified iff it is probable conditional on the all-things-considered non-doxastic seemings,²⁰ but find oneself left with a deeper question about *why* such beliefs are justified. EKJ may give the answer.

While EKJ can give such an underpinning for factoral theories, it is strictly speaking neutral on factoral questions. The same holds for other foundational first-order theories. Just as consequentialism in ethics needs to be supplemented with a theory of final value to make firm predictions about cases, so epistemic consequentialism needs to be supplemented with

²⁰For such theories, see Huemer (2007) and Conee and Feldman (2004).

a theory of fundamental epistemic goodness (e.g., Veritism) to make firm predictions about cases. EKJ's factoral neutrality will play an important role in what follows. As we will see later, I think we can explain why a certain kind of epistemic consequentialism seemed plausible by noting that it is extensionally equivalent to a plausible factorally supplemented version of Epistemic Kantianism, where only the latter provides a satisfying normative explanation.

2.6 Why Believe $J=R$?

Since Epistemic Kantians offer a foundational first-order theory rather than a factoral theory, their most important task is to show why their explanation of the factoral data outperforms the epistemic consequentialist's. Still, it had better be clear in the first place that Epistemic Kantians respect the factoral data. But is it even true that all justified beliefs manifest respect for accuracy, as $J=R$ claims? Without some defense of $J=R$, it is hard to see why we should be interested in EKJ, and Epistemic Kantianism as a whole.

When $J=R$ is properly understood, it shouldn't be controversial. The business of manifesting recognition respect for accuracy needn't be understood in a psychologically taxing manner. I took it to be equivalent to a kind of evidence-responsiveness. If so, $J=R$ should not be more controversial than:

$J=C$: S's belief that p is epistemically justified iff this belief is favored by the subject's evidence and the subject complies with this evidence.

$J=C$ is not, it is worth admitting, terribly informative. $J=C$ leaves it open what constitutes compliance with the evidence. One can take different views here, depending on one's independent case-by-case intuitions. It is, for example, consistent with $J=C$ and $J=R$ to hold that one complies with evidence in believing that p iff one believes that p as a manifestation of one's competence to respond to *objective evidence* in the indicator reliabilist's sense. But this flexibility is a boon, not a burden, for $J=R$. Just take the case-by-case intuitions about justification that seem strongest and read 'compliance' and 'respect' so that $J=C$ seem $J=R$ true, as they should.

Now, as I said before, the defender of Epistemic Kantianism will suggest that whether a belief counts as complying with the evidence should itself be normatively explained by the

more fundamental demand of respect for accuracy. Still, it is legitimate to defend this view by noting that

J=C is plausible on a case-by-case basis when understood in a theory-neutral way,

and that

the best explanation of why J=C is extensionally plausible is that J=R and indeed Epistemic Kantianism are true.

And J=C is plausible when understood in a theory-neutral way. It is consistent with all the data that motivate evidentialism and undermine simple versions of reliabilism (e.g., the clairvoyance and new evil demon problems, to be further discussed in §4). But it is also consistent with certain counterexamples to simple versions of evidentialism like the improper basing objections from Turri (2010) and Goldman (2012: Introduction). For it is one thing *merely to conform* to the demands of the evidence by happening to believe what it recommends, and another to *comply*. Compliance plausibly requires manifesting dispositional sensitivity to the relevant evidence-for relation between one's evidence and the doxastic attitude that it supports, and this is missing in improper basing cases. Defenses of these claims can, at any rate, be found elsewhere.²¹

Hence, I take the case for J=R to be straightforward, though it must build upon independent factoral and meta-normative theorizing about justification and reasons. The case, in short, is just this. Given the earlier gloss on respect, J=R is equivalent to the equivalence suggested by sophisticated reasons-based theories of justification. These theories are independently defensible. So, J=R is defensible parasitically.

3 Why Epistemic Kantianism? The Desiderata Revisited

Let's now consider in more detail why we should take Epistemic Kantianism seriously. The first and most general point in its favor is that Epistemic Kantianism seems uniquely poised to capture the desiderata from §1.3 on a foundational theory of epistemic justification. In defending this claim, I will begin by showing first how Epistemic Kantianism captures the

²¹See, e.g., Lord (2018) and Mantel (2018).

stock desiderata—viz., extensional adequacy, explanatory power, simplicity, and compatibility with wider constraints like naturalism. I will then explain how Epistemic Kantianism can explain why many epistemologists have been wrongly attracted to epistemic consequentialism by *deriving* key consequentialist desiderata—reliability and the minimization of expected inaccuracy—and by giving a better explanation of why these desiderata matter than consequentialists can themselves give.

Epistemic Kantianism Better Captures Stock Desiderata than Competing Views. For starters, as we have partly seen, Epistemic Kantianism comports with intuitions about epistemic justification to at least the degree to which any sophisticated reasons-based account would. It does so because respect for accuracy coincides extensionally with compliance with evidence, which is what the sophisticated reasons-based theorist thinks about justification. But it may outperform the average reasons-based account in virtue of the fact that the requirement of respect transparently excludes bad etiologies. A belief that fits the evidence and is based on it but improperly so—say, via use of a bad inference rule—wouldn't intuitively count as manifesting respect for accuracy.

Again, given the way it was understood in the preceding section, respect for accuracy requires forming beliefs in ways that manifest sensitivity to evidence-for relations. For each form of respect coincided with a kind of *compliance* with the evidence, which is itself understood as akin to what Sosa (2007, 2015) calls 'aptness': it requires not just forming the attitude that is favored by the evidence, but doing so *as a manifestation of a disposition to do so*, where this disposition itself amounts to a kind of sensitivity to the evidence. Accordingly, Epistemic Kantianism can absorb some attractive predictions of virtue epistemology while avoiding an instrumentalist rationale for them. Forming beliefs as a manifestation of sensitivity to the evidence matters not *fundamentally* because it is reliable in certain worlds, but rather because it constitutes respect for accuracy.

Besides capturing key intuitions behind internalist and externalist views, Epistemic Kantianism provides a parsimonious explanation of these intuitions. On my view, there is one fundamental epistemic norm—respect accuracy!—and one fundamental epistemic value—accuracy.

Epistemic Kantianism gives a unifying reason for caring about heeding the evidence and for caring about forming beliefs in an epistemically virtuous way: these things matter in virtue of manifesting respect for accuracy, which in turn matters because it is the response fundamentally demanded by the value of accuracy.

This fact deserves emphasis. Consider the internalist side first. As it stands, defenders of evidentialist and other reasons-based views in epistemology have not offered a plausible story about why respecting evidence matters. It is hard to believe that it *fundamentally* matters: we seem to care about respecting evidence only derivatively, as a manifestation of a concern for something more fundamental. Evidentialists and other reasons-based theorists have perhaps been wary of giving more of a story here because they haven't grasped a truth-oriented way of explaining the significance of evidence-responsiveness that doesn't hasten collapse into reliabilism. But without a more fundamental rationale, their theory seems at best to be a factoral account, in need of foundations. Epistemic Kantianism provides these foundations.

A similar point holds for virtue epistemology. Notwithstanding the title of their theory, virtue epistemologists have not shown the distaste for systematic theorizing and principles that characterizes so much work in virtue ethics. Virtue epistemology of the ambitious form defended by Zagzebski (1996) and Sosa (2007) is no anti-theory position: on the contrary, it has in these ambitious forms always been motivated by its ability to give a unifying and elegant solution to a wide range of traditional problems in epistemology. Yet virtue epistemology hasn't been adequately motivated as a foundational theory. Both Sosa and Zagzebski understand epistemic virtues by means of a certain relation to the fundamental value of truth: for Sosa, the relation is broadly instrumental (see Sosa (2007: Ch.4)), while for Zagzebski it is motivational (see, e.g., Zagzebski (2003)).

Epistemic Kantianism provides a more fundamental rationale for the role of epistemic virtue: believing from epistemic virtue just is the way to respect truth, and it matters for that reason. This story seems more satisfying than an unsystematic appeal to virtue, especially when the virtue invoked is so transparently linked to the more fundamental value of truth. One could be forced to accept such an appeal if virtue epistemologists went the way of virtue ethicists, opposing the search for a foundational theory, commending sheer sensitivity to

relevant factors, and recommending the independent study of virtue. But arguments for this position seem sorely lacking. It has not taken off in the literature.²² Virtue epistemology is fruitful as a factoral theory, but needs deeper foundations unless a compelling argument against the search for foundations emerges.

In addition to being extensionally attractive and providing a deeper explanation than similarly extensionally attractive theories, Epistemic Kantianism is also consistent with wider attractive outlooks. It is, for example, not incompatible with naturalism. The factors that make for respectful belief-formation needn't be non-natural. The view also needn't be psychologically implausible, like responsibilist virtue epistemology allegedly is, in virtue of its emphasis on 'love' of truth.²³ For respect for accuracy needn't be understood as requiring anything like *passion* for or *love* of the truth. Instead, a plausible necessary and sufficient criterion for manifesting respect for accuracy is just the kind of compliance with evidence emphasized by independently defensible reasons-based views.

This response does raise a question worth addressing, however. Although it is not psychologically implausible to understand *justified belief* by appeal to evidence and reasons, one might doubt that all epistemology stands in the space of reasons. So even if one agrees that Epistemic Kantianism gets justification right, one might reasonably wonder how much of epistemology it can illuminate. One might accordingly doubt whether good cognition in general has much to do with respect for accuracy.

My response is to stress that not all of epistemology is obviously normative, including the central relation of knowing. Epistemic Kantianism is not an epistemology in the traditional sense, but rather a foundational first-order theory of epistemic normativity. We shouldn't expect such a theory to yield analyses of seeing that *p*, remembering that *p*, intuiting that *p*, or other factive mental states. If knowledge is just a more general factive mental state, the same could go for it. For if knowledge is a more general factive mental state, it's no clearer that it partly consists in justified belief than that seeing that *p* partly consists in justified

²²Even Baehr (2011)—who rejects the ambitious views of Sosa and Zagzebski—doesn't suggest abandoning systematic epistemology: he just thinks virtue's role in systematic epistemology is auxiliary rather than fundamental (though very important nonetheless). Only Code (1984, 1987) and Kvanvig (1992) experimented with the more radical sort of virtue epistemology, but this was before Zagzebski's and Sosa's projects came to fruition.

²³For this criticism of Zagzebski (1996)'s virtue epistemology, see especially Sosa (2001, 2015: Ch.2)

belief.

On the other hand, if one finds it compelling that there is a kind of knowledge that is normatively constituted, one should find it less obvious that such knowledge *isn't* a standing in the space of reasons. But then there is no objection to including this knowledge—*reflective knowledge*—within Epistemic Kantianism's remit. So, if the objector is worried about our take on knowledge, we can offer a two-pronged response. Either knowledge is normative or it isn't. If it isn't, one shouldn't expect Epistemic Kantianism to say anything about it. If it is, it's unclear why Epistemic Kantianism will have more difficulty saying plausible things about it than about justified belief. Kantians have a clear picture of reflective knowledge: it is belief whose accuracy manifests strong respect for accuracy.

The Special Desiderata. So much for Epistemic Kantianism's ability to capture the stock desiderata of extensional adequacy, simplicity, and explanatory power, and its advantages over the competition. The more important point is that Epistemic Kantianism captures the two more special desiderata mentioned in §1.3. Firstly, it captures the Deontic-Evaluative Alignment thesis. On Epistemic Kantianism, epistemically justified attitudes are also the epistemically best attitudes available. But this fact holds for a different reason than consequentialists assume: epistemically justified attitudes are the epistemically best available because they alone respect fundamental epistemic value. The Epistemic Kantian's basis for DEA is preferable to the epistemic consequentialist's, since hers also avoids the sorts of worries Berker (2013a,b) raised for attempts to tightly link epistemic value and justification for reasons mentioned in §1.3.

While Epistemic Kantianism avoids the bad predictions of non-gerrymandered epistemic consequentialisms, it does so while remaining consistent with the central importance of reliability and the minimization of expected inaccuracy. Indeed, Epistemic Kantianism *explains* their importance: they matter because they are entailed by manifestations of respect for accuracy. This fact allows Epistemic Kantianism to capture the other desideratum mentioned in §1.3, which was the ability to explain epistemic consequentialism's appeal. It seemed appealing because reliability and the minimization of expected accuracy are centrally impor-

tant; indeed, justified belief might require little more than both. This feature is, as far as I know, unique among existing non-consequentialist epistemologies. Responsibilist virtue epistemology lacks this feature if it is a genuine alternative to Epistemic Kantianism. And if evidentialism and coherentism are understood as foundational non-consequentialist epistemologies (as Berker (2013a) suggests), they also fail to satisfy this desideratum. On neither view is it clear why these properties should matter, or why many have been misled into favoring epistemic consequentialism.

Let me explain this point in more detail, beginning with the view's ability to explain the significance of reliability. The derivative significance of reliability follows from at least three parts of the account. Firstly and most obviously, in virtue of understanding reflective knowledge as belief that is accurate *in virtue of manifesting respect for accuracy*, the account predicts that reliability is necessary for knowledge. For it is not possible for a belief to be accurate *in virtue of* manifesting strong respect for accuracy without this belief being non-accidentally true, and hence reliably formed (in a broad sense).

Secondly, in virtue of understanding *ex post* justified belief as belief that manifests respect for accuracy, which is in turn understood in terms of *compliance* with evidence, the account entails that justified belief is itself non-accidentally successful relative to the evidence. This result is what enables the account to get better predictions about cases of improper basing than some (e.g., Goldman (2012: Introduction)) assume evidentialists get. Of course, this is not yet to say that justified belief must be produced by a process that is reliable across modal space or even in the world of belief-formation (rather than in the actual world, as Comesaña (2002)'s indexical reliabilism suggests). The reliability here is reliability in conforming to evidence, and whether this in turn entails Goldman-style process reliability in the world of belief-formation depends on factorial theorizing about the nature of evidence and what it takes to possess evidence.²⁴ Still, a clear upshot of cases like Huck Finn's is that strong respect for a value in general does require objectively likely success in doing what is required by this value in appropriate, non-misleading circumstances. Hence, Epistemic Kantianism's account of justified belief plausibly will—given the relevant first-order intuitions about respect—yield

²⁴For two accounts that would have this result, see Schmidt (forthcoming) and Sylvan (2015).

reliability in normal circumstances.

Finally, in addition to being able to secure reliability as an upshot and hence explain away the appeal of reliability-based theories, it is worth noting that Epistemic Kantianism gives a better explanation of the significance of reliability than reliabilists themselves offer. To be sure, the property of reliability that attaches to a belief-forming *process* can be explained by an accuracy-based axiology that takes true belief as a value that is to be promoted rather than respected. But when reliabilists emphasize reliability in analyzing justification and knowledge, what they really care about is a backwards-looking property of a doxastic attitude rather than the process itself: namely, the *ancestry* of the doxastic attitude in a reliable process.²⁵ At this stage, a puzzle from Zagzebski (1999) known as the *swamping problem* emerges for reliabilism. If truth is fundamentally to be promoted and the significance of reliability is merely instrumental to this more fundamental value, why should *being formed* by a reliable process *add* value to a belief if that belief is *already* true? The fact that the belief is already true would seem to *swamp* whatever value the fact of reliable ancestry might have contributed on its own. This problem simply doesn't arise for Epistemic Kantianism. For as Jones (1997: 425) in effect noted in foretelling Zagzebski's lesson, the problem turns entirely on thinking of the value of true belief as merely 'to be promoted', and of thinking of derivative significance relative to this value in instrumental terms. When reliable ancestry does add value—which is not always, as Bonjour (1980)'s unwitting clairvoyants suggest!—it is as a manifestation of respect.

Having seen how Epistemic Kantianism derives the significance of reliability, it is even easier to see how it derives the significance of minimizing expected inaccuracy. One couldn't have sufficient *apparent* evidence for a doxastic attitude if a competing attitude would less risk inaccuracy given that evidence. So, given that manifesting respect for accuracy constitutively requires complying with apparent evidence, it will also require minimizing the expected inaccuracy of the belief that is formed. Hence, there is a clear reason why this central value of epistemic consequentialism should matter.

Moreover, Epistemic Kantianism provides a better explanation of why it has real signif-

²⁵Goldman (2015) emphasizes this fact in responding to Berker (2013a)'s discussion of his views.

icance. Of course, there is a clear reason why minimizing *actual* inaccuracy should have real value on an accuracy-first view that takes truth as a value ‘to be promoted’. But we are here considering the significance of minimizing *expected* disvalue. Plausibly, however, it is not merely an apparently good thing from the epistemic point of view to avoid disrespecting accuracy by risking inaccuracy, but an actually good thing. Hence, it is not clear how an accuracy-first view that takes truth as a value ‘to be promoted’ can explain all the relevant intuitions. The natural strategy is to agree with the consequentialist view it at the factorial level but provide a different kind of explanation at the foundational level. According to Epistemic Kantianism, the reason why minimizing expected inaccuracy is a genuinely good thing from the epistemic point of view is that this is equivalent to manifesting respect for accuracy. Since it is right in a perspective-transcendent sense to respect values that demand respect, it follows that by doing the perspective-dependently right thing according to the aforementioned epistemic consequentialist, we are thereby doing something perspective-transcendently right. And it is clearly good to do what is perspective-transcendently right.

4 Further Arguments

Having given a general reason to take Epistemic Kantianism seriously, I will now give some more targeted arguments for preferring the view to epistemic consequentialism.

4.1 Argument from the Objective Significance of the Subject’s Perspective

The first argument I will present turns on Epistemic Kantianism’s uniquely compelling ability to explain the *significance of the subject’s perspective* in epistemology, by giving a foundational rationale for believing in ways that are justified by factors supervenient on this perspective: forming beliefs in such ways constitutes respect for accuracy, which is the fundamental epistemic value.²⁶ I will first argue that epistemic consequentialism cannot fully explain the

²⁶As I hope will become clear, I don’t use the term ‘perspective’ quite in the way that Bergmann (2006) carves out in articulating the main intuition behind internalism; in particular, there is no immediate connection to higher-order attitudes. Indeed, the notion of perspective that appears below could be given the externalist construal of Miracchi (2017, forthcoming).

significance of the subject's perspective. I will then argue more briefly that Epistemic Kantianism provides an elegant explanation of how the subject's perspective could have the sort of significance it seems to have, which is not merely subjective but rather objective or—as I will say in a more Kantian vein—perspective-transcendent. I take this to be a strong reason to prefer Epistemic Kantianism, though some consequentialists may be willing to abandon what I take to be the commonsense view about the significance of perspective.

A disclaimer is in order before I proceed. In giving this first argument, I will be assuming something reminiscent of what is claimed by some internalists in epistemology (though it is weaker in ways I will explain). But I will *not* be assuming that epistemology is internalist at the foundational level (nor will I be assuming internalism at all, for reasons that will become clear). Indeed, I suspect that any foundationally internalist epistemology would involve a deeply unsatisfying sort of navel-gazing that doesn't respect the world-directedness of inquiry. We care about heeding the evidence that is within our subjective perspective because we care about something that transcends that perspective—viz., truth. Offering internal justification-making features at the foundational level would fail to respect this fact. Still, there is a way to reconcile the fact that justification is *at the factorial level* perspective-dependent with its being oriented toward truth, which is not perspectival except in special cases. The way to do this is to understand the perspective-transcendent value toward which justification is oriented as a value to be respected. For respect is achieved by suitably responding to perspective.

It should come as no surprise that the solution to the problem of reconciling the perspectival character of justification with its perspective-transcendent significance is Kantian. Both Kant's theoretical philosophy and his practical philosophy were devoted to solving problems with this structure, though Kant's task in these cases was to explain how knowledge and normativity could be both 'spontaneous' and objective, and spontaneity is not an important consideration in my story. One could even understand the argument of this section as a kind of transcendental argument that takes the objective significance of perspective for granted as a fact to be explained, and seeks to show how this is possible by taking the value of accuracy as a value fundamentally to be respected. The *form* of the argument is hence more

reminiscent of Kant's theoretical philosophy (esp. the Transcendental Deduction),²⁷ though the *matter* is inspired by his practical philosophy, when understood in the value-based way suggested by Wood (1999).

Preliminaries: Justification and the Subject's Perspective . The subject's perspective matters in epistemology in a way that requires explanation on any view that seeks to explain justification in a wholly truth-oriented way at the foundational level. The strongest version of this idea is best illustrated by the lessons of two stock objections to simple reliabilist accounts of justified belief: Cohen (1984)'s 'new evil demon' problem and BonJour (1980)'s clairvoyance problem. As Cohen suggested, it is tempting to claim that intrinsic mental duplicates—e.g., you and your counterpart in another world who is deceived by a Cartesian demon, and whose belief-forming processes are unreliable *in situ*—would be justified in holding the same doxastic attitudes. And as BonJour suggested, subjects with beliefs formed by equally truth-conducive processes don't necessarily seem to be justified in believing the same propositions to the same degree; for example, Norman the unwitting clairvoyant, who believes on the basis of no evidence that I am in New York City, would not seem to be justified in believing that I am there to the same degree as a normal perceiver who sees me there. Both cases suggest in different ways that justification is intimately tied to the subject's perspective: you and the demon victim seem similarly justified because your perspectives are indistinguishable, while Norman and a normal perceiver would not seem equally justified because of a difference in perspective.

Now, there are stronger and weaker ways of reacting to these intuitions. The strongest way is to adopt an internalist, reasons-based account of justification according to which (i) justification is grounded solely in the intrinsic properties of the 'non-factive' mental states that a subject shares with their counterpart in radical skeptical scenarios, and (ii) justification requires reasons for belief, where these reasons are understood as corresponding to (contents of) non-factive mental states such as perceptual experiences. A slightly weaker way would

²⁷If 'objective' were read as *objectively normative* rather than *about objects*, Kant's description of his problem as that of explaining how certain parts of the subject's perspective 'could have objective validity' in *Critique of Pure Reason* (A90/B122) would be more than just formally reminiscent of our problem.

be to drop the grounding thesis in (i) and simply uphold a supervenience thesis. This way is weaker because it is compatible with versions of reliabilism that index the sort of reliability required by justified belief to normal environments in the actual world, such as Comesaña (2002)'s indexical reliabilism and Sosa (1993)'s virtue reliabilism. These versions still claim that it is the reliability that this kind of belief-forming process has in the actual world that at least partly grounds the justification, rather than any intrinsic feature of the subject's overall mental state. It is also compatible with versions of virtue epistemology that ground justification in the 'inner seat' of a competence (see Sosa (2017: Ch.13)) that *would* yield reliability *when* installed in a suitable environment and when the subject is in suitable 'shape'. This view still claims that it is the fact that the seat is the seat *of such a competence* that grounds justification. A further way in which (i) could be weakened is to say that whether the subject has *enough justification to justify outright belief* supervenes on non-factive mental states, but to allow that the subject in the good case may have *even more justification*, as factive mentalists about justification like McDowell (1995) and Pritchard (2012) hold.

There is a yet weaker idea that one could invoke in defending Epistemic Kantianism. Instead of claiming that sufficient overall justification supervenes on the subject's perspective, one could merely claim that apparent reasons within the subject's perspective are (i) capable of giving real justifications of significant weight and not mere excuses or merely subjective reasons (*contra* Littlejohn (2012) and Williamson (forthcoming)), and (ii) capable of defeating the justification that would otherwise be provided by reliable processes. One could then allow that such justifications may fall short of establishing the all-things-considered objective permissibility of believing, relative to all the perspective-transcendent reasons. As Gardner (2007) noted in the philosophy of law, even if we distinguish between justification and excuse, we might deny that justifications are always sufficient to establish objective permissibility; Gardner hence allows for the possibility of justified wrongdoing that is more than merely excusable wrongdoing. In a similar vein, one might perhaps agree that it is always epistemically wrong to believe falsely or on the basis of unreliable processes, but insist that this wrongdoing can be justified and not merely excusable; one might perhaps also agree that there is a sense in which it is always correct to believe what is true, but insist that believing

in this way would be unjustified (and not merely show the subject to be blameworthy) if one lacked apparent evidence.

To avoid limiting the appeal of this argument to internalists, I will understand the objective significance of the subject's perspective in this final way:

OSP: The subject's perspective has objective normative significance in two ways:

- *Positive Significance:* If a subject S's apparent evidence recommends holding a doxastic attitude D toward p, then S has a real justification, not an excuse or merely subjective reason, for D(p)-ing. This is true even if S is in a skeptical scenario in which believing on the basis of the apparent evidence is not reliable *in situ*.
- *Negative Significance:* If S's apparent evidence fails to recommend D(p)-ing, then even if S has a reliable process available to her (e.g., clairvoyance) that would produce D(p), S would not be justified in D(p)-ing.

I will also assume that the Negative Significance half of OSP entails that apparent reasons are required for having a justification, since I take it that if one lacks apparent reasons for ϕ -ing, one's apparent reasons trivially fail to recommend ϕ -ing.

Many epistemologists would want to vindicate OSP, and those who haven't have felt an obligation at least to explain away the intuitions in favor of OSP by appealing to some normative category distinct from justification, like excusability (Littlejohn (2009); Williamson (forthcoming)), rationality (Littlejohn (2012)), or mere 'weak' justification (Goldman (1988)). And again, OSP is not incompatible with some externalist views about justification: there remains a question about the factorial ground of perspectival justification. Indexical reliabilists, virtue reliabilists, competence-based virtue epistemologists, and theorists who understand reasons to include factive mental states like seeing that p or the facts at which they are directed will not appeal purely to intrinsic features of a subject's non-factive mental states to explain these facts.

Some explanatory questions linger for all theorists (though they are easily answered by Kantians, as we will see). For theorists who embrace OSP, there is a question about *how* perspective can necessarily give real justifications rather than mere excuses. This question is particularly pressing if the fundamental epistemic value is accuracy: in cases where one's perspective is simply not a good guide to the truth, why think it could give justifications

rather than mere excuses? It is most pressing for reliabilists and virtue epistemologists like Sosa who think accuracy is the fundamental epistemic value and then understand other values in *instrumental* terms relative to this more fundamental value. But the question has broader importance. The role of the subject's perspective in epistemology distinguishes it in an important way from ethics, where it is harder to believe that something other than the mere subjective rightness or excusability of an action could be perspectival in quite this way. Even non-consequentialists will allow, after all, that actual consequences have some bearing on the actual justifiedness of an act, and that normative properties of interest to *foundational* normative theory don't supervene simply on intrinsic properties of the agent's mind and bodily movements. Of course, one could deny—as Goldman (1980) did—that the normative properties of interest to *foundational* normative epistemology really are perspectival. But even these theorists need to explain why normative properties that are perspectival have the kind of significance they do—and I don't think epistemic consequentialists can fulfill this explanatory task.

Epistemic Consequentialism Doesn't Adequately Explain OSP. I will now argue that epistemic consequentialism doesn't adequately explain—or have adequate resources to explain away—OSP. I will be focusing initially on Veritist versions of epistemic consequentialism. Because of its unpopularity, I will not be focusing on direct epistemic consequentialism. I will just be focusing on versions of indirect epistemic consequentialism. As a reminder, indirect epistemic consequentialism about justification claims in its most general form that:

EC_i The fundamental normative explanation of why justified beliefs are justified is that they are indirectly linked to producers of final epistemic value.

The standard version of EC_i takes the relevant 'indirect link' to be the *proximately caused by* relation, and takes the 'producers' to be belief-forming processes or methods (see Goldman (1986)), though one could easily imagine a version that took the producers to be larger faculties (e.g., perceptual systems): justified beliefs are those formed by truth-conducive processes, methods, or faculties. Let's call this standard version *standard* EC_i.

Standard EC_i provides no rationale for OSP, and indeed calls it into question. There are

several reasons why this is true. Firstly, in its standard form, EC_i requires justified beliefs to be linked to *real* producers of fundamental epistemic value *in the world in which the belief is held*, and OSP is inconsistent with such a requirement: a belief can be justified in a world even if it is not linked to such a producer in that world, at least if the fundamental epistemic value is truth. Secondly, in its standard form, EC_i does not require the producers in question to be or to depend on apparent reasons. Thirdly, it is natural to take EC_i to raise a doubt about why perspective should necessarily matter, given these facts: what essential role *could* it play in the explanation of justified belief, if the fundamental work is done by apparatus that can work independently of perspective?

Now, there are several ways in which standard EC_i could be revised in response to these points. But as we will see, these revisions seem *ad hoc*, and it is unclear what motivation there could be for including them in a *foundational* normative theory. I take it that this is part of why *standard* EC_i avoided such revisions. Now, perhaps—as we will consider later—one could try to explain away the significance of perspective by deeming it important only for non-foundational normative theory—e.g., in a theory of ‘decision-making procedures’ rather than ‘right-making characteristics’, in the language of Bales (1971) that Goldman (1980) once invoked, or in the epistemic analogue of a theory of subjective rightness in Smith (2010)’s sense. Making this move would, however, involve abandoning the attempt to directly explain OSP by EC_i itself.

A first way in which EC_i could be revised is to complicate the ‘indirect link’. Perhaps the link shouldn’t be between a belief held in a world and a producer of final epistemic value *in that world*, but rather between a belief held in a world and a producer of final epistemic value *in some distinguished world*. This is the sort of proposal that a friend of indexical reliabilism might make: if the kind of reliability that matters is reliability in normal circumstances in the actual world, then the link would have to be complicated in this way for indexical reliabilism to be subsumed under EC_i .

Unfortunately, it is difficult to buy this claim as a piece of *fundamental normative explanation*. Consider an analogy. When I ask for the fundamental normative explanation of why I ought to perform an action in this world, it would be strange to receive an answer referring

me to some *other* world and its properties. Why should I care about the goings-on of that world? How could those goings-on fundamentally bear on what I should do here? This kind of concern is especially pressing given that the world inhabited by the relevant cognitive agent—the demon world—is very far away. In justifying our actions, we do, of course, attend to mere possibilities in asking questions like ‘What if everyone did that?’, and one might even try to build an appeal to the reason brought out by the answer to such a question in a foundational normative theory in ethics. But the appeal of the reason is severely diminished if the world in which everyone did that is radically unlike ours, as the ‘demon world’ is supposed to be.

Of course, there might sometimes be a good rationale for thinking that the goings-on of a distant possible world do *bear* on the question of what I should do or believe here. Indeed, I might agree that it is a necessary condition for a belief to be justified in any world that the process by which it is formed be reliable in the actual world. But this is because it falls out of my theory of strong respect for accuracy together with an independently plausible account of the evidence to which one must respond in respecting accuracy. It is not because this necessary condition belongs in the statement of the foundational theory.

A revised version of EC_i that incorporated it would hence be very unsatisfying. Perhaps one could be pressured into accepting it if there were no alternative. But it is very hard to believe that this version of EC_i provides a *good foundational explanation* of the justifiedness of the beliefs of someone in a skeptical scenario. This point is strengthened by the fact that the most plausible contemporary versions of reliabilism are not offered as essentially consequentialist theories: the rationales that these versions give for restrictions of this kind do not involve commitment to the fundamental normative explanatory significance of facts about reliability in other worlds.

This point is the most pressing reason for thinking that even a non-standard version of EC_i will be unable to explain OSP. But notice that even setting this problem aside, there remains a problem about the role of reasons. Here again defenders of EC_i could stipulate that reasons are always the relevant ‘producers’ of the doxastic attitudes in question. But this would leave us with a question: what consistent consequentialist rationale can be given for

restricting the producers to reasons in this way? Without an answer to this question, the revision seems clearly *ad hoc*.

There are two other kinds of revisions that should be mentioned. Firstly, one might imagine someone who is attracted to *subjective* consequentialism being tempted to stipulate that the relevant producers must only be *expected* producers in the world of belief-formation. Qualified in this way, the view would avoid the issues raised by the demon world, and might explain intuitions about unwitting clairvoyants without invoking reasons by saying that the clairvoyant's belief isn't formed by a process that is an expected producer of true beliefs. Unfortunately, however, it is again difficult to understand how this response helps with the fundamental question, rather than just pushing the question back a step. For it is difficult to see why a *foundational* normative theory should be appealing to expected value rather than actual value: this is a notion that figures only in ethical analogues of theories of decision procedures rather than right-making characteristics, and we are considering theories of right-making characteristics. The significance of expected reliability seems on a par with the significance of perspective. So this theory in effect builds perspective in at the foundational level. Perhaps, again, we might be forced to do this. But it would be preferable if we could give some deeper explanation of the significance of perspective.

The same point undermines the other revision worth mentioning, which is to revise the underlying theory of fundamental epistemic value. An obvious way to sidestep these questions would be to take justification of the sort that is perspective-dependent and add it to the list of fundamental epistemic values. But it is hard to believe that justification is a fundamental epistemic value: we seem to care about it because we care about something else (though on my view this 'because' isn't merely instrumental).

Now, there are other ways to reject Veritism than to add justification to the list of fundamental epistemic values. One could add knowledge to the stock, for example. But doing so will not help with the problem raised by the demon world. Knowledge is factive. For this reason, the relevant processes in the demon-world will not be conducive to fundamental epistemic value if knowledge is the fundamental epistemic value.

Two Further Moves Blocked. I conclude that epistemic consequentialism lacks a satisfying direct explanation of OSP. At this stage, there are two further moves the epistemic consequentialist could make, apart from rejecting OSP (as, admittedly, some epistemic consequentialists may be happy to do). One move would be to say that the intuitive appeal of OSP confuses justification of the kind analyzed by a foundational theory of right-making characteristics with some secondary normative property—perhaps a property that belongs to the theory of decision-making procedures or the theory of subjective rightness. Goldman (1980) anticipated a move of this kind in arguing against internalism about justified belief.

My main trouble with this response is that it ultimately seems to be another way of deferring the important question, which we can raise again in slightly modified terms. Suppose we agree with the critic that it is not justification but rather some other normative property that is perspective-dependent in the manner of OSP. Let's call this property 'rationality'. There is now a descendant of the intuition about the demon world that remains problematic for the modified version of OSP. The modified intuition is that there is something of *perspective-transcendent normative importance* about having this property of rationality even in the demon world, not something merely of *apparent* or *subjective* normative importance. Views that simply write off epistemology's perspective-dependent normative property as a mere analogue of subjective rightness in ethics fail to explain this residual intuition: what is perspective-dependently right in epistemology also seems to have perspective-transcendent importance, even in the demon world.

Now, the version of the response that appeals instead to decision procedures can explain in the actual world why doing what is subjectively right in general has objective significance: this is objectively the best way for a cognitively limited subject to bring it about that she holds the objectively right attitudes. But this claim is not true in the demon world, nor does it seem relevant that the demon worlder is using methods that *would* be good guides if she *were* in a different possible world. Ultimately, one might simply deny that rationality necessarily has perspective-transcendent normative significance. I would, however, prefer a theory that captures the intuition that it does—and this is why I prefer Epistemic Kantianism, as we will see.

This point also addresses one other question that the neutral reader might have had after considering the complaint that epistemic consequentialism fails to provide a direct explanation of the intuition. Remember that the key point behind that complaint was that the revised versions needed to address the intuition didn't seem plausible candidates for *foundational* normative theories. At this point, one might question our starting assumption that epistemic consequentialism is best understood as a foundational normative theory and say that the objection has no force if it is considered instead as a factoral theory. This claim would, however, be moot given that the aim of this paper is to defend Epistemic Kantianism as a foundational normative theory. If the best I can show is that it is the only foundational normative theory that directly explains OSP, I will have done my work.

Epistemic Kantianism Explains OSP. So let's consider how Epistemic Kantianism explains OSP. The story is extremely simple. According to Epistemic Kantianism, the fundamental normative explanation of why justified beliefs are justified is that they manifest respect for accuracy, which is the fundamental epistemic value. Respect for accuracy, in turn, coincides with believing in accordance with one's apparent evidence, which is carried over to the demon world. So it is clear according to Epistemic Kantianism why there is a genuine justification provided by apparent reasons in this world: the fundamental epistemic value of accuracy calls in an objective, not merely subjective, way for respect, and hence giving this respect has perspective-transcendent significance.

4.2 Argument from the Swamping Problem

I take the fact that epistemic consequentialism fails to fully explain the significance of perspective to provide sufficient reason to reject it, and the fact that Epistemic Kantianism explains the significance of perspective to provide strong reason to accept it. But the intuitions behind PCJ and the significance of perspective aren't sacrosanct. Contravening them mainly represents a significant cost of epistemic consequentialism.

But I think there is a more fundamental argument against epistemic consequentialism, which is that it is either false or unmotivated given an independently defensible view about

epistemic value. Since Epistemic Kantianism is both compatible with and motivated by this view, we have reason to accept it.

The view about epistemic value that I have in mind involves two ideas. One is to follow a long line of axiologists in denying that all derivative value is instrumental value ('Instrumentalism about Derivative Value (IDV)'). I won't fully rehearse why we should do so in the epistemic case, since this has recently been done at great length by Sylvan (2018). But I'll give a brief summary of his reasoning.

It is plausible that justification has only *derivative* value relative to a more fundamental epistemic value like accuracy. But if justifying processes have *only instrumental* value relative to accuracy, is unclear why being justified would make a belief better if it is already true. As epistemic value theorists have grown fond of saying, the value contributed by the belief's already being true *swamps* the instrumental value that attaches to being justified relative to accuracy. For this reason, the problem of explaining how justification could add value to true belief if it is only derivatively valuable is a version of the *swamping problem* mentioned in the previous section.²⁸

As Jones (1997) emphasized in the earliest statement of this problem, it is not merely a problem about knowledge. A true belief's being justified makes it epistemically better regardless of whether it is Gettiered. We need to explain that. It is hard to see what that explanation can be assuming IDV. But if we reject IDV, we can solve the swamping problem and keep our epistemic axiology simple. The thought, in short, is that IDV is the real source of the swamping problem.

To reject IDV is to deny that all derivative epistemic value is merely instrumental. But it is not obvious that any further form of derivative epistemic value will help with the swamping problem. As Sylvan (2018) suggests, a model from Hurka (2001) seems needed. According to Hurka, ways of valuing more fundamental values are themselves derivatively but non-

²⁸Understood as a problem about the greater value of knowledge over true belief given a JTB+ account of knowledge and an instrumental account of the value of the J and the +-factor, this problem was made famous by Zagzebski (1999) as a problem for reliabilism and given the name by Kvanvig (2003). But it was anticipated in a more general form by Jones (1997) and has come again to be appreciated as a broader problem through the work of Pritchard (2010, 2011). It is not just a problem for reliabilism but for all views that take the value of justification and other necessary conditions for knowledge beyond true belief to have merely instrumental value. In fact, as Sylvan (2018: 383) notes in responding to an attempt to dismiss the problem by Carter and Jarvis (2012)'s, it is a problem for all views that take derivative epistemic value to be merely instrumental.

instrumentally valuable. Plausibly, manifestations of these ways of valuing have the same kind of value: if it is non-instrumentally good to value X, it is also non-instrumentally good to manifest that valuing by acting or thinking in further ways (e.g., if it is non-instrumentally good to value equality, it also seems non-instrumentally good to manifest that valuing by treating others as equals).

This kind of derivative value isn't swampable. Beneficence performed not as a PR stunt but rather for the sake of beneficence is better, though the caring about beneficence that it manifests is only derivatively valuable relative to the value of beneficence. Similarly, a true belief that manifests respect for truth is better than a true belief that doesn't, though respect for truth is good only derivatively, relative to the value of truth. As Sylvan (2018) suggests, then, it seems that the best way to keep our axiology simple is to reject IDV and accept Hurka's model of derivative value.

This isn't obviously to reject epistemic consequentialism, as Sylvan noted: Hurka wanted to help consequentialists preserve the idea that virtue is non-instrumentally valuable. But I want to question whether we can use Hurka's view to solve our problem without also rejecting epistemic consequentialism. I think we cannot. To see why, we need to look more closely at the structure of his view, and how it should be generalized.

Note that Hurka doesn't talk about ways of valuing that are *called for* or *demand*ed by the relevant value, or that meet some other normative bar. What he rather does is suggest a recursive principle on which 'love' of the non-instrumentally good is itself non-instrumentally good. But this principle is obviously insufficiently general. Love is just one response that can be called for by value. And others seem non-instrumentally good in virtue of being called for by value. Plausibly, it is all and only the ways of valuing that are called for or demanded by the relevant value that derive non-instrumental value from that value. The addition of a normative qualifier is crucial, since not every pro-attitude toward the valuable is valuable. Some pro-attitudes make no sense when held toward some values. Consider something desirable, like pleasure. You could also have the pro-attitudes of admiration and worship toward it. But it doesn't seem good to *admire* or *worship* pleasure. Doing so seems confused and unjustified.

The next important point is that we cannot solve the swamping problem in Sylvan's way if the only ways of valuing we recognize are instrumental. One might count a disposition to produce something as a way to value it. Yet if a good attitude is produced by a mere disposition to produce good attitudes, it wouldn't thereby be better as such. Indeed, if production were the only response demanded by a value, it wouldn't seem to matter from this value's point of view how it is produced. This conclusion is suggested by Zagzebski's point about coffee. It is because good cups are merely to be produced that a good cup from an unreliable coffeemaker is just as good as one from a reliable coffeemaker.

Accordingly, if we want to use Hurka's model, we must invoke non-instrumental ways of valuing. For this reason, Sylvan appealed to the non-instrumental relations of *commitment* and *respect*. Even if one doesn't invoke respect, one will need some non-instrumental response in this ballpark to solve the swamping problem with the help of Hurka's model. On the face of it, however, to demand such a response just is to assume that there is a non-consequentialist demand. If that impression is accurate, it seems that if we reject IDV in the way needed to solve the swamping problem, we should also abandon epistemic consequentialism.

One might think this impression isn't accurate, and that consequentialists could accept my claims. Can't consequentialists just say that you *ought to promote the state of affairs of having respect for accuracy*? They could, but this order of explanation is implausible: the demand to bring it about that one values fundamental value is not a *fundamental* demand. Such a demand, if there is one, is a consequence of the more fundamental demand to respond to value only in justified ways. Matters would be different if the epistemic consequentialist denied that Hurkan value is *derivative*. If it were fundamental, there wouldn't be a problem, since there wouldn't be a more fundamental requirement of respect. But it isn't, so there is a problem.

Let's put this argument more officially. I'll call it the *Scanlonian Argument*, since it is inspired by Scanlon (1998: Ch.2)'s argument against consequentialism from the nature of value:

1. IDV is false: not all derivative value is instrumental value.

2. Hurka's model should be extended to epistemology as an illustration of (1): if X is fundamentally epistemically valuable, then valuing X is non-instrumentally but derivatively epistemically valuable.
3. If (1) and (2), then some fundamental epistemic values call for non-instrumental responses (e.g., respect).
4. If fundamental epistemic value calls for non-instrumental responses, then there is a demand to value fundamental epistemic value which is not instrumentally grounded.
5. If there is a demand to value fundamental epistemic value which is not instrumentally grounded, epistemic consequentialism is false.
6. So, epistemic consequentialism is false.

As I've explained, (1) and (2) have been defended compellingly elsewhere. I just explained why we should accept (3) and (4). And (5) follows from the way epistemic consequentialism was defined earlier.

It is a further step to get Epistemic Kantianism. But Epistemic Kantianism is among the simplest views consistent with the facts that drive this argument. These facts point to a demand to value fundamental epistemic value in some non-instrumental way. A demand of respect fits the bill. It would be unparsimonious to insist that the demand that the Scanlonian Argument tells us to acknowledge is a different one, since ours provides a full explanation of the facts that drive the argument. So, there is reason to favor Epistemic Kantianism given the Scanlonian Argument.

4.3 Argument from the Asymmetries

A final argument for Veritist Epistemic Kantianism is that it correctly predicts certain asymmetries in epistemic value and justification.

4.3.1 The Axiological Asymmetry

Consider how Epistemic Kantians should understand the value of doxastic states. Their core claim is that accuracy is a value fundamentally to be respected. How do they then derive conclusions about the derivative epistemic goodness or badness of doxastic states? Here is a natural suggestion:

Derivation through Respect: A cognitive state is *pro tanto* derivatively epistemically good if it manifests respect for accuracy, and *pro tanto* derivatively bad if it manifests disrespect for accuracy.

Derivation through Respect makes some important predictions. It predicts that

(a) an accurate belief in *p* won't be epistemically valuable insofar as it fails to manifest respect for accuracy vis-à-vis *p*,

and

(b) an inaccurate belief in *p* won't necessarily be epistemically disvaluable if it doesn't manifest disrespect for accuracy vis-à-vis *P*.

These predictions are plausible. Suppose *S* has overwhelming apparent evidence to doubt whether *p*, *S* believes *p* despite the evidence, but *p* turns out luckily to be true. Is there anything good about *S*'s belief? 'No' is a more intuitive answer than 'Yes'. When accurate believing would also be utterly careless believing, it wouldn't be good believing in the relevant sense. Epistemic Kantianism correctly predicts this conclusion.

If this conclusion doesn't strike one as obviously correct, one is, I suspect, confusing what is good and epistemic with what is epistemically good. True beliefs are often good: they get us to Larissa. Indeed, there is something plausibly valuable in accurately rather than inaccurately representing reality. Since true beliefs are epistemic items, I also agree that they are good and epistemic. But it doesn't follow that they are always epistemically good. For my own part, once I distinguish these questions, it is highly intuitive that making a careless, reflectively lucky guess is not good believing at all.

One might wonder how these claims are consistent with an accuracy-first approach. How can we consistently permit or even require forming inaccurate beliefs if accuracy is the fundamental value and we ought to respect this value? Well, the Epistemic Kantian thinks that we ought to constrain our doxastic deliberation by respect for accuracy. Yet from the fact that

we ought to constrain our deliberation by respect for a value *V*,

it does not follow that

we ought *always* to promote *V* or avoid promoting the corresponding disvalue.

This inference fails when we consider values that are fundamentally to be respected. In such cases, if the apparent evidence misleadingly suggests that we would act against the relevant value by A-ing but conform by B-ing, respect for the value will demand B-ing rather than A-ing even if B-ing would fail to promote the value. Respect requires constraining one's deliberation by the value. It does not straightforwardly require promoting the value—not in all cases. Consider what the value of commitment to a relationship requires in the following sort of case:

(DISGUISE) A and B have agreed to have a monogamous relationship. But A worries that B would cheat if B got the chance. A decides to test this hypothesis. With the help of some extraordinary costuming, A manages to dress up like a totally different person on whom B would have an instant crush. Disguised, A has been showing up around B's workplace to make advances. B believes on the basis of this misleading evidence that this is a fascinating person distinct from A. B now seems to be having a date with this distinct person when A had planned to be out of town....

B plausibly manifests inadequate commitment for the relationship into which A and B contracted. Part of the reason why is that B fails to constrain their practical deliberation by the value of commitment to the contract. Imagine furthermore that A has a lookalike who has been playing A's role while A has been playing the distinct inamorata. If B later tries to be faithful to A, B might then fail to promote the value of commitment to one's relationships. If B ought to constrain their practical reasoning by the value of commitment, because doing so is necessary for having the kind of commitment demanded by the relationship, and B oughtn't actually to do what, objectively speaking, would promote commitment to a relationship. For this reason, it is implausible to claim that B is relevantly lacking in commitment if B instead does attempt to manifest commitment to A and fails only due to the misleading evidence about which person is A.

I don't think our intuitions are simply conflating the subjective 'ought' with a more objective 'ought' here. What there is most objective, value-based reason for B relative to the value of commitment to do is respect this value. To do that, B must govern their deliberations by the value of commitment, and so do what would ironically fail to promote commitment in the extended case. If Epistemic Kantianism is true, a parallel conclusion holds in epistemology. To respect accuracy, one must constrain one's deliberations by the value

of accuracy. Since respect is the response fundamentally demanded by accuracy's epistemic value, one ought sometimes not to promote accuracy. This will be true when promoting accuracy would require disrespecting accuracy, and so carelessly failing to have one's doxastic deliberation constrained by the value of accuracy.

4.3.2 The Deontic Asymmetry

There is a related argument worth discussing separately, which I previewed in §2.1. As we saw earlier, there is pressure to reject

Strong T-Obligation: We ought epistemically to produce accurate doxastic attitudes.

and replace it with something like

Weak T-Obligation: We ought epistemically to ensure the accuracy of our existing doxastic attitudes.

Now, is this obligation best explained in instrumental terms, so that the response fundamentally linked with accuracy's epistemic value is indeed *promotion*—just promotion of a conditional state of affairs (believing only if true) rather than an unconditional one (believing the true)? This is implausible as a bedrock explanation. The analogy with ethics is suggestive. Recall the need to replace

Strong P-Obligation: We ought to cause people to exist who are treated well.

with something like

Weak P-Obligation: We ought to treat existing people well.

Our desire to switch isn't brute. But it also isn't best explained in instrumental terms. It can and should be explained by a reversal in the priority structure among pro-responses connected with personhood's value. It is the fact that personhood more fundamentally calls for respect than promotion that explains the asymmetry between our reasons to bring about states of affairs in which people are treated well and our reasons to treat existing people well. This reversal is unsurprising. If the fundamental response associated with personhood's value were the teleological one of promotion, we wouldn't expect this asymmetry. Since this

asymmetry exists, we should expect a corresponding asymmetry in the priority structure among the responses connected with personhood's value.

Now we can see an argument for Epistemic Kantianism. It is odd to think that accuracy is fundamentally to be promoted given the asymmetry to which we have drawn attention. This asymmetry calls out for explanation just like Narveson's asymmetry in population ethics. One might try to explain it by denying that accuracy is an epistemic value at all. One might also try to take the asymmetry as a bedrock fact. The first move would be a mistake. The second should be avoided unless we lack a clear alternative. It is better, then, to deny that accuracy's epistemic value is to be understood teleologically. So we should seek a non-consequentialist explanation. Epistemic Kantianism provides one, which is a reason to believe it.

5 Concluding Remarks

Let's conclude by taking stock and recalculating the score in the debate.

Although epistemic non-consequentialism's core claim is negative, my goals have been positive. I have developed an alternative to epistemic consequentialism that shares the unity and simplicity of the latter's most familiar version, and agrees with this version that accuracy is the fundamental epistemic value. Where Epistemic Kantianism parts company is in its conception of what this value fundamentally demands—viz., respect. As I have argued, this departure bears significant fruit: it helps explain asymmetries in epistemic value and duty, and shows how a perspectival kind of justification might derive its normativity from the value of accuracy despite not being truth-conducive in all worlds. This departure is also independently justified: to solve the swamping problem, we already need an axiology which sees accuracy not merely as a value 'to be promoted', and this admission points toward Epistemic Kantianism.

As we've seen, adopting Epistemic Kantianism is not only *compatible* with the importance of reliability and the minimization of expected inaccuracy, but *explains* their importance. Even if the demon-worlder equally respects accuracy, respect might still require the

indexical reliabilist's sort of reliability, and require forming beliefs only if the evidence suggests one wouldn't thereby risk their inaccuracy. The importance of both ideals follows given a factoral gloss on respect in terms of responsiveness to apparent evidence and an independently defensible account of the nature of apparent evidence.

The belief that abandoning consequentialism would require abandoning these ideals might explain why the main reaction to the anti-consequentialist backlash has been to revise epistemic consequentialism. But this belief is indefensible if one avoids conflating foundational first-order theorizing with factoral theorizing or metaphysical analysis, and notices that non-accidentality is also a centerpiece of one non-consequentialist tradition (viz., Kantianism). If we avoid such conflation and see the possibility of giving Kantian foundations for these ideals, there is little reason to worry that adopting epistemic non-consequentialism will leave us with buyer's remorse. Epistemic Kantianism shares the attractions epistemic consequentialists claim for their side: it provides a simple and unified explanation of our epistemic intuitions, and preserves the link between epistemic value and justification. But by understanding the link between epistemic value and justification in a fundamentally non-instrumental way, Epistemic Kantianism avoids the worries that have driven some away from epistemic consequentialism, and explains the significance of perspectival justification.

For these reasons, Epistemic Kantianism arguably stands as the superior option. At the very least, we see that epistemic non-consequentialism needn't be a reactionary last resort but a highly appealing view, as this paper has aimed to show.

6 References

- Ahlstrom-Vij, K. and Dunn, J. 2014. 'A Defence of Epistemic Consequentialism.' *Phil. Quarterly* 64: 541-551.
- Ahlstrom-Vij, K. and Dunn, J. Forthcoming. 'Is Reliabilism a Form of Epistemic Consequentialism?' *American Phil. Quarterly*.
- Alston, W. 2005. *Beyond Justification*. Ithaca: Cornell UP.
- Anderson, E. 1993. *Value in Ethics and Economics*. Cambridge, MA: Harvard UP.
- Anderson, E. 1999. 'What is the Point of Equality?' *Ethics* 109: 287-337.
- Armstrong, D. M. 1973. *Belief, Truth, and Knowledge*. Cambridge: CUP.
- Arpaly, N. 2003. *Unprincipled Virtue*. OUP.
- Baehr, J. 2011. *The Inquiring Mind*. OUP.

- Bales, R. E. 1971. 'Act Utilitarianism: Account of Right-Making Characteristics or Decision-Making Procedure?' *American Philosophical Quarterly* 8: 257-265.
- Baron, M. 1995. *Kantian Ethics Almost without Apology*. Ithaca: Cornell UP.
- Benson, P. 1987. 'Moral Worth.' *Phil. Studies* 51: 365-382.
- Beddor, R. 2017. 'Justification as Faultlessness.' *Philosophical Studies* 174: 901-926.
- Bergmann, M. 2006. *Justification without Awareness*. OUP.
- Berker, S. 2013a. 'Epistemic Teleology and the Separateness of Propositions.' *Phil. Review* 122: 337-393.
- Berker, S. 2013b. 'The Rejection of Epistemic Consequentialism.' *Phil. Issues* 23: 361-387.
- Berker, S. Forthcoming. 'The Unity of Grounding.' *Mind*.
- BonJour, L. 1980. 'Externalist Theories of Empirical Knowledge.' *Midwest Studies in Philosophy* 5: 53-73.
- BonJour, L. 1985. *The Structure of Empirical Knowledge*. Cambridge MA: Harvard UP.
- Bykvist, K. and Hattiangadi, A. 2007. 'Does Thought Imply Ought?' *Analysis* 67: 277-85.
- Carnap, R. 1947. *Meaning and Necessity*. University of Chicago Press.
- Carter, J. A. and Jarvis, B. 2012. 'Against Swamping.' *Analysis* 72: 690-9.
- Code, L. 1984. 'Toward a 'Responsibilist' Epistemology.' *Philosophy and Phenomenological Research* 45: 29-50.
- Code, L. 1987. *Epistemic Responsibility*. Hanover, NH: University Press of New England.
- Comesaña, J. 2002. 'The Diagonal and the Demon.' *Philosophical Studies* 110: 249-266.
- Conee, E. and Feldman, R. 2004. *Evidentialism*. OUP.
- Cohen, S. 1984. 'Justification and Truth.' *Philosophical Studies* 46: 279-95.
- Darwall, S. 1977. 'Two Kinds of Respect.' *Ethics* 88: 36-49.
- Darwall, S. 2006. *The Second-Person Standpoint*. Cambridge MA: Harvard UP.
- Dean, R. 2006. *The Value of Humanity in Kant's Moral Theory*. OUP.
- Dreier, J. 1993. 'The Structure of Normative Theories.' *The Monist* 76: 22-40.
- Dretske, F. 1981. *Knowledge and the Flow of Information*. Cambridge, MA: The MIT Press.
- Ewing, A. C. 1947. *The Definition of Good*. New York: Macmillan.
- Fine, K. 2002. 'The Varieties of Necessity' in Gendler, T. and Hawthorne, J. (eds.) *Conceivability and Possibility*. OUP.
- Fine, K. 2012. 'A Guide to Ground' in Correia, F. and Schnieder, B. (eds.) *Metaphysical Grounding*. Cambridge: Cambridge UP.
- Foley, R. 1993. *Working without a Net*. OUP.
- Foley, R. 2004. 'A Trial Separation of the Theory of Knowledge and the Theory of Justified Belief' in Greco, J. (ed.) *Ernest Sosa and his Critics*. Blackwell.
- Fumerton, R. 1995. *Metaepistemology and Skepticism*. Lanham: Rowman and Littlefield.
- Gardner, J. 2007. *Offences and Defences*. OUP.
- Geach, P. 1956. 'Good and Evil.' *Analysis* 17: 33-42.
- Gibbons, J. 2013. *The Norm of Belief*. OUP.
- Goldman, A. I. 1967. 'The Causal Theory of Knowing.' *The Journal of Philosophy* 64: 357-372.
- Goldman, A. I. 1975. 'Innate Knowledge' in Stich, S. (ed.) *Innate Ideas*. University of California Press.
- Goldman, A. I. 1976. 'Discrimination and Perceptual Knowledge.' *Journal of Philosophy* 73: 771-791.
- Goldman, A. I. 1979. 'What Is Justified Belief?' in Pappas, G. (ed.) *Justification and Knowledge*. Dordrecht: D. Reidel.
- Goldman, A. I. 1980. 'The Internalist Conception of Justification.' *Midwest Studies in Philosophy* 5: 27-51.
- Goldman, A. I. 1986. *Epistemology and Cognition*. Cambridge, MA: Harvard UP.
- Goldman, A. I. 1988. 'Strong and Weak Justification.' *Philosophical Perspectives* 2: 51-69.
- Goldman, A. I. 1999. 'Internalism Exposed.' *Journal of Philosophy* 96: 271-293.
- Goldman, A. I. 2002. 'The Unity of the Epistemic Virtues' in *Pathways to Knowledge*. OUP.

- Goldman, A. I. 2011. 'Toward a Synthesis of Evidentialism and Reliabilism?' in Dougherty, T. (ed.) *Evidentialism and its Discontents*. OUP.
- Goldman, A. I. 2012. *Reliabilism and Contemporary Epistemology*. OUP.
- Goldman, A. I. 2015. 'Reliabilism, Veritism, and Epistemic Consequentialism.' *Episteme* 12: 131–143.
- Graham, P. 2014. 'A Sketch of a Theory of Moral Blameworthiness.' *Philosophy and Phenomenological Research* 88(2): 388-409.
- Grimm, S. 2009. 'Epistemic Normativity' in Haddock, A., Millar, A. and Pritchard, D. (eds.) *Epistemic Value*. OUP.
- Guyer, P. 2000. *Kant on Freedom, Law, and Happiness*. Cambridge: CUP.
- Hare, R. M. 1981. *Moral Thinking*. OUP.
- Herman, B. 1981. 'On the Value of Acting from the Motive of Duty.' *Phil. Review* 66: 359-382.
- Herman, B. 1993. *The Practice of Moral Judgment*. Cambridge, MA: Harvard UP.
- Herman, B. 2007. *Moral Literacy*. Cambridge, MA: Harvard UP.
- Huemer, M. 2007. 'Compassionate Phenomenal Conservatism.' *Philosophy and Phenomenological Research* 74: 30-55.
- Hurka, T. 2001. *Virtue, Vice and Value*. OUP.
- Joyce, J. 1998. 'A Nonpragmatic Vindication of Probabilism.' *Philosophy of Science* 65: 575–603.
- Joyce, J. 2009. 'Accuracy and Coherence: Prospects for an Alethic Epistemology of Partial Belief' in Huber and Schmidt-Petri (eds.) *Degrees of Belief*. Springer.
- Jones, W. 1997. 'Why Do We Value Knowledge?' *American Phil. Quarterly* 34: 423-439.
- Kagan, S. 1992. 'The Structure of Normative Ethics.' *Phil. Perspectives* 6: 233-242.
- Kagan, S. 1997. *Normative Ethics*. Westview Press.
- Kagan, S. 1999. 'Rethinking Intrinsic Value.' *Journal of Ethics* 2: 277–97.
- Kamm, F. 1992. 'Non-Consequentialism, the Person as an End-in-Itself, and the Significance of Status.' *Philosophy and Public Affairs* 21(4): 354-389.
- Kant, I. 1781/1999. *Critique of Pure Reason*. Cambridge: CUP.
- Kant, I. 1785/2012. *Groundwork for the Metaphysics of Morals*. Cambridge: CUP.
- Kant, I. 1797/1996. *The Metaphysics of Morals* in McGregor, M. (tr./ed.) *Practical Philosophy*. Cambridge: CUP.
- Korsgaard, C. 1983. 'Two Distinctions in Goodness.' *Philosophical Review* 92: 169–95.
- Korsgaard, C. 1996. *Creating the Kingdom of Ends*. Cambridge: CUP.
- Korsgaard, C. 2008. *The Constitution of Agency*. OUP.
- Kvanvig, J. 1992. *The Intellectual Virtues and the Life of the Mind*. Rowman and Littlefield.
- Kvanvig, J. 2003. *The Value of Knowledge and the Pursuit of Understanding*. Cambridge: CUP.
- Littlejohn, C. 2009. 'The Externalist's Demon.' *Canadian Journal of Philosophy* 3: 399-434.
- Littlejohn, C. 2012. *Justification and the Truth Connection*. Cambridge: CUP.
- Lord, E. 2018. *The Importance of Being Rational*. OUP.
- Lyons, J. 2009. *Perception and Basic Beliefs*. OUP.
- Mantel, S. 2018. *Determined by Reasons: A Competence Account of Acting for a Normative Reason*. London: Routledge.
- Markovits, J. 2014. *Moral Reason*. OUP.
- McDowell, J. 1995. *Mind and World*. Cambridge, MA: Harvard UP.
- McHugh, C. 2012. 'The Truth Norm of Belief.' *Pacific Philosophical Quarterly* 93: 8–30.
- Miracchi, L. 2017. 'Perspectival Externalism is the Antidote to Skepticism.' *Episteme* 14: 363–379.
- Miracchi, L. Forthcoming. 'Competent Perspectives and the New Evil Demon Problem' in Dorsch, F. and Dutant, J. *The New Evil Demon*. OUP.
- Moore G. E. 1903. *Principia Ethica*. Cambridge: CUP.
- Narveson, J. 1973. 'Moral Problems of Population.' *The Monist* 57: 62-86.
- Nozick, R. 1974. *Anarchy, State, and Utopia*. Cambridge, MA: Harvard UP.
- Parfit, D. 2011. *On What Matters*. OUP.
- Pettigrew, R. 2015. *Accuracy and the Laws of Credence*. OUP

- Pritchard, D. 2010. 'Knowledge and Understanding' in Pritchard D., Millar A., Haddock A. (eds.), *The Nature and Value of Knowledge: Three Investigations*. OUP.
- Pritchard, D. 2011. 'What Is the Swamping Problem?' in Reisner, A. and Steglich-Petersen, A. (eds.) *Reasons for Belief*. Cambridge: CUP.
- Pritchard, D. 2012. *Epistemic Disjunctivism*. OUP.
- Rabinowicz, W. and Rønnow-Rasmussen, T. 2000. 'A Distinction in Value: Intrinsic and For Its Own Sake.' *Proceedings of the Aristotelian Society* 100: 33–51.
- Rashdall, H. 1907. *The Theory of Good and Evil: A Treatise on Moral Philosophy*. Oxford: Clarendon Press.
- Raz, J. 1999. *Engaging Reason*. OUP.
- Raz, J. 2011. *From Normativity to Responsibility*. OUP.
- Scanlon, T. M. 1998. *What We Owe to Each Other*. Cambridge, MA: Harvard UP.
- Scheffler, S. 1982. *The Rejection of Consequentialism*. OUP.
- Schellenberg, S. 2013. 'Experience and Evidence.' *Mind* 122: 699–747.
- Schmidt, E. Forthcoming. 'Possessing Epistemic Reasons: The Role of Rational Abilities.' *Philosophical Studies*.
- Schroeder, M. 2015. 'Knowledge is Belief for Sufficient (Objective and Subjective) Reason.' *Oxford Studies in Epistemology* 5: 227–252.
- Smith, H. 2010. 'Subjective Rightness.' *Social Philosophy and Policy* 27: 64–110.
- Sosa, E. 1991. *Knowledge in Perspective*. Cambridge: CUP.
- Sosa, E. 1993. 'Proper Functionalism and Virtue Epistemology.' *Nous* 27: 51–65.
- Sosa, E. 2001. 'For the Love of Truth?' in Fairweather, A. and Zagzebski, L. (eds.) *Virtue Epistemology: Essays on Epistemic Virtue and Responsibility*. OUP.
- Sosa, E. 2007. *Apt Belief and Reflective Knowledge, Volume 1*. OUP.
- Sosa, E. 2015. *Judgment and Agency*. OUP.
- Sosa, E. 2017. *Epistemology*. Princeton UP.
- Sylvan, K. 2015. 'What Apparent Reasons Appear to Be.' *Phil. Studies* 172: 587–606.
- Sylvan, K. 2018. 'Veritism Unswamped.' *Mind* 127: 381–435.
- Stratton-Lake, P. 2000. *Kant, Duty, and Moral Worth*. London: Routledge.
- Sutton, J. 2007. *Without Justification*. Cambridge MA: MIT Press.
- Temkin, L. 1997. *Inequality*. OUP.
- Thomson, J. J. 2008. *Normativity*. Chicago: Open Court.
- Turri, J. 2010. 'On the Relationship between Propositional and Doxastic Justification.' *Philosophy and Phenomenological Research* 80(2): 312–326.
- Wedgwood, R. 2013. 'The Right Thing to Believe?' in Chan, T. (ed.) *The Aim of Belief*. OUP.
- Whiting, D. 2010. 'Should I Believe the Truth?' *Dialectica* 64: 213–24.
- Williamson, T. Forthcoming. 'Justifications, Excuses, and Skeptical Scenarios' in Dutant, J. and Dorsch, F. (eds.) *The New Evil Demon*. OUP.
- Wood, A. 1997. 'The Final Form of Kant's Practical Philosophy.' *Southern Journal of Philosophy* 36: 1–20.
- Wood, A. 1999. *Kant's Ethical Thought*. Cambridge: CUP.
- Zagzebski, L. 1996. *Virtues of the Mind*. Cambridge: CUP.
- Zagzebski, L. 1999. "From Reliabilism to Virtue Epistemology." *Proceedings of the Twentieth World Congress of Philosophy, Volume 5: Epistemology*. Bowling Green, Ohio: Philosophy Documentation Center.
- Zimmerman, M. J. 2001. *The Nature of Intrinsic Value*. Lanham, MD: Rowman and Littlefield.