

Reliabilism without Epistemic Consequentialism

Abstract

This paper argues that reliabilism can plausibly live without epistemic consequentialism, either as part of a non-consequentialist normative theory or as a non-normative account of knowledge on a par with certain accounts of the metaphysics of perception and action. It argues moreover that reliabilism should not be defended as a consequentialist theory. Its most plausible versions are not aptly dubbed ‘consequentialist’ in any sense that genuinely parallels the dominant sense in ethics. Indeed, there is no strong reason to believe reliabilism was ever seriously intended as a form of epistemic consequentialism. At the heart of its original motivation was a concern about the necessity of *non-accidentality* for knowledge, a concern quite at home in a non-consequentialist or non-normative setting. Reliabilism’s connection to epistemic consequentialism was an accretion of the ’80s, and a feature of only one of its formulations in that decade.

1 Introduction

Recent literature conveys the impression that epistemic consequentialism has long been a dominant view about justified belief.¹ One source of this impression is the prominence of reliabilism in the post-Gettier history of epistemology. While reliabilism has had many critics, it retains many adherents and remains a dominant perspective on the nature of justified belief and knowledge; indeed, it is arguably the leading form of externalism. Accordingly, if reliabilism were a form of epistemic consequentialism, one could reasonably conclude that the latter has been a major force in traditional epistemology since the late ’60s and early ’70s.

But such thinking would be misguided. Reliabilism in its early formulations was not offered as a consequentialist theory, or indeed as a *normative* theory at all. While it did take on a relevantly normative formulation in one classic work nearly twenty years after its emergence (viz., Goldman (1986)), this formulation was an accretion, and lacks a claim to dominance. Moreover, reliabilism’s apparently consequentialist formulations aren’t really analogous to consequentialism in ethics, nor are they best defended in the way in which consequentialist positions are best defended.

Such, at any rate, are some main claims to be defended below. More briefly, my goal is as follows. After documenting the case for my historical claims, I will argue that whether or not reliabilism was ever seriously intended as consequentialism, it

(C1) can naturally be defended without epistemic consequentialism, and

(C2) should be defended without epistemic consequentialism, since its most plausible forms are not properly regarded as consequentialist.

¹See especially Berker (2013a).

I will begin with some preliminary clarifications in §2, defend (C1) and (C2) in §3–5, outline a positive approach in §6, and take stock and conclude in §7.

2 Consequences and the Structure of Normative Theory

Before I proceed, I want to briefly explain what kinds of theories I would count as consequentialist, since my arguments will rely on some assumptions here that are familiar in normative ethics but less familiar in epistemology.

Consequentialism is, I take it, a first-order normative theory, not a meta-normative theory.² Since there are different kinds of first-order theory, it is worth considering in more detail what kind consequentialism is. Following Kagan (1992, 1997), we should distinguish *factoral* and *foundational* normative theories. To understand the point of the first kind of theory, note that many normative statuses (rightness, value, justification...) even within a domain (prudential, moral, epistemic...) are a function of various relevant factors. Factoral theorists seek, in Kagan (1997: 17)'s words, to 'articulat[e] these various normative factors, and discover how they interact' to determine statuses like rightness, value and justification. In other words, factoral theories seek to answer the following question of some normative status *S* in a domain *d*:

(Q1) **The Normative Factors Question:** What factors bear on whether a target of evaluation in *d* has *S*? What, in other words, are the *S-relevant factors*?

Although factoral theories often posit several factors, they can be monistic. One might, for example, claim that the only factor that bears on whether one's act is prudentially right is its contribution to one's pleasure/pain ratio. This claim can be offered factorally rather than foundationally because one could conceivably go further, adopting a more fundamental but still normative explanation of why one's pleasure/pain ratio is the only prudentially relevant factor. For example, one might defend hedonism about prudential rightness on the threefold basis of

- (a) agent-relative act consequentialism about prudential rightness,
- (b) a welfarist view about prudential goodness, and
- (c) a hedonist account of welfare.

At the foundational level, the theorist offering this position is a consequentialist about prudential rightness and a welfarist about goodness-for. This theorist happens also to believe that welfare is grounded in pleasure, and hence thinks pleasure is the only factor relevant to prudential rightness.

Despite not being meta-normative positions, foundational normative theories do more than answer (Q1). They seek to answer:

(Q2) **The Normative Foundations Question:** What is the fundamental normative explanation of why the *S-relevant factors* are *S-relevant factors*?

²Though see Copp (2009, 2014) for a different approach.

It is natural to assume that we can distinguish normative explanations from other kinds of explanations, just as we can distinguish normative necessity and other kinds of necessity (as Fine (2002) does).³ Hence the relevant answer to (Q2) will differ from the answer provided by meta-normative theories like the Humean theory of normative reasons, which seek a metaphysical grounding in naturalistic terms of some normative property or relation.

It is one thing to *analyze* prudential rightness in terms of welfare maximization and another to say that the most fundamental *normative explanation* of why certain acts are right is that they are welfare-maximizing. Foundational normative ethics aims to provide the most fundamental normative explanations of why various normatively relevant factors are the relevant ones. But not all argumentation in normative ethics occurs at this level. Much just concerns which factors are morally relevant. Consider partiality: can the fact that I am in some relationship with X affect how weighty a moral reason I have to benefit X? It is unclear that answering this question requires having already taken a stand on the correct foundational normative theory.⁴

I take consequentialism in its most interesting form to be a foundational first-order theory, not a factorial theory. There remain, of course, many varieties of consequentialism—direct vs. indirect, maximizing vs. satisficing, welfarist vs. non-welfarist, etc. While I suspect the only coherent way to be a consequentialist is to be a direct consequentialist, I will not rely on this suspicion, keeping various indirect versions—e.g., rule and motive consequentialism—on the table.

When consequentialism is understood in this way, it is possible for some views to *look* consequentialist from a certain angle without *really being* consequentialist. One way to see this point is to suppose for the sake of argument that Parfit (2011) is right that the best forms of rule consequentialism, Kantianism, and contractualism make the same predictions. Even if Parfit were right, there would remain a question about which theory provides the best foundational explanation. One can imagine holding that while Kantianism and rule consequentialism extensionally coincide, the deeper reason why the acts that rule consequentialism deems right are right is Kantian.

While I am unsure Parfit is right about the extensional convergence of these theories, I suspect there is a compelling case for a relevantly similar conclusion. It seems reasonable to expect that the end result of factorial theorizing will *underdetermine* foundational theorizing even if our factorial inventory is highly constrained. This underdetermination probably won't leave the foundational options wide open. But to expect no underdetermination between at least one consequentialist theory and at least one non-consequentialist theory seems optimistic. If life turns out even a

³This claim needn't beg the question against Berker (forthcoming)'s unity of grounding thesis. It isn't crucial to put it as a point about two kinds of grounding or even two kinds of explanation, which might stand to grounding much like it stands to causation. All I need is a distinction between theories that focus on two different relations, one of which needn't be a grounding relation. We can surely distinguish (i) theories that inquire after whether the property of rightness is analyzable and what its analysis might be and (ii) theories seeking a fundamental answer to the question of why right acts are right, with that question understood normatively. If one must then claim that the analysis relation isn't a grounding relation or that the last question isn't a grounding question, so be it.

⁴Of course, methodological generalists would say that we must do the foundational work first. But one needn't be a Dancy-style particularist to prefer *methodological* particularism, insisting that we should give a foundational story only after doing the factorial work.

little hard, we will find ourselves choosing between a consequentialist theory and a non-consequentialist theory that from a distance may sound consequentialist simply because it agrees at the factorial level with the consequentialist theory.

This fact plays a central role below. The evidence for some forms of reliabilism may look like evidence for epistemic consequentialism, since it is also predicted by epistemic consequentialism. But the best foundational explanation of this evidence and best underpinning for reliabilism is non-consequentialist.

My final assumption should be obvious but is arguably neglected in epistemology. The assumption is that consequentialism is a *normative theory*, in a broad sense of ‘normative’ that includes both theories of rational decision procedures and theories of right-making characteristics. Note that there are all manner of theories in philosophy that ground certain kinds of facts about an X in facts about X’s effects, or in facts about the effects of something else. Since effects are consequences, these theories ground certain kinds of facts in facts about consequences. Many of these theories are obviously not in the same ballpark as consequentialism. Consider, for example:

- *Causal theories of property identity*—e.g., the view that F and G are identical iff F and G confer the same causal powers on all particulars that instantiate them.
- *Causal theories of perception*—e.g., the (admittedly oversimplified) view that a visual experience is a seeing of x iff x causes that visual experience.
- *Causal theories of intentional action*—e.g., the (admittedly oversimplified) view that a bodily movement is an intentional action iff it is caused by an intention.
- *Response-dependent theories of color*—e.g., the view that X is red iff X causes red-experiences in observers in normal conditions.

It might be heuristically useful to analogize these theories with forms of consequentialism. But it would be a mistake to frame this analogy as the discovery that consequentialism is a popular view in the metaphysics of color, property identity, perception, and action. Such framing would be confused partly because color, perception, and action seem non-normative. As a result, these causal theories aren’t in the same business as consequentialist theories. But even if there were some sense in which color and action, say, were normative, it still wouldn’t follow that the explanations given by causal theories of color and action would be normative explanations. For the relation between perception and its causal constitution as well as between action and its causal constitution is not normative, while the relation between an act’s rightness and what makes it right in first-order terms is normative.

These points are important in assessing whether reliabilism is rightly viewed as a form of consequentialism. Reliabilism in its earliest formulations was a theory about the nature of knowledge. Knowledge might just be a highly generic factive mental state. Even if this state were in some sense normative, a causal theory of its constitution analogous to the causal theory of perception would not be a consequentialist theory. For the constitution relation in play needn’t be a normative relation of the kind in play in first-order ethics.

A final comment is in order. Occasionally one hears it said with a nod to Rawls that consequentialism in its most general form is simply the claim that value is prior

to rightness. This is wrong. One can be a value-firster without being a consequentialist. While putting value before rightness is part of being a consequentialist, consequentialism involves a further claim about the way that rightness is grounded in value: it is grounded in value by means of instrumental relations, broadly understood to include constitutive means. In other words, consequentialism in its most general form is the more specific claim that our sole fundamental duty is to *promote* the good, where ‘promote’ means *bring about*. As Nozick (1981: 429–30), Anderson (1993), Scanlon (1998: Ch.2), Raz (2011: Ch.11), Parfit (2011 v.1: Ch. 20), and others observe, there are many other relations one can properly bear to value other than promotion. One can *respect* and *honor* value, for example.

As I see it, the best non-consequentialist theories are not brute deontological views on which we have certain duties that are disconnected from value. They are rather views that preserve a connection between value and rightness but claim that the fundamentally fitting response to certain values is not promotion but rather (e.g.) respect. Kant himself may have held such a view, putting the value of autonomy at the bottom in ethics, and giving a non-instrumental story about how our duties are explained by this value (e.g., because they are ways of respecting it).⁵

3 Reliabilism’s Life without Consequentialism (and Occasional Dalliance with It)

I turn now to a historical argument for (C1). Here in outline are the main points this section will make. Reliabilism lived for more than ten years as a non-normative theory. It was treated in Goldman (1979) as similar in aim to various first-order ethical theories. But Goldman (1979)’s theory could be endorsed factorally or foundationally, and if factorally endorsed it could be embedded within a non-consequentialist foundational view—e.g., virtue epistemology, which was supposed to be analogous to one main alternative to consequentialism in ethics.⁶ While Goldman (1986) later proposed an analogy between his project and the rule consequentialist’s, this was just one moment in the history of reliabilism, and much of Goldman’s later work didn’t build on the analogy. Anyway, as I will suggest, Goldman (1979, 1986)’s views aren’t analogous to indirect consequentialism: the analogy is flawed in ways that cannot be ignored if we are serious about our comparisons with ethics. Finally, other reliabilists have continued to hold a more traditional, non-normative form of reliabilism. So although Goldman’s contributions are preeminent, one shouldn’t forget that reliabilism receives notable treatment in other hands.

I turn to document these points in more detail.

⁵For some interpretations of Kant on which value (in the broad sense including a creature’s ‘status’) is at the bottom of the explanatory chain, see, e.g., Dean (2006), Kamm (1992), and Wood (1999).

⁶Goldman (1993: 274) suggests this framing.

3.1 Pre-1979 Reliabilism: Non-Accidentality and the Constitution of Knowledge

Reliabilism started out as a theory of knowledge, not a theory of justification.⁷ What the early versions of reliabilism in Armstrong (1968, 1971), Unger (1968), Dretske (1971), and Goldman (1975) have in common is an attempt to capture the non-accidentality condition on knowledge. While it is standard after Pritchard (2012) to distinguish anti-luck conditions from ability conditions,⁸ and reliability conditions are arguably best understood as versions of the latter, this distinction is new.⁹ Instead of taking it to show that non-accidentality approaches are not necessarily reliabilist, we could take it as an invitation to distinguish different types of reliabilist theory—say, process reliabilism on the one hand and ‘counterfactual reliabilism’, to use Vogel (2000)’s term, on the other. Indeed, there is plausibly a sense of ‘accident’ for which it is necessarily true that it’s no accident when reliable thinkers’ beliefs turn out true. This fact remains even if environmental luck cases show there to be some other condition on knowledge, negatively connected to another sense of ‘accident’.

If one is a reliabilist of the classic stripe, it would be coherent for one to either withhold belief in epistemic consequentialism or reject it. The epistemic consequentialist would derive her reliability condition on knowledge from three things: the assumption that knowledge implies justified belief, an indirect epistemic consequentialist account of justified belief, and a veritist epistemic axiology. The traditional reliabilist needn’t accept the reliability condition for such reasons. She can be a reliabilist simply because she thinks knowledge is non-accidentally accurate belief, and reliability provides the best unpacking of such non-accidentality.

There is no necessary tie between embracing a view that sees non-accidentality and the main—or even the only—thing one must add to true belief to get knowledge and embracing epistemic consequentialism. Indeed, a close parallel in ethics to early reliabilism’s emphasis on non-accidentality is provided by interpretations of Kant that see *non-accidental conformity to moral law* as the defining feature of moral worth, and the feature missing in examples like the shopkeeper case, which Kant (1785/2012: 4.398) describes as a case of ‘fortunate’ (i.e., accidental) rightness.¹⁰ If one sees reliability as important because it is a way of capturing this more fundamentally significant constraint—a plausible interpretation of the paradigm non-consequentialist—one can easily be a reliabilist without being a consequentialist.

One can also be a reliabilist of the classic stripe without offering reliabilism as a normative theory of any kind. Armstrong (1968: Ch.9)’s first presentation of his reliabilism about knowledge appeared in a book on the metaphysics of mind, where he treated knowledge a mental state and saw perceiving to be a special case

⁷Ramsey (1931b: 193–198) provides an interesting exception in giving a reliabilist account of reasonable belief, and there is evidence that he was an epistemic consequentialist (see esp. p.196). But as Goldman and Beddor (2015) observe, Ramsey’s early reliabilism ‘attracted no attention at the time and apparently did not influence reliability theories of the 1960s, 70s, or 80s.’

⁸Though see Littlejohn (2014).

⁹Goldman and Beddor (2015) describe Goldman (1976)’s response to the fake barns case as ‘counterfactual reliabilism’, however.

¹⁰For some different versions of this interpretation, see Herman (1981), Benson (1987), Baron (1995), and Stratton-Lake (2000).

of knowing (see (1968: Ch.10) and cf. (1961)). On this kind of view, one could hold that knowledge is no more normative than the factive mental states that are its determinates. If so, the condition of non-accidental arrival at the truth will just be a generalization of the condition of non-accidental veridicality on seeing that *p*. It no less a mistake to claim that this kind of view is an example of consequentialism in epistemology than it is to say that the causal theory of properties is a version of consequentialism in metaphysics.

3.2 1979 Reliabilism

It must be acknowledged that reliabilism did in the hands of Goldman (1979) become a normative theory in a broad sense. Goldman was explicit from the first page:

The term ‘justified’, I presume, is an evaluative term, a term of appraisal. Any correct definition or synonym of it would also feature evaluative terms. I assume that such definitions or synonyms might be given, but I am not interested in them. I want a set of substantive conditions that specify when a belief is justified. This might be defined in other ethical terms or phrases, a task appropriate to meta-ethics. The task of normative ethics, by contrast, is to state substantive conditions for the rightness of actions. Normative ethics tries to specify non-ethical conditions that determine when an action is right. [...] Analogously, I want a theory of justified belief to specify in non-epistemic terms when a belief is justified.¹¹

On the same page, Goldman mentions act utilitarianism as an example of a theory that gives ‘substantive conditions’ for rightness. But he doesn’t suggest that his view is analogous to act utilitarianism or any other utilitarianism in more than the general sense clarified in the above passage. We are not getting the overt modelling on consequentialism here that we get in Goldman (1986).

Nonetheless, one might think that the normative theory Goldman (1979)’s view most strongly resembles is a monistic indirect consequentialism—say, rule utilitarianism or, closer yet, motive utilitarianism. For this reason, one might not find it misleading to say that indirect epistemic consequentialism received its first detailed formulation in Goldman (1979). But even this claim lacks plausibility. Goldman (1979)’s theory is in key ways not parallel to indirect consequentialism.

There are at least four features of Goldman (1979)’s account that support this verdict, on which I’ll comment briefly in outline and then at greater length:

- A. The account appeals only to a highly restricted class of reliable processes, which is not what one would expect in a theory analogous to any familiar consequentialism in ethics.¹² If these restrictions have ethical analogues, the analogues are agent-centered restrictions against maximizing the good in certain ways. Furthermore, these restrictions lack a plausible indirect consequentialist rationale.

¹¹Goldman (1979: 1).

¹²Goldman (1979: 13–14).

- B. The account counts only certain highly restricted outputs of processes in the measure of their reliability, which is not what one would expect in a theory analogous to any familiar consequentialism in ethics.¹³
- C. Goldman's account of defeat has no analogue in any familiar consequentialism in ethics. While it is a more purely reliabilist account of defeat than the one found in some of his later work (see, e.g., Goldman (1986) and Goldman (2011)), it is not thereby more consequentialist, but less so.
- D. The asymmetry the account yields in the conditions for justified basic belief vs. justified non-basic belief has no real analogue in consequentialism in ethics.¹⁴

Let me now elaborate on each, taking (A) first. Goldman (1979: 12–3) notes that there are different ways one could understand the *extent* of the relevant reliable belief-forming processes. Belief-forming processes could be regarded as originating outside the organism or within the organism. Goldman observes that our ordinary attributions of justified belief only seem to take into account the reliability of processes typed as starting within the organism, with 'regist[r]ation and transform[ation] [of] the stimulation that reaches him' (13).

It is hard to see how this restriction follows from foundational consequentialist principles. While indirect consequentialists face something like the generality problem, since rules and motives can be described with more or less precision, a restriction of this type has no rule or motive consequentialist precedent. Rule and motive consequentialists directly evaluate rules and motives in terms of their consequences—no holds barred—and then evaluate acts by looking at the rules with which they conform or the motives that cause them. If process reliabilism were analogous to rule or motive consequentialism, it would directly evaluate processes by the consequences of their use, and then indirectly evaluate beliefs by looking at the processes that cause them. But there is no reason to think that processes typed as starting within the organism are more likely to maximize its ratio of true to false beliefs than processes typed as starting outside the organism.

Let's now consider (B). While the reliability of a process is determined by looking at certain properties of its consequences, only certain consequences matter—viz., only the 'immediate' outputs of the process, which Goldman (1979: 13–14) takes into account in his clauses (6A) and (6B). This restriction lacks precedent in the consequentialist tradition in ethics. Direct consequentialists take all future consequences of an act to be relevant, and indirect consequentialists take all consequences of whatever they directly evaluate in consequentialist terms—rules or motives—into account.

Ahlstrom-Vij and Dunn (2014) try to use the presence of the immediacy restriction in reliabilism to answer Berker (2013a)'s argument against epistemic consequentialism, insisting that reliabilism is both a direct and an indirect consequentialist theory. But they beg the question here by presupposing without argument that reliabilism is a form of epistemic consequentialism,¹⁵ and by using 'direct' in a sense not used in the consequentialist literature in ethics. Because they assume that reliabilism

¹³Goldman (1979: 12–13).

¹⁴Goldman (1979: 13–14).

¹⁵To be fair, Berker (2013) himself wrongly paints reliabilists as epistemic consequentialists.

is a form of consequentialism, they assume that the immediacy restriction is available to the consequentialist in virtue of its availability to the reliabilist. But it is *not* available to reliabilists for consequentialist reasons, if one uses ‘consequentialist’ in a sense that parallels its sense in ethics. While Ahlstrom-Vij and Dunn (forthcoming) now acknowledge that it is an open question whether reliabilism is a kind of consequentialism, they again err by assuming that consequentialism is merely the claim that rightness is somehow or other explained by value, a claim insufficient to distinguish consequentialism from paradigmatically non-consequentialist theories.

A third feature of Goldman’s account that is hard to derive from consequentialist foundations is its theory of defeat. Of course, the very idea of defeat is not foreign to consequentialism: an action can be right in one circumstance but wrong in another if it has different consequences in the two. Examined from a distance, adding or subtracting consequences might resemble adding or subtracting information. But the particular way in which reliabilists model defeat is not what one would expect if reliabilism were given consequentialist foundations. Goldman (1979)’s account of defeat is especially interesting to consider, since it is his only purely reliabilist account of defeat. He moved in (1986) to allow beliefs and experiences to be defeaters irrespective of ancestry, and continued in (2011) to give experiences this role.

Goldman (1979)’s account of defeat is that a belief produced by a reliable process has *prima facie* but not *ultima facie* justification if there is an alternative reliable process available to the believer that would have resulted in a different doxastic attitude (disbelief or suspension). There is no direct or indirect consequentialist analogue of this account of defeat. There is no direct consequentialist analogue of this conception of defeat, since direct consequentialism is forward-looking and the general value of the propensities of act-producing processes are swamped in rightness evaluation by the particular consequences of the token act. Perhaps there is a vague analogy between the relevance of alternative processes and the relevance of alternative strategies open to the agent at the time of action. But note that if two strategies that recommend different acts would have equally good consequences, act consequentialism permits either; by contrast, if one reliable process would recommend belief and an equally reliable process would recommend suspension of judgment, process reliabilism doesn’t permit either belief (and rightly so).¹⁶

For a related reason, there is also no indirect consequentialist analogue of Goldman (1979)’s conception of defeat. Consider rule consequentialism first. If two rules that recommend different acts A and A* in some circumstance C are such that their acceptance would have equally good consequences, or if two rules are equally reliable guides to promoting the good, it is permissible to do either A or A* in C. Something similar goes for motive consequentialism: if two motives are equally reliable guides to promoting the good, or are such that acting on either would tend to produce similar amounts of goodness, then acting on either is permissible. But again, if one reliable process would recommend belief and an equally reliable process would recommend suspension of judgment, process reliabilism (rightly) doesn’t permit both.

Curiously, the less purely reliabilist stories Goldman later tells about defeat are more similar to what consequentialist foundations would encourage. Note that the

¹⁶Cf. Littlejohn (2012: 77), who notes that epistemic consequentialism fails to secure this prediction but maligns Goldman because he assumes Goldman is a consequentialist.

consequentialist can help herself to the structural distinction between subjective and objective rightness. She might hold that justified action requires subjective rightness, and hence must have good expected consequences, or be in accordance with a rule whose acceptance would have good expected consequences. Goldman's later treatment of beliefs and experiences as potential defeaters might be read in a similar way: if some beliefs or experiences can influence expected reliability, they can thereby impact epistemic justification. But these proposals about defeat in epistemology are often regarded as *ad hoc*: as a thoroughgoing externalist view, reliabilism only explains why *actual* reliability matters, not perceived or apparent reliability.¹⁷

A final feature of Goldman (1979)'s account that doesn't straightforwardly follow from consequentialist foundations is its (modest) foundationalist structure, and specifically the asymmetry it yields between the justification conditions for basic and non-basic beliefs. Of particular interest is the requirement in Goldman's recursive account that non-basic beliefs not be based on unjustified beliefs on pain of inheriting this unjustifiedness. The analogue of this requirement in the practical case would be a requirement that non-basic actions not be based on wrongful actions. Now consider the practical basing relation of immediate interest to the consequentialist, which is the means-end relation. If consequentialism were strongly parallel to Goldman's reliabilism, we would expect the consequentialist to claim that a non-basic action cannot be right if it is performed via wrongful means.

But consequentialists shouldn't like this claim. While a consequentialist can allow that a killing of an innocent person is *pro tanto* bad and *prima facie* wrong in virtue of realizing an intrinsically bad state of affairs, a killing can become right *ultima facie* if its effects are good enough. While some consequentialists set high thresholds, they will eventually allow tradeoffs. The thinking here is the reverse of what seems sensible in the epistemic case. The all-things-considered deontic evaluation of the basis belief is prior to the evaluation of the based belief, while the all-things-considered deontic evaluation of the means action is posterior to the evaluation of the goal action.

Note again that I am not saying that there won't be heuristically interesting analogies between the foundationalist's structure for justified beliefs and some normative structures that consequentialists don't frown upon. There is, for example, an interesting analogy between regress arguments for foundationalism and regress arguments for intrinsic value. But what we need is not some indirect analogy between a theory of epistemic justification and a theory of value, but rather an analogy between a theory of epistemic justification and a theory of *right action*. For there to be an analogy strong enough to say that the same view appears in both domains, we need more than rough resemblances between unlike normative categories. What we need are parallels between like normative categories—i.e., justified belief and justified action.

Goldman did draw an analogy between reliabilism and rule consequentialism in *Epistemology and Cognition*. But despite this fact, many of the disanalogies to which I've drawn attention remain: Goldman (1986) keeps similar restrictions on the kinds of processes and outputs that matter, and preserves an asymmetry between basic and non-basic beliefs like that in Goldman (1979). Given those remaining disanalogies,

¹⁷See, e.g., Greco (2010: 158)

Goldman's view is more similar in the end to moderate deontology.¹⁸

3.3 Non-Goldmanian '80s Reliabilism

Epistemology and Cognition was not the only memorable event for reliabilism in the '80s. A different reliabilism—indicator reliabilism—also flourished in that decade. While the indicator reliabilism of the time was overshadowed by Goldman's approach, insights from '80s indicator reliabilist theories appear in Goldman's recent views (e.g., Goldman (2011)). As I'll argue shortly, if indicator reliabilist theories are counted as consequentialist, then views in ethics that aren't consequentialist will be wrongly classified as consequentialist.

Besides indicator reliabilism, the '80s also featured reliabilist theories of knowledge that—like Goldman (1967)'s causal theory—denied that justification is necessary for knowledge and are consistent with thinking of knowledge as no more normative than factive mental states like seeing that *p*.¹⁹ So while *Epistemology and Cognition* was a classic moment in the history of reliabilism, it would be wrong to think that reliabilism took a definite turn toward consequentialism in the '80s.

Let's further consider indicator reliabilism, and then briefly consider the clearer case of '80s reliabilism about knowledge. Foreshadowed by some externalist accounts of knowledge in the '70s,²⁰ indicator reliabilism about justification received subtly different formulations in Swain (1981a, 1981b, 1985) and Alston (1985, 1988). Alston's account is more intuitive, so I'll take it as the paradigm. Alston's account is a two-level account that first identifies justified beliefs with beliefs based on adequate reasons, and then identifies adequate reasons with reliable indicators. I do not think that the structure of this theory automatically recommends a consequentialist interpretation. If incorporating reliability in an account of justification in this way were sufficient for making the account consequentialist, then ethical theories that are not consequentialist will be wrongly classified as consequentialist.

I will take an appealing combination of views in ethics to illustrate this point. I find it attractive to combine (a) an account of practical justification in terms of possessed objective reasons with (b) an account of objective reasons for action that analyses them in terms of objective evidence and rightness *a la* Thomson (2008) and Star (2015). This two-level account is structurally like Alston's: (a) is parallel to the first layer of his view, and (b) is parallel to the second layer, with objective evidence being naturally construed as the same thing as a reliable indicator.

It is obviously possible to hold (a) and (b) while being a deontologist. Thomson is a paradigmatic deontologist, and she accepts (b) and could accept (a). And Star

¹⁸While that claim might sound strange, don't forget that deontologists agree that there is a strong reason to promote the good: they just think this reason can be defeated if the only way to promote the good would be by violating restrictions.

¹⁹I have especially in mind Dretske (1981). While Foley (1987) famously defended an internalist theory of rational belief, he also recommended divorcing the theory of knowledge and the theory of justified belief and indicated his sympathies for a reliabilist theory of knowledge.

²⁰See Dretske (1971) and Armstrong (1973)'s 'thermometer view', a name unwittingly hearkening back to an indicator reliabilist theory of meaning by the same name that Price (1953: 185-97) attacked. Sellars (1956: §VII.31) approvingly alludes to Price's attack just before anticipating and attacking the indicator reliabilist account of knowledge in §VIII.

contrasts his view with indirect consequentialism but embraces (a) and (b). Consistency is possible here for many reasons, one of which is that rightness can then be understood in terms of conformity to deontological standard(s). The same goes for epistemology. The truth norm could be understood as a deontological standard of the same sort. Just as (a) and (b) are compatible with moral deontology, as one sees in Thomson, so Alston's view is compatible with a truth first epistemic deontology, which I'll describe further in the next section.

The more fundamental point here is the simple one emphasized earlier that one intuitively attractive theory can be underpinned by different normative foundations. Indicator reliabilism, in particular, can be underpinned by quite different conceptions of how truth governs belief. A non-consequentialist way to think of how truth governs belief is to think that there is an objective norm of correctness for belief—it is correct to believe *p* only if *p*—and that good epistemic reasons for belief are objective evidence that the belief that *p* would be in conformity with this norm.

Of course, one might frame indicator reliabilism in a gerrymandered consequentialist way. One could take the fundamental epistemic norm not to be a deontological norm of correctness but rather a consequentialist principle that says 'Promote a high ratio of true to false beliefs!', and one could then preserve the epistemic analogues of (a) and (b). But this fact merely shows that a plausible non-consequentialist view in epistemology can be extensionally equivalent to a gerrymandered consequentialist view. The case-based evidence for indicator reliabilism doesn't decide between non-consequentialist and gerrymandered consequentialist framings.

One might try to bracket indicator reliabilism on the grounds that it is not a purely reliabilist view, but rather a hybrid with internalist and externalist elements. It was intended as such a hybrid in Alston (1988). But the mere fact that a view has internalist and externalist elements is not a sufficient reason for disqualifying it as a form of reliabilism. Goldman (1986)'s theory is reliabilist if any theory is. But it allows internalist factors to defeat externalist *prima facie* justification. There is admittedly an important distinction between views on which reliability is merely necessary for justification and views on which it is necessary and sufficient for justification. But the former views are full-fledged externalisms if the reliability required is reliability vis-à-vis external world propositions. In any case, the most popular versions of reliabilism about justification don't hold that reliable formation is necessary and sufficient for justified belief across the board. Even Lyons (2009)—perhaps the most 'mad-dog' reliabilist around—rejects this view for non-basic beliefs.

So much for how indicator reliabilism could be given non-consequentialist foundations. I'll conclude with a brief mention of '80s reliabilism about knowledge. Dretske's *Knowledge and the Flow of Information* was at least as important an event for externalist epistemology as Goldman's *Epistemology and Cognition*. Dretske (1981: Ch.4) continued in the footsteps of the early Goldman, denying that knowledge should be analyzed in terms of justification, and instead offering a synthesis of earlier causal accounts of knowledge with insights from the indicator reliabilist tradition. This view is a further reminder that reliabilism can be understood as a non-normative theory, and hence not a consequentialist theory.

3.4 Virtue Reliabilism

Another theory often classified as a version of reliabilism is the reliabilist virtue epistemology of Sosa (1991, 2007, 2015) and Greco (2010); indeed, when this view was first introduced at the end of ‘The Raft and the Pyramid’, the section title used to introduce it was ‘Reliabilism: An Ethics of Moral Virtues and an Epistemology of Intellectual Virtues’. But this view is not a version of epistemic consequentialism.

If this view is intended not only as an alternative to other forms of reliabilism, but as genuinely analogous to virtue ethics (hence earning its name), it is best read as having two levels, one foundational and the other factoral. At the foundational level, the account is the epistemic analogue of virtue ethics. As a result, it is analogous to an alternative to the other major foundational options, which are epistemic consequentialism and deontology. At the factoral level, the account is extensionally equivalent to a form of reliabilism in which person-level dispositions worthy of the title ‘competences’ rather than potentially subpersonal processes are the key justifiers. It is then a first-order matter that the normative properties identified as crucial at the foundational level—viz., virtues—line up with the non-normatively specifiable reliable dispositions which do the work at the factoral level.

If this is the correct interpretation of reliabilist virtue epistemology, it is not a form of epistemic consequentialism. Now, some critics—e.g., Dancy (1995) and Zagzebski (1996)—have insisted that the conception of virtue with which Sosa and Greco work is so thin that the account turns out to be consequentialism in disguise. But this criticism is not persuasive, especially if one examines the most recent articulations of the view. It is natural to see Sosa and Greco as inspired by the Aristotelian thought that the fundamental normative property is attributive goodness, which has performance normativity as a special case. It is natural to see them as holding in addition that as a substantive matter, manifestations of unreflective dispositions make for good performances in this sense. Perhaps this view will be extensionally equivalent to something like motive consequentialism, but the reason for the ‘thinness’ of its conception of virtue is not some underlying consequentialism, but rather that this is just plausible on its face when one reflects on some paradigm cases of knowledge (including perceptual knowledge and animal knowledge).

4 The Many Non-Consequentialist Bases for Reliabilism

So far I’ve given a historical argument for (C1). I want now to abstract away from history and describe five broader ways to give reliabilism non-consequentialist foundations. I’ll then explain in §5 why reliabilists should seek some such foundations.

4.1 Knowledge (Sort of) First

The first strategy makes explicit what seemed implicit in reliabilism’s early history. It proposes that reliability is important in virtue of being the key to the analysis of knowledge. While this strategy puts knowledge axiologically first, it allows that knowledge is analyzable, perhaps adopting a simple analysis along the following lines:

S knows that p =_{df-metaphysics} (i) S has a true belief that p , (ii) this belief was

formed by a reliable process, and (iii) the very accuracy of the belief manifests the reliability of the process.

Although it takes reliability to be axiologically posterior to knowledge, this view counts as a version of reliabilism in virtue of treating knowledge as metaphysically explained by reliability. This combination of metaphysical priority and axiological posteriority is familiar and perfectly coherent. By analogy, consider how dabs of paint in a good painting get value by constituting it, so that even if they seemed worthless in isolation, they derive value from their role in a whole.

This view could be defended in a non-normative or a normative form. A non-normative version would view knowledge as simply a determinable factive mental state, itself no more normative than more determinate factive mental states like seeing that *p* and remembering that *p*. The reliabilist analysis offered by this kind of view would be in the same ballpark as (e.g.) causal theories of intentional action. This view is consistent, however, with taking knowledge to be something fundamentally epistemically good, from which reliability derives value as a key constituent. After all, some candidates for intrinsic value are not normatively constituted. Consider pleasure. Just as whatever constitutes pleasure will be valuable in virtue of constituting something intrinsically valuable, so will whatever constitutes knowledge be valuable in virtue of constituting something fundamentally epistemically valuable.

If one wants to explain the value of justified belief, one could add another idea familiar from knowledge-first epistemology: namely, the idea that the fundamental epistemic norm is the knowledge norm. This norm can itself be understood in a deontology-friendly way, as the following correctness norm (where ‘correct’ means *objectively permissible*):

KNC: It is correct to believe that *p* if and only if you know that *p*.

Justification could then be viewed not as a means to true belief but rather as a status had when one is in apparent conformity from one’s own perspective with (KNC). Alternatively, if one is willing to deny that there are justified false beliefs, one could view justified belief as coinciding for first-order normative reasons with knowledgeable belief. This isn’t to say that justification is a constituent of knowledge, but rather that knowledge is itself a justification-making characteristic analogous in role to the right-making characteristics one finds in normative ethics.

4.2 Truth Deontologically First

A second option now becomes apparent. One could put truth rather than knowledge first deontologically, and derive the significance of reliability from the fact that it is the key constituent of *compliance* rather than mere conformity to the truth norm:

TNC: It is correct to believe that *p* if and only if *p* is the case.

One wants not merely to conform to norms, but to comply with them. The rationale might be the non-consequentialist one that the law calls for one to act *from* it, not merely *in accordance with* it. This is a stronger relation than conformity, which can

be accidental. As I noted earlier, this rationale figures centrally in Herman (1981) and Baron (1995)'s interpretations of Kant on moral worth, and isn't foreign to the non-consequentialist tradition. Provided one offers a psychologically undemanding account of acting 'from' a norm—a long-standing goal of the Kantian literature on moral worth—one can derive accounts of justified belief and knowledge that are extensionally equivalent to familiar reliabilist accounts, but with non-consequentialist foundations. Knowledge could be viewed as consisting in conformity with the truth norm as a manifestation of respect for the truth norm, where such respect could be a psychologically undemanding state that Huck Finn has for Jim's humanity. This account would make doubly reliabilist predictions, firstly in virtue of the way it understands respect, and secondly in virtue of the fact that the manifestation relation required for knowledge implies non-accidentality.

4.3 Virtue First

A third option is to take the analogy with virtue ethics as an alternative to consequentialism seriously while retaining externalist themes, defending a kind of synthesis of responsibilism and Aristotelian skill-based virtue theory. Zagzebski (1996) came close to offering such a theory. While some aspects of her view are worth dropping, an attractive view can be constructed on the foundations to which she drew attention.

Zagzebski (1996: §§4.1.2-4.1.3) distanced herself from an earlier internalist virtue responsibilist—viz., Montmarquet (1986)—by arguing that virtue has a 'success component'. The rationale is straightforward: virtues are *excellences*, and one cannot have an excellence to ϕ unless one tends to succeed in ϕ -ing when in suitable environments and in suitable shape. Not just any kind of reliable belief-forming disposition merits the title of 'excellence', however. Accordingly, Zagzebski also imposes a motivational requirement on knowledge. Although it is familiar to object to this account on the grounds that it overintellectualizes—well, overpsychologizes—justified belief and knowledge,²¹ Zagzebski observed that the motivational constraint could be understood in a sufficiently undemanding way that the account could turn out extensionally equivalent to sophisticated reliabilist views:

[M]y definition of knowledge in terms of acts of intellectual virtue can be interpreted in a more externalist way than I have intended, more like Plantinga. This could be done by modifying the motivational element in my account of virtue, making it weaker and farther removed from conscious awareness and control, although I do not think the internalist aspect can be eliminated entirely. The resulting notion of virtue might look a lot like proper function. [...] I mention this, not because I think it is a good idea to bend it [in this way], but because I want to indicate that what I propose here can be adapted by those whose intuitions about knowledge are... more strongly externalist... than mine.²²

Precisely such a bending of the view was, in effect, wrought by Greco (1999)'s agent reliabilism and Sosa (2015)'s updated reliabilist virtue epistemology, which were in-

²¹See, e.g., Baehr (2011).

²²Zagzebski (1996: 329-330).

tended to reconcile responsibilist intuitions with externalist modesty about the psychology of knowing. Provided that one construes the biconditionals offered by agent reliabilists factorally rather than foundationally, one can then defend these biconditionals on the basis of Zagzebski's foundations together with psychological modesty about appropriate motivation. The advantage that this modified Zagzebskian approach would have over Greco/Sosa-style virtue epistemology would then be fine-grained: given the non-consequentialist foundations for its extensional predictions, it has a better solution to reliabilism's value problem.

4.4 Objective Reasons First

The penultimate approach I'll mention is an updated version of indicator reliabilism that profits from insights in recent work on reasons and rationality outside epistemology.²³ In ethics, it has become common to view the fundamental notion of a reason as that of a fact, often worldly, that counts in favor of some action or attitude. On one attractive view, justification of the sort that contrasts with excuse is then taken to be a function of the reasons so understood that one possesses.

The relation of possession can, as I've suggested elsewhere,²⁴ be analyzed in terms of a non-normative access condition together with a normative condition of competence with objective reasons. Since competence with objective reasons plausibly coincides extensionally with sensitivity to reliable indicators—understood to include worldly indicators, not just the mental states of classic indicator reliabilism—and access can be understood in a non-normative reliabilist fashion, this account will be extensionally equivalent to a sophisticated indicator reliabilism.

But the foundations of this account needn't be consequentialist. On the one hand, one could take objective reasons as primitive and see the truth that objective epistemic reasons for belief are truth-indicators as a synthetic *a priori* truth rather than a clue to their analysis. On the other hand, one could analyze objective reasons in any number of patently non-consequentialist ways that are consistent with the account's remaining an indicator reliabilism. For example, one could hold that objective reasons to ϕ in a domain are probability-raisers of the proposition that ϕ -ing would be correct (in a normative sense) in the domain. And one could then understand the norms of correctness in a domain deontologically.

While this account appeals to the probability that ϕ -ing is correct, it doesn't follow that it is a consequentialist account. Every complete normative theory must appeal to some such probability somewhere, since every normative theorist should recognize a distinction between objective correctness in ϕ -ing and ϕ -ing that is justified, praiseworthy, etc., where there is a probabilistic relation between the two. For example, Kantians can agree that there is a distinction between objective rightness and justification, viewing objective rightness as action in compliance with the Categorical Imperative and justified action as action that is likely relative to the indicators within one's perspective to be in compliance with the Categorical Imperative. The core distinction is a structural one neutral on substantive normative foundations.

²³This approach has precedent: McDowell (1994) and Fricker (2007) developed such an account of the epistemology of testimony, and generalizing their accounts suggests the view described below.

²⁴See Sylvan (2016) and Sylvan (Forthcoming-a).

4.5 Epistemic Kantianism

I turn to describe a final view, which is my preferred option.²⁵ I call this view Epistemic Kantianism, since it is built on an analogue of the Kantian thought that value calls fundamentally for *respect*, and only derivatively (if at all) for promotion. According to Epistemic Kantianism, the fundamental normative explanation of why justified beliefs are justified is that they manifest respect for fundamental epistemic value. Since I find it independently plausible that the sole fundamental epistemic value is accuracy, I will be considering the Veritist version of this view. The core of this view is hence:

Veritist Kantianism about Justified Belief: The fundamental normative explanation of why justified beliefs are justified is that they manifest respect for the fundamental epistemic value, which is accuracy.

Reliability enters the picture in virtue of its role in a factorial account of what it takes to manifest respect for accuracy:

Respect—Factoral Theory: One manifests respect for accuracy in holding a doxastic attitude *D* vis-a-vis *p* iff one's *D*-ing is sensitive to the objective evidence that one possesses bearing on whether *p*.

Objective evidence can in turn be given an indicator reliabilist analysis. If we take objective evidence to consist in facts that bear an objective probabilistic relation to the propositions on which this evidence bears, we can then understand the *possession* of such evidence in the doubly reliabilist way mentioned in the last section. Possessing a fact as objective evidence for *p* requires having access to this fact, which could be understood as one's hosting some factive mental state (e.g., seeing that *p*) itself to be analyzed in reliabilist terms, and also being sensitive to the evidence-for relation between this fact and one's belief, which could be understood in terms of a reliable disposition to form the attitudes that are objectively supported by such evidence.

Although reliability is central to the factorial account of respect for accuracy, the overall view does not collapse into a consequentialist view. This is because the more fundamental reason why these various kinds of reliability are important is that they are required for manifesting respect for accuracy. Respect remains the central notion in the fundamental normative story, despite the fact that the biconditional for justified belief entailed by this view is reminiscent of syntheses of evidentialism and reliabilism by Comesaña (2010) and Goldman (2011). Indeed, Epistemic Kantian foundations provide a deeper rationale for the otherwise *ad hoc*-seeming synthesis: both evidence—understood in an indicator reliabilist way—and process reliability—understood as partly constitutive of the access condition on reason possession—are crucial because both fall out of the normatively more fundamental requirement of respect for accuracy.

The idea of deriving a broadly reliabilist requirement from a more fundamental non-consequentialist requirement is familiar in the Kantian tradition. Again, non-accidentality is central to Kant's account of moral worth, but matters derivatively in

²⁵This view is defended at length in my (Ms).

virtue of its role as the factorial underpinning of worthy action. The broader idea that respecting value might itself be understood in terms of reliable sensitivity to objective reasons given by that value is familiar from Arpaly (2003), whose example of Huck Finn illustrates how genuine care or respect for a person needn't be understood as a matter of treating that person in the way one believes to be right, but rather as a matter of treating that person in a way that manifests sensitivity to the objective reasons given by that person's nature. Respect remains the normatively fundamental ideal, though what it takes to manifest it is something friendly to reliabilism.

4.6 Do These Accounts Really Count as Reliabilist?

We discover an embarrassment of riches if we seek non-consequentialist foundations for accounts that extensionally coincide with various reliabilisms. But one might still ask whether these accounts are reliabilist, perhaps with the hope of resisting my proof by illustration that reliabilism can have non-consequentialist foundations. While the raw materials for answering this question were provided earlier, let's walk through why these accounts merit the title of 'reliabilism'.

A broader point is in order first. Several of the foregoing accounts were framed as 'X-first' accounts, where $X \neq$ reliability. It doesn't follow from this mere fact that these accounts are not reliabilist. Reliabilists needn't think that reliability is absolutely first. In the weak, non-Williamsonian sense in which the first proposal was 'knowledge-first', the original versions of reliabilism were knowledge-first. Similarly, in the sense in which the third proposal was 'virtue-first', Greco's agent reliabilism and Sosa's new version of virtue reliabilism are virtue-first. And none of the alleged examples of consequentialist reliabilism are reliability-first views either. On these views, *the value of accuracy* comes first. It is only when we consider how epistemic norms derive from the value of truth that we encounter reliability.

Let's now wade more carefully through the reasons why the foregoing accounts are genuine reliabilisms. As I noted at the outset, reliabilism can be offered as a non-normative theory or a normative theory, and if it is offered as normative theory, it can be offered factorally or foundationally. We can take each possibility in turn.

A non-normative version of reliabilism would propose that a certain kind of reliable ancestry is the key element in the constitution of knowledge, itself understood as no more normative than perception. If this view deserves the name 'reliabilism'—as the history suggests it does—then the first proposal deserves the name 'reliabilism'.

If a factorial theory can deserve the name 'reliabilism', then the second through fifth accounts all straightforwardly deserve the name of 'reliabilism' as well. Factoral theories are consistent with multiple foundations. In virtue of predicting that the same factors matter to justification as uncontroversially reliabilist theories, the second through fifth accounts count as factorally reliabilist.

It is only if reliabilism is understood exclusively as a foundational normative theory that the foregoing accounts might not merit the name. But this exclusive criterion prevents some *classic reliabilist views* from counting as genuine reliabilisms. For some classic reliabilist views weren't intended as foundational normative theories or emphasized different normative notions (e.g., virtue) at the foundational normative level. And again, while Goldman (1986) did offer reliabilism as a foundational nor-

mative theory akin to consequentialism, his view contained restrictions that aren't well-motivated from a consequentialist perspective. A generalization of this last point indicates why reliabilists ought to be non-consequentialists, as we will see.

5 Why Reliabilists Ought to Be Non-Consequentialists

5.1 The Argument from Elimination

Here in outline is an argument that reliabilists shouldn't be consequentialists:

The Argument from Elimination

1. Reliabilism is a kind of consequentialism only if it is a foundational normative theory.
2. But reliabilism should not be defended as a foundational normative theory: if it is defensible, it is only as a factoral theory or a non-normative theory.
3. So, reliabilism shouldn't be defended as a kind of epistemic consequentialism.
4. Moreover, epistemic consequentialism is inconsistent with the most plausible factoral and non-normative versions of reliabilism.
5. So, not only should reliabilism not be offered as a version of consequentialism, it should be offered in opposition to epistemic consequentialism.

I already defended (1) in §2 when I explained why consequentialism should be classified as a foundational normative theory. But here is a reminder of why (1) is true. Reliabilism can be offered in three formats. The factoral and non-normative formats are consistent with epistemic non-consequentialism. To be sure, to say that a version of reliabilism is consistent with non-consequentialism isn't to say that it is inconsistent with consequentialism in its most plausible formulations, as (4) claims. But if a theory *T* is consistent with both the affirmation and negation of another theory *T**, *T* cannot be a *version* of *T**. Hence, (1) is true.

(2) and (4) need defense. I previewed the defense of (4) in §3 when I noted that the restrictions that careful reliabilists have wanted to impose are not restrictions any honest-to-goodness epistemic consequentialism—i.e., one analogous to consequentialism in ethics—would contain. I will deepen this argument by abstracting away from the details of the particular reliabilist accounts I discussed there, arguing that this point holds for any plausible version of reliabilism.

I have not previewed the defense of (2), which is an importantly different claim from (4). (2) doesn't claim that reliabilism is implausible when framed as a form of epistemic consequentialism. It rather claims that reliabilism—setting the question of whether it is a kind of consequentialism aside—is implausible as a foundational theory. But the case for this claim is not unrelated to the case for (2): a foundational theory that puts reliability first would have the same difficulties motivating restrictions needed for factoral plausibility as epistemic consequentialism.

5.2 In Defense of (2)

Let's consider what a foundational first-order reliabilism would look like. Like consequentialism, this view could take direct and indirect forms. A direct first-order foundational reliabilism about a certain epistemic status S (e.g., justification) of a certain potential target of evaluation T (e.g., a doxastic attitude) in a world w would make the following schematic claim:

Schematic Direct Foundational Reliabilism: The fundamental normative explanation of whether target T has status S in w is that T is reliable in w .

Where the relevant targets of evaluation are doxastic attitudes, no prominent reliabilists have favored such a view. This view doesn't claim that it is the reliability of the belief-forming process or method that explains why the beliefs are justified, but rather the reliability of the *beliefs* formed by this process (whatever that might mean), independently of the method by which they are formed.

Perhaps the closest relative of such a view would be a simple counterfactual reliabilism, since in its simple form counterfactual reliabilism does ask us to consider the behavior of the belief in various counterfactual scenarios. But simple views of this kind face obvious counterexamples, like Nozick's (1981)'s grandmother case and Goldman (1976, 1983)'s dachshund case, as well as in-principle problems like closure failure of the most elementary sort (see Kripke (2011: Ch.7)). The former problems are why Nozick relativized to methods, and the latter why Roush (2005) has abandoned the Nozickian account for non-basic knowledge and offered a recursive account which is Nozickian only at the foundational level.

It is no surprise that reliabilists about justified belief and knowledge do not appeal to the reliability of beliefs but rather the reliability of the processes, methods, dispositions, etc., which generate beliefs. So if they were to commit to foundational reliabilism, it would have to take the following different form:

Schematic Indirect Foundational Reliabilism: The fundamental normative explanation of whether a target T has status S in w is that T^* is reliable in w , where T stands in an indirect relation R to T^* (e.g., T is produced by T^*).

While there are extensional concerns with this theory, the deeper concerns center around its claim of fundamental normative explanation.

My core concern is hardly without precedent, though it hasn't been put in quite the way I'll put it. As others have noted,²⁶ the way that some prominent reliabilists address objections like Bonjour (1980)'s clairvoyance objection and defeat objections seem ill-motivated if reliability is what matters fundamentally.²⁷ Some of these writers accordingly suggest that we ditch reliabilism for a different kind of theory. But such ditching is advisable only if reliabilism is a foundational theory, something

²⁶See, e.g., Greco (2010: 158) and Beddor (2015).

²⁷Bergmann (1997) nicely documents this strand in the literature. But it has seen further iterations. Goldman (2011), for example, allows that mere experiences are defeaters even when they are misleading (and indeed, at least conceivably, not reliable indicators at all), an allowance that parallels his (1986) allowance of mere beliefs as defeaters (which is also allowed by Plantinga (1993) and Nozick (1981)).

which, as we've seen, it needn't be. The refined theories prominent reliabilists have offered are defensible as factoral theories. Since externalists have tended to put a premium on commonsense intuition, often contenting themselves with a description of the epistemic concepts revealed in our ordinary practice of attribution,²⁸ it is hard to see why they wouldn't be willing to rest content with limiting their ambitions to the factoral level. So, it is challenging to be a foundational reliabilist if one makes consistency with ordinary intuition a desideratum, though factoral reliabilism of the sort defended by central reliabilists is not so challenged.

Let's consider the concern in a bit more detail. Two most familiar problems for simple forms of reliabilism about justified belief are the clairvoyance and new evil demon problems. The former can be addressed extensionally by requiring not just real reliability but apparent reliability from one's perspective. The latter can be addressed extensionally in two non-revisionary ways: by indexing reliability to the actual world, or by allowing that apparent reliability from one's perspective suffices for some kind of normative status, though perhaps weaker than full-fledged justification (cf. Goldman 1988). If we prefer foundational reliabilism, it is unclear why one would favor either of these moves. If reliability provides the most fundamental explanation of why a belief has normative standing, apparent reliability ought to make for apparent justification, not any real kind of justification. Similarly, the fact that a process is unreliable in the believer's world should prevent it from being justified in that world, whatever might be said of the process in other worlds. Finally, if a process is reliable in the believer's world, it is unclear why it also must be apparently reliable from the believer's perspective to qualify for full-fledged normative status: everything needed for a foundational explanation of its normative status is already in place. One could revise the story, holding that the complete explanation of how a belief acquires its normative status must advert to apparent reliability. But that story is implausible as a foundational story. If apparent reliability has significance, its significance is derivative, not foundational.

5.3 In Defense of (4)

The argument for (4) turns on similar points, though the problems for giving reliabilism consequentialist foundations are more pronounced than the problems for a foundational reliabilism that doesn't claim the 'consequentialist' moniker. For while reliabilism is the property of epistemologists and can perhaps be framed however its leading defenders fancy, consequentialism is not. Given consequentialism's roots in ethics, an epistemological theory worthy of the name 'consequentialism' should parallel theories by that name in ethics. Otherwise 'consequentialism' threatens to become a mere homonym in ethics and epistemology; even if epistemologists aren't bugged by this threat, it is unclear why we should bother introducing a misleading new label for a theory that already has a name. But the stronger one makes the anal-

²⁸The Moorean tendency among many externalists to reject out of hand theories that imply skepticism is an important illustration. While Goldman (1993) envisaged two projects—descriptive epistemology, which seeks to understand our 'epistemic folkways', and normative epistemology, which seeks to refine them—his later work on methodology (e.g., Goldman and Pust (1998) and Goldman (2007)) suggests that he takes philosophy's proper task to be (more or less) a description of our ordinary concepts as revealed by our practices of attribution.

ogy between consequentialism in ethics and consequentialism in epistemology, the farther away ones move from a theory reliabilists will gladly endorse.

Some of the problems here have already been well-documented in the literature attacking epistemic consequentialism.²⁹ I'm not sure we should be persuaded by Selim Berker's much-discussed arguments if understood as arguments against familiar views in epistemology. But Berker does persuasively show that theories of epistemic justification that closely parallel consequentialism in ethics will at best be highly revisionary, revisionary enough to fail the desideratum of continuity with ordinary attribution that Goldman (1993) imposes on 'scientific' epistemology. This combination of thoughts recommends my conclusion. While epistemic consequentialism is probably false, reliabilism remains unscathed because it never plausibly was consequentialist in the first place. As Goldman once said,³⁰ while he did affirm consequentialism at one point, he never affirmed it *loudly*.

Even if reliabilism had been loudly promulgated as a consequentialist view in some formulations, it shouldn't have been: a firm analogy with consequentialism in ethics won't secure the predictions reliabilists have wanted or motivate the restrictions they have imposed. Ahlstrom and Dunn (2014)'s response to Berker provides a nice illustration of this point, as I noted earlier. They are right that Goldman's process reliabilism isn't impugned by Berker's arguments. But again, these restrictions have no parallel in consequentialist theory in ethics, and look more like the agent-centered restrictions of deontology. No principled consequentialist would hold the disvalue of violating a restriction to be so great that the value of a certain set of consequences could never outweigh it. Yet only such a claim makes consequentialist sense of Goldman's restrictions (which, while plausible, is not sustainable on genuinely consequentialist foundations).

But we needn't invoke Berker-style arguments to defend (4). We need only consider where epistemic consequentialist foundations will lead the reliabilist *vis-à-vis* clairvoyants, demon worlders, and the value of justified and knowledgeable belief. While veritist epistemic consequentialists might exploit the seemingly theory-neutral distinction between subjective and objective rightness to address some intuitions about clairvoyance cases and the demon world, there are others they will be ill-equipped to address. At best they will be able to claim that the clairvoyant believes in a way that is subjectively wrong though objectively right, and that the demon worlder believes in a way that is subjectively right through objectively wrong. But they will not have the resources to explain why we think the demon-worlder is *genuinely better off from an epistemic point of view* in virtue of believing rationally.

While subjective rightness may be a kind of rightness, expected value is not a kind of value. To a veritist epistemic consequentialist, the demon-worlder's rational worldly belief-forming processes should seem bad all-things-considered, since they only reliably produce false beliefs. A related claim holds *vis-à-vis* the clairvoyant: to a veritist epistemic consequentialist, there shouldn't seem to be anything bad about the clairvoyant's belief-forming dispositions, since they reliably produce true beliefs. Even if the veritist epistemic consequentialist could find something damning to say about the clairvoyant's belief-forming processes and something good to say about the

²⁹See Berker (2013a-b), Fumerton (2001) and Firth (1981).

³⁰In person in 2010.

demon-worlder's belief-forming processes (in her world), the core of the swamping problem would remain vis-à-vis our intuitions about the value of the *products* of these processes. For as Jones (1997) noted in giving the swamping problem one of its earliest formulations, the problem is a general problem for epistemic instrumentalism and other theories founded on its basis, not a problem for reliabilism as such. Indeed, if reliabilism is offered factorally and given virtue-theoretic or deontological foundations, it won't face the problem at all.

Now, there is a familiar move that consequentialists in ethics make to address intuitive deficiencies of this kind: they expand their theory of intrinsic value, moving from a monistic axiology like hedonism to a pluralist axiology like Moore's. But the analogue of this move in epistemology threatens to undermine the reliabilist's explanatory ambitions vis-à-vis the nature of justification, rationality, and knowledge. It is hence not a move that a reliabilist would be well-advised to make. Suppose, for example, we expand our account of intrinsic epistemic value to include rational belief, justified belief, and knowledge. If we do so and still seek to derive reliabilism from epistemic consequentialism, we will end up with a correspondingly expanded conception of reliability. Reliability won't simply be truth-conducivity but conductivity to a plurality of intrinsic epistemic values. If we are serious about deriving justification from consequentialist foundations, our account of justified belief will have to be correspondingly expanded. A justified belief by pluralist epistemic consequentialist lights will be one that is formed by a process that is reliable in the expanded sense. The result is that we are left with a circular account of the properties reliabilism promised to analyze.

I conclude, then, that if reliabilists want to avoid implausibility, they will need to seek non-consequentialist foundations. Hence, not only should reliabilism not be offered as a version of epistemic consequentialism, but it should also be founded positively on a kind of non-consequentialism.

6 A Non-Consequentialist Reliabilism

I turn to consider in greater detail what such a non-consequentialist reliabilism might look like. The aim is not to provide conclusive reasons for preferring this view to all competing epistemologies. The aim is rather (a) to give a detailed illustration of a view built on non-consequentialist foundations that deserves the name 'reliabilism', and (b) to suggest that there is at least as much reason to accept it as there is to accept a consequentialist reliabilism.

6.1 The Structure of the View: Epistemic Metaphysics and Normativity

The view has two parts, one of which is itself further divided. On the one hand, it offers a metaphysics of epistemology. This metaphysics is non-consequentialist in the trivial sense that it isn't a normative theory, and *a fortiori* isn't a consequentialist theory. And it is reliabilist in two ways. Firstly, it gives as central a place to reliability in its metaphysical analyses, as some early reliabilists did. Secondly, together with some innocuous normative assumptions, the view entails factoral reliabilism.

This metaphysics of epistemology consists in a bi-level picture I've recommended elsewhere.³¹ On this picture, epistemology divides into a non-normative and a normative tier, where the former provides a foundation for the latter. Within the latter tier, all normativity is analyzed in terms of *objective reasons*, which are facts that constitute objective evidence for the correctness of attitudes and acts. Not all objective reasons are relevant to the analysis of the most familiar forms of epistemic normativity, however. Only ones to which we have *access* are relevant. Access I take to be a factive mental relation analyzable in the spirit of the non-normative analyses of knowledge favored by Dretske and the early Goldman. I take the accessing of a fact and the knowing of it to be the same thing, and take knowledge to be a non-normative relation partly for this reason. In this way, knowledge amounts to a precondition for standing in the space of reasons rather than a standing in the space of reasons like justified and rational belief.

Now let me describe the reliabilist analyses that fall out of this picture, taking the higher—or normative tier—first. Within the normative tier, the account's commitments are meant to integrate with a larger reasons-first approach familiar in metaethics that analyses all normativity in terms of reasons. This tier contains the following analyses of *ex ante* and *ex post* epistemic justification:

(R_{ea}) There is *ex ante* justification for S to believe that p =_{df-metaphysics}

- (i) there are objective normative reasons to believe that p,
- (ii) S possesses these reasons, and
- (iii) there are no stronger objective normative reasons to disbelieve p or suspend judgment on p to which S stands in the relation of possession.

(R_{ep}) S has an *ex post* justified belief that p =_{df-metaphysics}

- (i) conditions (i-iii) in (R_{ea}) obtain, and
- (ii) S believes that p in compliance with the objective reasons she possesses.

These analyses look a bit like evidentialist analyses. But a kind of factoral reliabilism falls out of them in virtue of their invocation of

1. the relation of *being an objective reason* to believe p
2. the relation of *being in compliance with an objective reason* to believe p

A natural understanding of (1) is indicator reliabilist in spirit. As a substantive normative matter, it is plausible that the factor that makes a fact F into an objective reason to believe p is the obtaining of an indication relation between F-type facts and the truth of p-type propositions. Because it is factoral, the claim that indication relations make for objective reason-relations is like the claim that pleasure makes for well-being. How we explain this factoral truth is a further question. But as I'll note in the next subsection, we needn't explain it in a consequentialist way.

³¹See Sylvan (forthcoming-a).

A natural understanding of (2) is an ability-based reliabilist one. It is again plausible that the factor that makes—as a substantive normative matter—S count as being in compliance with an objective reason to believe p is S’s manifesting her attunement with this indication relation, where S is attuned with an indication relation iff S is disposed to believe what is indicated to be true, provided that there are no equally strong or stronger indicators to the contrary.

There can be excellent indicators to which we are not attuned, given their informational or rational limitations. Moreover, we can be attuned to certain kinds of indicators and these indicators can favor believing p even though we do not manifest our attunement by responding with belief to these indicators. *Compliance* with a reason requires not just conformity in doxastic attitude with the indications, but also that one’s conformity manifest one’s attunement to these indications.

In addition to having a factoral role with respect to compliance, attunement has an analytical role to play with respect to possession. For possessing an objective reason to believe p given by some fact q is not simply a matter of having access to the fact that q. For example, it may be that q is only an objective reason to believe p because there is an extremely complicated proof from q to p, one beyond the logical acumen of most people. Unless one is also attuned to this proof—itself a kind of conclusive indication relation—one will not possess the specific reason to believe p that is provided by q. Hence, it seems that an adequate analysis of possessing a normative reason will not only require access but also attunement. Compliance then goes beyond possession because one can, as noted earlier, be attuned to a general kind of reason-relation without manifesting this attunement by forming the relevant doxastic attitude (e.g., even if one has the logical acumen to make out the proof, one might simply not go ahead and do the proof and form the belief that p).

To possess an objective reason is also a partly non-normative matter of having access to the fact that gives the reason. Here we encounter another place where the account has factorally and non-normatively reliabilist implications. It gives a non-normative, reliabilist analysis of access in the spirit of early Goldman and Dretske, as follows:

S accesses the fact that p iff

- (a) some cognitive system of S’s (e.g., her perceptual system) veridically represents the fact that p,
- (b) this representation manifests the reliability of that cognitive system, and
- (c) the sheer veridicality of this veridical representation manifests the reliability of that cognitive system.

Having access to a fact is then a matter of being in a position to access the fact. The notion of ‘being in a position’ to access something is of course vague. But so is the notion of possession, so there is no problem. This reliabilist account of access is also offered as an analysis. But it is a non-normative analysis. Hence, it is not strictly speaking part of the reasons-first theory—not as such, at any rate. So this theory is trivially non-consequentialist for the same reason why, e.g., the causal theory of properties is not a consequentialist theory.

Now, the metaphysics of epistemology just described is extensionally like the evidentialist reliabilism of Comesaña (2010) and Goldman (2011). The main difference is that it is more fine-grained, giving it advantages I'll describe shortly. Since those views are forms of reliabilism at the factoral level, the account just described is also a reliabilism at the factoral level. Nonetheless, this view is compatible with any number of foundational normative theories. There is an embarrassment of riches, and it would take another paper to decide on any specific foundational normative underpinning for this factoral reliabilism.

Here I will just briefly rehearse my preferred foundations, which I previewed in §4.5. These foundations are provided by Epistemic Kantianism, which I would develop in a specifically Veritist form. Epistemic Kantianism holds that the fundamental normative explanation of why justified beliefs are justified is that they manifest respect for fundamental epistemic value, which is accuracy. What does it take to manifest respect for accuracy? At this point, the metaphysical story I've just been telling becomes subordinate to Epistemic Kantianism, in virtue of helping to provide a factoral story about respect for accuracy. Respect for accuracy requires, as a factoral matter, complying with objective epistemic reasons, which are truth-indicators. It is for this reason that justified belief itself requires such compliance. R_{ea} and R_{ep} hence fall out of Veritist Epistemic Kantianism when combined with a plausible factoral account of what it takes to manifest respect for accuracy.

On this overall view, the *reason why* reliability matters in epistemology is that it is thoroughly bound up with objective reasons and their possession, which in turn play a factoral role in understanding respect for accuracy. To be sure, if one is mainly interested in factoral explanation or in the non-normative grounds of normative properties, one might see things a bit differently: one might say that on the contrary, reliability is what makes up both sensitivity to reasons and respect. But we must distinguish different uses of 'make' and different styles of explanation. Both angles can be legitimately adopted. Choosiness is justified only if we have specified the kind of explanation being sought. It's just that the analogy with ethics, if intended seriously, makes some decisions for us: consequentialism offers a distinctive kind of explanation, and in epistemology that explanation is wrong-headed.

6.2 ≥ Epistemic Consequentialism

I turn to explain why there is at least as much reason to accept this kind of view as there is to accept any epistemic consequentialist view. I'll divide the evaluation into case-based reasons for theory choice and more fundamental explanatory reasons.

6.2.1 Case-Based Reasons

The only strong case-based argument I can see for epistemic consequentialism would involve exploiting the mistaken assumption that reliabilism is a version of it and then thieving reliabilism's case-based advantages. But our view has all the advantages of the familiar reliabilist views and more, so it is at least as good as epistemic consequentialism on this front.

Let's consider some of these advantages. As we've seen, the account is triply externalist, and has the many of the attractions of familiar externalist views. Recall

the three sources of externalism: the view is externalist in using at its core the notion of an objective reason, in understanding the notion of possessing such a reason in part in broadly reliabilist terms, and in understanding the notion of compliance with reasons also in part in broadly reliabilist terms. The first dose of externalism allows us to follow reliabilists in avoiding the conflation of justification and excuse. The second and third doses allow us to follow reliabilists in avoiding bad internalist predictions about cases, including the speckled hen, cases of checkered etiology (like Siegel (2012)'s cases of wishful seeing and Greco (2005)'s cases of incompetent storage), and cases of incompetent basing.

The speckled hen can be understood as follows. True, since experience reliably indicates reality, the fact that one has an experience as of P can constitute an objective reason to believe that P. But having access to an objective reason and being sensitive to the indication relation in virtue of which it is such a reason here come apart. Upon seeing a speckled hen, one is not automatically sensitive to the indication relation between the fact that one is having an experience of the sort one is having and all the propositions represented by the experience. So, one does not possess all the objective reasons that exist in virtue of one's having the experience. Admittedly, one might say: 'Look, you *have* the experience, and the fact that you're having the experience is an objective reason to believe those precise propositions. So, you *have* that reason.' While this speech has some pull, it owes to equivocation. 'Have' means two different things in the speech.

In cases of checkered etiology, two things might be happening. On the one hand, one might simply lack an objectively good reason for belief. Consider Greco's example of Sophie, who came to believe that Dean Martin is Italian by being told so by a source she knew to be unreliable. Sophie has now forgotten the source but continues to believe that Dean Martin is Italian, because she seems to remember that this is so. Sophie's original reason was bad. The fact that her current seeming has this source undercuts its status as an objectively good reason. While the internalist is right that reasons have to be possessed to justify, reasons also have to be objectively good to justify. And the factors that make a reason good or bad needn't themselves be accessible: they are enabling conditions, not further reasons. Similar claims could be made about Siegel's wishful seeing examples. But more can be said. If it seems to one that P partly because one desires that P, this seeming cannot *manifest* a well-functioning capacity of a perceptual system.

Cases of incompetent basing can be addressed straightforwardly by our account, given that it understands ex post justification in terms of compliance, not mere conformity. Consider the following case from Turri (2010):

(SPURS WIN) Ponens and F. A. Lacy know that the Spurs will win if they play the Pistons, and know that the Spurs will play the Pistons. Thus, they both possess sufficient reasons to believe that the Spurs will win. They both infer that the Spurs will win. But Ponens uses *modus ponens*, while Lacy uses the *modus profusus* rule: for any P, Q, and R, infer R from P & Q.

Lacy possesses a good reason to believe that the Spurs will win. But he fails, I claim, to comply with this reason. He is insensitive to the deductive relation between this reason and his belief. This is obvious, given that he is following a crazy rule.

While our view has the same advantages over internalist views that orthodox reliabilism has, it also has advantages over orthodox reliabilism. Orthodox reliabilism struggles to explain why unwittingly reliable clairvoyants are unjustified. This is painfully true if orthodox reliabilism is given consequentialist foundations. On such a view, epistemic virtues are reliable abilities to hit the mark of truth, and the difference between justified belief and knowledge is a difference between a mere exercise of this ability that is compatible with luck and a genuine manifestation. But there is no reason why an unwittingly reliable clairvoyant couldn't manifest such a truth-instrumental ability in forming her clairvoyant beliefs.

6.2.2 Explanatory Satisfactoriness

Sometimes I hear it said that although epistemic consequentialism might fail to accommodate our case-based intuitions, it at least has the virtue of explaining why we value epistemic justification and other normative epistemic statuses. In this way, epistemic consequentialism might be defended in the way act consequentialism is defended: while it is counterintuitive, it is more explanatorily satisfying and does a better job of keeping us in touch with what matters (*viz.* value). But these claims are, I believe, backwards. Epistemic consequentialism makes a mystery of the actual epistemic value of virtually everything. Our view, by contrast, fits in with a better account of derivative epistemic value that also illuminates the value of internalist statuses like rationality.

If it takes an epistemic consequentialist form, reliabilism is cornered into giving an instrumental account of the epistemic value of justification, rationality, and knowledge if it does not claim that they are fundamentally epistemically good (which is implausible). I have argued elsewhere that it is instrumentalism, not epistemic value monism, that leads to the swamping problem.³² If rationalizing and justifying competence have merely instrumental value relative to true belief, is unclear why being justified or rational would make a belief better if it is already true. This problem is not merely a problem about knowledge. A true belief's being justified makes it epistemically better even if it is Gettiered.

We need to explain that. It is hard to see what explanation if at the foundational level we are consequentialists and we want to keep our list of fundamental epistemic values short and simple. By contrast, the view I've sketched meshes well with a different, non-instrumental story about the derivative epistemic value of rationality and justification. Being evidence-sensitive is a way to be truth-oriented that isn't merely instrumental: being evidence-sensitive has derivative value from a truth-oriented point of view in virtue of the fact it is a way to *respect* the truth norm. It is generally plausible that being correctly responsive to the reasons that bear on whether one is conforming to a norm is a way of respecting that norm. Given the link between reasons and respect, our view can exploit this fact to explain why rational belief and justified belief have epistemic value as a special case of the more general truism that it is better to conform to norm by respecting it than by failing to respect it. A foundationally consequentialist view that keep its epistemic axiology simple cannot exploit this model in a coherent way. For one can manifest a reliable ability

³²See Sylvan (2012, Forthcoming-b).

to hit the mark of truth in a blind and irrational way that couldn't constitute respect for accuracy.

7 Concluding Remarks

Of course, there remains a lot more work to do in developing a version of non-consequentialist reliabilism, but the main goal of this paper has just been to explain why epistemologists attracted to reliabilism should head in this direction. Let's briefly take stock of how this goal has been achieved.

I pursued this goal in two stages, beginning with the simpler task of explaining why reliabilism's fate needn't depend on epistemic consequentialism's. In §2, I laid the foundations for the divorce of reliabilism and epistemic consequentialism by drawing attention to a distinction between three kinds of theorizing: non-normative theorizing about the metaphysical constitution of epistemic properties, on the one hand, and normative theorizing at the factoral and foundational levels, on the other hand. Epistemic consequentialism is an example of the last kind of theorizing, and reliabilism would be independent of epistemic consequentialism if it were understood as an example of either of the other two kinds of theorizing. I then argued in §3 that as a matter of historical fact, reliabilism has not primarily been a foundational normative theory. It started out as a piece of non-normative theorizing about the metaphysics of knowledge, and when it took a normative turn at the end of the 1970s, it wasn't immediately under the rubric of foundational normative theorizing. It was only in one admittedly important contribution to reliabilism in the 1980s that reliabilism took a foundational normative form allegedly similar to consequentialism. But even this form, as I turned to explain, wasn't analogous to any familiar version of consequentialism.

Of course, one might think that the historical record is a record of failure, and that the fact that reliabilism often wasn't a version of epistemic consequentialism says little about the theoretical format in which it should be couched. But as I noted in §4, there are many possible non-consequentialist foundations for reliabilism, and some of them are more attractive than the supposedly consequentialist foundations offered in Goldman (1986). Furthermore, there is a deeper argument, as we saw in §5, for restricting one's attention to these kinds of options. Reliabilism *shouldn't* be offered as a version of consequentialism, since the only version of reliabilism that would coincide with a recognizably consequentialist view would be a foundational first-order version which fails to provide principled foundations for the factors that common-sense intuition suggest to be important. Epistemic consequentialism is simply not a good basis for the most plausible versions of reliabilism. Hence, reliabilists ought to work in opposition to epistemic consequentialists. To do so, they will need to seek a partnership with some foundational non-consequentialist outlook. Here there is an embarrassment of riches, as I had noted in §4, and there is also a strong case to be made for one specific outlook that I outlined in §6 (though it would take another paper to single out it out as the best non-consequentialist reliabilism).

8 References

- Ahlstrom-Vij, K. and Dunn, J. 2014. 'A Defence of Epistemic Consequentialism.' *Philosophical Quarterly* 64: 541-551.
- Ahlstrom-Vij, K. and Dunn, J. Forthcoming. 'Is Reliabilism a Form of Epistemic Consequentialism?' *American Philosophical Quarterly*.
- Alston, W. P. 1985. 'Concepts of Epistemic Justification.' *The Monist* 68(1): 57-89.
- Alston, W. P. 1988. 'An Internalist Externalism.' *Synthese* 74: 265-83.
- Anderson, E. 1993. *Value in Ethics and Economics*. Cambridge, MA: Harvard University Press.
- Armstrong, D. M. 1961. *Perception and the Physical World*. London: Routledge and Kegan Paul.
- Armstrong, D. M. 1968. *A Materialist Theory of the Mind*. London: Routledge.
- Armstrong, D. M. 1973. *Belief, Truth and Knowledge*. Cambridge: Cambridge University Press.
- Arpaly, N. 2003. *Unprincipled Virtue*. Oxford: Oxford University Press.
- Baehr, J. 2011. *The Inquiring Mind*. Oxford: Oxford University Press.
- Baron, M. 1995. *Kantian Ethics Almost without Apology*. Ithaca: Cornell University Press.
- Beddor, R. 2015. 'Process Reliabilism's Troubles with Defeat.' *Philosophical Quarterly* 65: 145-159.
- Benson, P. 1987. 'Moral Worth.' *Philosophical Studies* 51: 365-382.
- Bergmann, M. 1997. 'Internalism, Externalism, and the No-Defeater Condition.' *Synthese* 110: 399-417.
- Berker, S. 2013a. 'Epistemic Teleology and the Separateness of Propositions.' *Philosophical Review* 122: 337-393.
- Berker, S. 2013b. 'The Rejection of Epistemic Consequentialism.' *Philosophical 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.
- Copp, D. 2009. 'Toward a Pluralist and Teleological Theory of Normativity.' *Philosophical Issues* 19: 21-37.
- Copp, D. 2014. 'Indirect Epistemic Teleology Explained and Defended' in Fairweather, A. and Flanagan, O. (eds.) *Naturalizing Epistemic Virtue*. Cambridge: Cambridge University Press.
- Dean, R. 2006. *The Value of Humanity in Kant's Moral Theory*. Oxford: Oxford University Press.
- Dretske, F. 1971. 'Conclusive Reasons.' *Australasian Journal of Philosophy* 49: 1-22.
- Dretske, F. 1981. *Knowledge and the Flow of Information*. Cambridge, MA: The MIT Press.
- Fine, K. 2002. 'The Varieties of Necessity' in Gendler, T. and Hawthorne, J. (eds.) *Conceivability and Possibility*. Oxford: Oxford University Press.
- Firth, R. 1981. 'Epistemic Merit, Intrinsic and Instrumental.' *Proceedings and Addresses of the American Philosophical Association* 55: 5-23.
- Foley, R. 1987. *The Theory of Epistemic Rationality*. Cambridge MA: Harvard University Press.
- Fricker, M. 2007. *Epistemic Injustice*. Oxford: Oxford University Press.
- Fumerton, R. 2001. 'Epistemic Justification and Normativity' in Steup, M. (ed.) *Knowledge, Truth, and Duty*. Oxford: Oxford University Press.
- Goldman, A. I. 1967. 'A Causal Theory of Knowing.' *Journal of Philosophy* 64: 357-372.
- 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. 1983. 'Review of Philosophical Explanations.' *Philosophical Review* 92(1): 81-88.
- Goldman, A. I. 1988. 'Strong and Weak Justification.' *Philosophical Perspectives* 2: 51-69.
- Goldman, A. I. 1986. *Epistemology and Cognition*. Cambridge, MA: Harvard University Press.
- Goldman, A. I. 1993. 'Epistemic Folkways and Scientific Epistemology.' *Philosophical Issues* 3: 271-285.
- Goldman, A. I. 2007. 'Philosophical Intuitions: Their Target, Their Source, and Their Epistemic Status.' *Grazer Philosophische Studien* 74: 1-26.
- Goldman, A. I. 2011. 'Toward a Synthesis of Evidentialism and Reliabilism?' in Dougherty, T. (ed.) *Evidentialism and its Discontents*. Oxford: Oxford University Press.
- Goldman, A. I. and Beddor, R. 2015. 'Reliabilist Epistemology' in Zalta, E. (ed.) *The Stanford Encyclopedia of Philosophy (Winter 2015 Edition)*. URL = <<http://plato.stanford.edu/archives/win2015/entries/reliabilism>>
- Goldman, A. I. and Pust, J. 1998. 'Philosophical Theory and Intuitional Evidence' in Depaul, M. and Ramsey, W. (eds.) *Rethinking Intuition*. Lanham: Rowman & Littlefield.
- Greco, J. 1999. 'Agent Reliabilism.' *Philosophical Perspectives* 19: 273-296.
- Greco, J. 2005. 'Justification Is Not Internal' in Steup, M. and Sosa, E. (eds.) *Contemporary Debates in Epistemology*. Oxford: Blackwell.
- Greco, J. 2010. *Achieving Knowledge*. Cambridge: Cambridge University Press.
- Harman, E. 2011. 'Does Moral Ignorance Exculpate?' *Ratio* 24(4): 443-468.
- Herman, B. 1981. 'On the Value of Acting from the Motive of Duty.' *Philosophical Review* 66: 359-382.
- Jones, W. 1997. 'Why Do We Value Knowledge?' *American Philosophical Quarterly* 34: 423-439.
- Kagan, S. 1992. 'The Structure of Normative Ethics.' *Philosophical Perspectives* 6: 233-242.
- Kagan, S. 1997. *Normative Ethics*. Westview Press.
- 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. 1785/2012. *Groundwork for the Metaphysics of Morals*. Cambridge: Cambridge University Press.
- Kornblith, H. 2008. 'Knowledge Needs No Justification' in Smith, Q. (ed.) *Epistemology: New Essays*. Oxford: Oxford University Press.
- Kripke, S. 2011. *Philosophical Troubles*. Oxford: Oxford University Press.
- Littlejohn, C. 2014. 'Fake Barns and False Dilemmas.' *Episteme* 11: 369-389.
- Lyons, J. 2009. *Perception and Basic Beliefs*. Oxford: Oxford University Press.
- McDowell, J. 1994. 'Knowledge by Hearsay' in Matilal, B. K. and Chakrabarti, A. (eds.) *Knowing from Words*. Dordrecht: Kluwer.
- Montmarquet, J. 1986. 'Epistemic Virtue.' *Mind* 96: 482-97.
- Nozick, R. 1981. *Philosophical Explanations*. Cambridge, MA: Belknap Press.
- Parfit, D. 2011. *On What Matters*. Oxford: Oxford University Press.
- Price, H. H. 1953. *Thinking and Experience*. London: Hutchinson & Co.
- Pritchard, D. 2012. 'Anti-Luck Virtue Epistemology.' *Journal of Philosophy* 109: 247-279.
- Ramsey, F. P. 1931a. 'Knowledge' in Braithwaite, R. B. (ed.) *Foundations of Mathematics and Other Logical Essays*. London: Routledge and Kegan Paul.
- Ramsey, F. P. 1931b. 'Truth and Probability' in Braithwaite, R. B. (ed.) *Foundations of Mathematics and Other Logical Essays*. London: Routledge and Kegan Paul.
- Raz, J. 2011. *From Normativity to Responsibility*. Oxford: Oxford University Press.
- Roush, S. 2005. *Tracking Truth*. Oxford: Oxford University Press.

- Scanlon, T. M. 1998. *What We Owe to Each Other*. Cambridge, MA: Harvard University Press.
- Sellars, W. 1956. 'Empiricism and the Philosophy of Mind' in Feigl, H. and Scriven, M. (eds.) *Minnesota Studies in the Philosophy of Science*. Minneapolis: University of Minnesota Press.
- Siegel, S. 2012. 'Cognitive Penetrability and Perceptual Justification.' *Nous* 46: 201–222.
- Sosa, E. 1991. *Knowledge in Perspective*. Cambridge: Cambridge University Press.
- Sosa, E. 2007. *Apt Belief and Reflective Knowledge, Volume 1*. Oxford: Oxford University Press.
- Sosa, E. 2015. *Judgment and Agency*. Oxford: Oxford University Press.
- Stratton-Lake, P. 2000. *Kant, Duty, and Moral Worth*. London: Routledge.
- Star, D. 2015. *Knowing Better*. Oxford: Oxford University Press.
- Swain, M. 1981a. 'Justification and Reliable Belief.' *Philosophical Studies* 40: 389–407.
- Swain, M. 1981b. *Reasons and Knowledge*. Ithaca: Cornell University Press.
- Swain, M. 1985. 'Justification, Reasons, and Reliability.' *Synthese* 64: 69–92.
- Sylvan, K. 2016. 'Epistemic Reasons I: Normativity.' *Philosophy Compass* 11(7): 364–376.
- Sylvan, K. Forthcoming-a. 'Knowledge as a Non-Normative Relation.' *Philosophy and Phenomenological Research*.
- Sylvan, K. Forthcoming-b. 'Veritism Unswamped.' *Mind*.
- Sylvan, K. Ms. 'An Epistemic Non-Consequentialism.'
- Thomson, J. J. 2008. *Normativity*. Chicago: Open Court.
- Turri, J. 2010. 'On the Relationship between Propositional and Doxastic Justification.' *Philosophy and Phenomenological Research* 80: 312–326.
- Unger, P. 1968. 'An Analysis of Factual Knowledge.' *Journal of Philosophy* 65: 157–170.
- Vogel, J. 2000. 'Reliabilism Leveled.' *Journal of Philosophy* 97: 602–623.
- Wood, A. 1999. *Kant's Ethical Thought*. Cambridge: Cambridge University Press.
- Zagzebski, L. 1996. *Virtues of the Mind*. Cambridge: Cambridge University Press.
- 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.