



Value Promotion and the Explanation of Evidential Standards

Tricia Magalotti¹ 

Received: 3 February 2021 / Accepted: 23 December 2021 / Published online: 12 February 2022
© The Author(s) 2022

Abstract

While it is commonly accepted that justified beliefs must be strongly supported by evidence and that support comes in degrees, the question of how much evidential support one needs in order to have a justified belief remains. In this paper, I consider how the question about degrees of evidential support connects with recent debates between consequentialist and deontological explanations of epistemic norms. I argue that explaining why strong, but not conclusive, evidential support is required for justification should be one explanandum that such theories seek to explain. Furthermore, I argue that foundational theories that appeal to the promotion of epistemic value (especially consequentialism, but perhaps also some versions of epistemic deontology) are better suited to provide such an explanation.

1 Introduction

Many epistemologists agree that (a) for a belief that p to be justified, p must be supported by one's evidence and (b) support comes in degrees. But how much support one needs is a notoriously difficult question. Bonjour (2009) describes the problem in terms of arbitrariness: how does one specify a non-arbitrary degree of required evidential support? One approach would be to evade arbitrariness by endorsing one of the two limiting views, requiring either conclusive support or the mere favoring of p over $\sim p$ by the evidence. Another would be to dismiss the problem as a mere annoyance. We have an intuitive idea of where the bar for justification is set—relatively high but lower than conclusive support. No need, one might think, to worry about getting more precise than this. Or, one might claim that we should not expect to be able to identify some precise cutoff since the boundaries of the threshold are vague.

Tempting as these responses are, they fail to adequately apprehend Bonjour's concern. The troubling thing about this arbitrariness is *not* that we cannot identify

✉ Tricia Magalotti
tmagalotti@gmail.com

¹ Department of Philosophy, Stockholm University, Stockholm, Sweden

some particular degree of required evidential support. Even supposing that we could identify some such value, say by testing our intuitions, this value would remain arbitrary in a more important sense, namely in the sense of lacking a normative *explanation* for why it is *this* degree of evidential support, rather than some other degree, that is normatively required for justification and for knowledge.

The recent wave of enthusiasm regarding consequentialism and deontology in epistemology prompts a reconsideration of this explanatory aspect of Bonjour's challenge. Consequentialism and deontology are *foundational* theories of epistemic justification, theories about the fundamental explanation of why justified beliefs are justified. The benefits of bringing these two sets of issues together are twofold. From one end, examining the arbitrariness problem through the lens of foundational epistemological theories provides us with a structural starting point from which we can elaborate potential explanations. And, from the other end, the arbitrariness problem provides us with one of the explananda that should be addressed by foundational theories of epistemic justification.

In this paper I consider how consequentialist and deontological theories might explain why the evidential bar for justification is set where it is. I argue that consequentialism has a readily available solution to the arbitrariness problem, whereas the prospects for a deontological solution are dubious at best. The upshots are, first, a proposal about how to avoid problematic explanatory arbitrariness regarding the standards for epistemic justification and, second, a novel consideration in favor of foundational theories that emphasize value promotion.

The paper is structured as follows. In Sect. 2, I provide some background on epistemic consequentialism and deontology, focusing especially on the explanation of evidential norms. In Sect. 3, I summarize the arbitrariness worry and develop what I call the Arbitrary Standards Challenge, which any successful explanation of evidential norms must be able to meet. In Sect. 4, I outline a consequentialist solution to this challenge. In Sect. 5, I consider whether a deontological version of this solution is viable. In Sect. 6, I show how Sect. 3–5 together form an argument in favor of consequentialism (and potentially some versions of deontology). The conclusion is modest but valuable: there is an important fact about justification that is better explained by appeal to value promotion than by any “pure” deontological account (i.e., one without promotional duties).

2 Consequentialist and Deontological Explanations of Evidential Norms

I will begin with some background. This paper aims to explain the epistemic norm that requires us to have good evidence for our beliefs. Thus, what follows will be most interesting to those who are sympathetic with the idea that there is such a norm. While I will not argue for it here, it may be worth highlighting that the general idea is an intuitive one: whether one is justified in believing p depends on whether one has good evidence for that belief. I will call this norm *Evidential Necessity*. More precisely:

Evidential Necessity: A belief that p is justified for S only if believing that p fits S 's evidence.¹

A few clarifications: first, this modification of Feldman and Conee's (1985) classical evidentialism is basic in the sense that it omits optional commitments sometimes associated with evidentialism, for example mentalism or access internalism.² Second, it remains neutral with respect to the interpretation of evidential fit. And third, it states only a necessary condition on justification. Thus, it does not entail evidentialism about justification, given that the latter also purports to state a *sufficient* condition for justification. It is compatible also with other accounts, for example hybrids of evidentialism and reliabilism, that take forming beliefs on the basis of good evidence to be a necessary but not sufficient condition for justification.³

Above, I described consequentialism and deontology as foundational theories of epistemic justification. This terminology comes originally from Kagan's (1992) discussion of ethical consequentialism and deontology and has helpfully been applied to the epistemic domain by Sylvan (2020). Foundational theories are contrasted with factoral theories. A *factoral theory* of some normative status details the properties that are relevant to whether something has that normative status. In the case of epistemic justification, factoral theories include evidentialism, reliabilism, virtue epistemology, and some versions of epistemic utility theory (EUT). These theories aim to list the properties that a belief must have in order to qualify as epistemically justified. Evidential Necessity would be classified as part of a factoral theory of justification. *Foundational* theories of a normative status, in contrast, aim to provide the fundamental explanation for *why* those properties are relevant to the normative status. If our factoral theory of epistemic justification tells us that a belief is justified iff it is formed on the basis of strong evidence, then our foundational theory will provide an explanation of *why* it is the beliefs formed on the basis of strong evidence that are justified.

The *epistemic consequentialist* claims that the fundamental explanation of why justified beliefs are justified is that they are formed in ways that promote fundamental epistemic value. This is a form of rule rather than act consequentialism. The relevant consequences are those of forming beliefs in the *way* in which this belief was formed, not those of a particular *act* of belief formation. I will dismiss act forms of epistemic consequentialism from the start for two reasons. First, they lead to serious tradeoff problems. Sometimes believing in accordance with one's evidence requires believing in a way that does not promote epistemic value in that particular instance.⁴ And second, none of the representative consequentialist views from the literature are of the act variety.

¹ This is modified from Conee and Feldman's (2004) classical account of evidentialism.

² *Mentalism* is the view that the justificatory status of a subject's beliefs supervene on the subject's mental states, events, and conditions. *Access internalism* is the view that a subject knows that p only if she has access to that which justifies her belief (e.g., the evidence on which the belief is based).

³ See, for example, Comesaña (2010) and Goldman (2011).

⁴ See Berker (2013), Ahlstrom-Vij and Dunn (2014).

The standard representatives of epistemic consequentialism are reliabilism and EUT. Reliabilism is usually thought to entail consequentialism, but since reliabilism is a factorial theory and consequentialism is a foundational theory, it is better to say that consequentialism provides the best explanation for reliabilism's first-order principles. Reliabilism holds that a belief is justified just in case the process that produced it is reliable with respect to the production of true belief.⁵ The reliabilist's explanation of this factorial theory will appeal in a straightforward way to the promotion of epistemic value.⁶ It is *because* reliable processes are truth-promoting that they are justification-conferring. Examples of consequentialism in the domain of EUT include projects that aim to vindicate the normative legitimacy of certain epistemic principles by demonstrating their tendency to be truth-promoting.⁷ Here again, the normative status of a belief is taken to depend on the epistemic consequences of believing in a particular way.

Epistemic deontology is often defined as the denial of the claim that the fundamental explanation of the norms underlying epistemic justification is that believing in accordance with those norms promotes epistemic value. But, in order to compare the explanatory merits of a view with those of competing views, we need a more substantive characterization. Thankfully, Sylvan (2020) has outlined such a proposal in some detail. Sylvan distinguishes between two importantly different kinds of deontological views: brute and non-brute. According to *brute epistemic deontologists*, facts about which beliefs are justified are brute facts. There is no further explanation to be had of why, say, a belief is justified iff it is formed on the basis of good evidence or, alternatively, why a belief is justified iff it is formed using a reliable process.⁸

However, one could alternatively hold that there is a non-promotional explanation of these norms. This is the *non-brute* deontological position. Ethical Kantians explain the rightness or wrongness of actions in terms of *respect* for moral value (e.g., personhood). Sylvan's own epistemic version of Kantianism adopts a parallel structure: the fundamental explanation of why justified beliefs are justified is that they manifest *respect* for the fundamental epistemic value of accuracy. While Sylvan endorses *veritism*, the view that the only thing of fundamental epistemic value is accuracy, I wish to remain neutral regarding the nature of fundamental epistemic value. Accordingly, veritism is compatible with but not entailed by each of the three characterizations below. Here, then, are the three foundational views that I will be considering:

Epistemic consequentialism: The fundamental explanation of why justified beliefs are justified is that they are formed in a manner that promotes fundamental epistemic value.

⁵ See Goldman (1979).

⁶ More recently, Sylvan (2018) has argued that this is not the right understanding of reliabilism, as reliabilism should not be interpreted as a theory of justification. These issues are interesting but orthogonal to our present focus.

⁷ See Joyce (1998), Leitgeb and Pettigrew (2010) and Pettigrew (2013).

⁸ See Chisholm (1966) for an example of brute epistemic deontology.

Brute epistemic deontology: The fundamental explanation of why justified beliefs are justified is that they are in conformity with the norm of belief (where there is no deeper explanation for this norm).

(Non-brute) epistemic Kantianism: The fundamental explanation of why justified belief are justified is that they manifest respect for fundamental epistemic value.

Since we will be concerned here with the explanation of Evidential Necessity in particular, we can modify these principles accordingly:

Consequentialist EN: The beliefs that are justified for S are those that fit S's evidence because forming beliefs that fit one's evidence promotes fundamental epistemic value.

Brute deontological EN: The beliefs that are justified for S are those that fit S's evidence because forming beliefs that fit one's evidence is in conformity with the norm of belief (where there is no deeper explanation for this norm).

Kantian EN: The beliefs that are justified for S are those that fit S's evidence because forming beliefs that fit one's evidence manifests respect for fundamental epistemic value.

These foundational theories offer different explanations for the same epistemic norm. Interestingly, evidentialism about justification (and with it Evidential Necessity) seems to be compatible with both consequentialism and deontology.⁹ There are intuitive ways of explaining evidentialism in both the consequentialist and deontological frameworks. On the consequentialist side, it is plausible that believing in accordance with one's evidence is the best way we have to promote epistemic value, and that this could explain evidentialist principles. And on the deontological side, it also seems reasonable that believing in accordance with one's evidence manifests respect for truth, and that this could explain those same principles.

Accordingly, anyone who accepts Evidential Necessity as part of their factorial theory will need to consider, among other things, which of the foundational theories provides the best explanation for this principle. There are various aspects of a factorial theory that we may want a foundational theory to be able to explain. In the following section, we will see that one aspect that calls out for explanation is the degree of evidential support that is required for justification.

3 The Arbitrary Standards Challenge

In Sect. 2, we saw that the best foundational theory is the one that provides the best explanation of epistemic norms. In this section, I will argue that one norm that stands in need of explanation is, as BonJour's challenge suggests, how much evidential support is required for a belief to meet the justification condition on knowledge.

⁹ See Conee (1992) for a more consequentialist-leaning explanation and Feldman (2000) for a more deontological one.

A. The determinacy problem and the explanatory problem

By “support” I mean the following: whether a belief that p fits a body of evidence depends on the degree of support that a body of evidence lends toward believing p . A belief fits (or doesn’t) a body of evidence, and a body of evidence supports (or doesn’t) a belief that p . Support comes in degrees. We have stronger evidential support for believing that we will roll higher than a one than we do for believing that we will roll higher than a two on a normal die. This gives rise to what I call the Degree Question.

Degree Question: To what degree must a body of evidence support a proposition p in order for a belief that p to be justified by that evidence?

Neither of the two limiting answers, conclusive support nor mere support, seem correct. Conclusive support is too restrictive. For illustration, consider a probabilistic account of fit according to which in order for p to fit some body of evidence e , the conditional probability of p on e must be high.¹⁰ For this view to be plausible, we must interpret “high” as something lower than 1, as this Cartesian standard of certainty would rule out the vast majority of ordinary cases of knowledge. On the other hand, mere support is too permissive. One is not justified in believing that p when one’s evidence just barely supports p . In probabilistic terms, one’s belief that p would not meet the justification condition on knowledge if one’s evidence only supports p to degree 0.51. We would be exceedingly hesitant to count as knowledge a belief formed on the basis of such meager evidential support.¹¹

According to Bonjour (2009), if neither of these answers is plausible, then there is no non-arbitrary answer to the Degree Question and thus no non-arbitrary account of knowledge. Bonjour’s own prognosis is grim. Indeed, he wonders whether we were wrong to focus so much epistemological attention on the concept of knowledge:

How important could it be (and why) that the strength of one’s reasons or justification for a claim is above rather than below a line that cannot clearly be nonarbitrarily defined? (Bonjour, 2009, 42)

There are two distinct worries lurking in Bonjour’s question. The first concerns our inability to identify the precise degree of evidential support that is necessary for justification/knowledge. It lies somewhere between mere and conclusive support, but we are seemingly unable to locate precisely where. Call this the *determinacy problem*. We may never be in a practical position to identify this threshold. However, if we could identify some in principle method for locating it, this would at least demonstrate metaphysical determinacy. The second worry is, to my mind, much more

¹⁰ The *conditional probability* of A given B is the probability that A will occur given that B occurs. For example, the probability that I roll a 4 (on a normal 6-sided die) given that I roll an even number is $1/3$, whereas the unconditional probability that I roll a 4 is $1/6$.

¹¹ Feldman suggests, without robust defense, something along the lines of the mere support criterion for justification when he says that “you should believe when your evidence is supportive rather than neutral” (2000, 681). But another interpretation is that “supportive” is just whatever amount of support is required to reach the threshold, wherever that may be.

serious. Even if we can correctly identify the threshold, we lack a normative explanation for *why* it is there. This is the *explanatory problem*. And it is far more difficult to escape. For example, suppose (contrary to fact) that we correctly identified conclusive support as the appropriate threshold. Even so, we would lack an explanation for *why* a conclusively supported belief is justified and a non-conclusively supported belief unjustified. What gives one belief its special normative status? The same holds, *mutatis mutandis*, if we suppose that only mere support is required.

Here are some ways in which one might try to escape the explanatory problem, none of which I take to be successful:

Vagueness

Suppose that the boundary of the required support threshold is vague. Even if there is no precise cut-off for how much evidential support is required for justification, some beliefs will fall squarely within that range and other beliefs will fall squarely outside of it. We are still faced with the explanatory question of why some of those beliefs qualify and others do not. Why is the vague threshold where it is and not much higher or much lower?

Pragmatic encroachment

Suppose, as Fantl and McGrath (2009) suggest, that the evidential threshold for knowledge that *p* is whatever is required to make acting as though *p* is true practically rational. On this view, whether one has enough evidence to know that *p* depends, at least partially, on the practical stakes. If the practical stakes are very high, then this can result in a very high evidential threshold. Although this approach may appear to provide a determinate answer to the Degree Question, there are reasons for doubting that it succeeds. One major problem is this. Pragmatic encroachment theorists must hold either that the threshold required for knowledge is fully determined by pragmatic features of the situation or that it is not. If the pragmatic features do *not* fully determine the threshold, then pragmatic encroachment does not yet provide a solution, at least not on its own. As Brown (2014) demonstrates, if the pragmatic features of the situation *do* fully determine the degree, then the view seems to have unacceptable consequences in very low stakes situations.¹² In these situations, it will be rational to act as though *p* is true; and yet, it does not seem as though *S* knows that *p*. One response to this problem, as Brown notes, is to require some minimum degree of evidential support that is epistemically required in all cases regardless of the practical stakes. But this would reintroduce the degree problem at the level of the new across-the-board evidential requirement. Furthermore, even if pragmatic encroachment did provide an answer to the Degree Question, pragmatic encroachment itself is a highly contentious view, and so I take it that it would still be worth searching for a solution that is “purely epistemic,” so to speak.

Degreed notion of “justification”

Suppose that there is a non-binary sense of “justification” such that when one has stronger evidential support for *p*, one has more justification for *p*.¹³ Then there is no

¹² Brown also provides much more extensive argumentation against the pragmatic encroachment response to the degree problem than I give here.

¹³ See McCain (2014).

need to worry about where this threshold is. While it is plausible enough that there is a sense in which how much justification one has for believing *p* is directly correlated with how much evidential support one has for *p*, the legitimacy of this concept does not obviate the theoretical need for a binary concept of justification.

First, we need the binary concept of justification to make sense of justification as a necessary condition on knowledge. Even if we can have different amounts of justification, there needs to be some cutoff between those beliefs that qualify for knowledge and those that do not. Lest we also allow that knowledge admits of the degrees, this requires binary justification.

Second, it is often thought that in order for one belief to provide justification for another, the initial belief must itself be justified. Once again, this requires a binary notion. Either one is permitted to use their belief that *p* as a premise in further reasoning or not. And so again there must be some cutoff. How much justification does one's belief that *p* need to have in order to permissibly be used as a premise in reasoning? For both of these reasons, it appears that the binary notion of justification is so not easily discarded.

Explanationist fit

It may appear as though (*best*) *explanationist* accounts of evidential fit avoid the threshold problem. According to these accounts, a belief that *p* fits a body of evidence *e* (for a subject *S*) just in case *p* is part of the best explanation available to *S* for why *S* is in possession of *e*.¹⁴ One may see no need on this view to identify any threshold. Either *p* is part of the best explanation for why *S* is in possession of *e* or it is not. But, despite initial appearances, explanationist versions of evidentialism do not avoid the explanatory problem.

One serious objection to explanationism, identified by Byerly and Martin (2015), is that it is subject to a *bad lot problem*: sometimes the best explanation available to an epistemic subject for why they possess certain evidence is not a *good* explanation. If one has only bad explanations available, then believing the best one is not epistemically justified. One should rather withhold judgment. Explanationism, however, has the result that believing the best of a bad lot of available explanations *is* epistemically justified.¹⁵ The obvious solution to this problem, also suggested by Byerly and Martin (2015), is to require that, in addition to being the best explanation for why the subject is in possession of *e*, the explanation of which *p* is a part must also be a *good* explanation. I agree that this would be an improvement to the account. But, it introduces a threshold notion. How good does an explanation have to be in order for one to be justified in believing its parts? So plausible versions of explanationist evidentialism must still contend with the Degree Question.

In summary, while some of these suggestions may allay our worries regarding the determinacy problem, they do nothing to address the explanatory problem. So, the problem stands.

¹⁴ This view is endorsed by Conee and Feldman (2004). A modified version has been defended by McCain (2014).

¹⁵ See McCain (2015) for a response to this objection.

B. The challenge

To adequately address the arbitrariness worry, we must address both the determinacy problem and the explanatory problem. In addition, we should add an extensional adequacy constraint. A solution to the arbitrariness problem must deliver results that accord with our pretheoretical notions about how strong one's evidential support must be, namely somewhere between mere and conclusive support.¹⁶ Together these three conditions make up what I call the Arbitrary Standards Challenge.

The Arbitrary Standards Challenge (ASC) A satisfying answer to the Degree Question will meet the following three conditions:

1. Determinacy: the proposed solution specifies an in principle method for determining the support threshold for the justification condition on knowledge;
2. Extensional adequacy: the proposed solution's results accord with the ordinary intuition that the support threshold for justification falls somewhere between mere support and conclusive support; and
3. Explanatory success: the proposed solution provides a normative explanation of why the threshold for support is where it is.

We should take note of how the Degree Question and ASC align with our previous discussion of factoral and foundational theories of epistemic justification. An answer to the Degree Question will be part of a factoral theory. If, say, our answer to the Degree Question is that the conditional probability of p given one's evidence must be at least 0.85, then this will be a justification-relevant property noted in the factoral theory. A solution to ASC, however, requires that we go beyond the factoral theory to the foundational theory. Whatever answer one gives to the Degree Question must be supplemented with a foundational theory in order to satisfy the explanatory success condition. Only when we have a fundamental explanation of why justified beliefs are justified will we be able to satisfy condition (3) and thereby provide a solution to ASC. In the next section, I propose a solution.

4 A Consequentialist Solution to the Degree Problem

In this section I argue that consequentialism can provide a satisfactory solution to ASC. First, let us review. The Degree Question asks to what degree some evidence must support a proposition p in order for a belief that p to be justified by that evidence. According to ASC, a satisfactory response to this question will meet the following three conditions: (1) it specifies an in principle method for

¹⁶ I take it that these pretheoretical intuitions concerning the threshold for justification apply also to knowledge in the following sense: where we are pretheoretically inclined to ascribe unjustified belief, we are also pretheoretically inclined to deny that the person has knowledge. This suggests that the kind of justification targeted by ASC is the same kind of justification that is necessary for knowledge.

determining the support threshold for the justification (determinacy); (2) its results accord with the ordinary intuition that the support threshold for justification falls somewhere between mere support and conclusive support (extensional adequacy); and (3) it provides a normative explanation of why the threshold for support is where it is (explanatory success).

To bring the proposed solution into focus, it will help to keep in mind the characterization of consequentialist EN from Sect. 2. Consequentialist EN holds that beliefs are justified for S only if they fit S's evidence because forming beliefs that fit one's evidence promotes fundamental epistemic value. Facts about which ways of believing promote epistemic value ground facts about justification. ASC shows that this (factoral) account of what it takes for a belief to be justified is under-specified. It does not indicate to what degree a body of evidence must support a proposition in order for the evidence to justify a belief in that proposition. Thus, we need a more precise account.

The general point that I wish to make is that the appeal to the promotion relation endows consequentialism with a distinctive advantage. The Degree Question points to a lack of clarity at the factoral level. Consequentialist EN can appeal to the promotion relation in order to settle this uncertainty.

We can understand Evidential Necessity as the claim that the justified beliefs are those formed in accordance with the following general rule:

RULE_G: Form a belief that p only if your evidence strongly supports p.

The Degree Question presses us to clarify what is required for a body of evidence to "strongly" support p. In order to answer the Degree Question, we must provide a value for d in the following more precise rule:

RULE_p: Form a belief that p only if your evidence supports p to degree d.

The Degree Question asks which value should be filled in for d in RULE_p. Each possible value for d, when inserted into RULE_p, yields a different version of the rule. Call the version containing the correct value for d RULE_{pC}. The consequentialist says that the explanation for why the justified beliefs are the ones formed in accordance with RULE_G is that forming beliefs in accordance with RULE_G best promotes the epistemic good. By this same reasoning, the explanation for why the justified beliefs are the ones formed in accordance with RULE_{pC} and not some other version of RULE_p, RULE_{p*}, with a different value for d, is that forming beliefs in accordance with RULE_{pC} best promotes epistemic value. To find the answer to the Degree Question, then, we must find which degree of evidential support maximizes the epistemic utility of the rule:

Value Promotion (VP) Response: The degree of evidential support that is necessary for justification is whatever degree of evidential support, when substituted for d, makes it the case that believing in accordance with RULE_p best promotes fundamental epistemic value.

Notice that the VP response itself is a factoral theory, a theory about which properties a belief must have in order to be epistemically justified. As I will argue

below, the VP response alone is enough to satisfy conditions (1) and (2) of ASC. In order to satisfy condition (3), however, we must combine the VP response with a foundational theory in order to yield a fundamental explanation of the VP response.

First is the determinacy condition, which requires the specification of an in principle method for determining the support threshold for justification. (Recall that the method must be an in principle method, not necessarily a practically feasible one.) In order to know what degree of support is necessary for justification, the VP response suggests that we must determine what value, when substituted for d , maximizes the epistemic value of believing in accordance with $RULE_p$.

Suppose that we could replicate a subject's mind for the purposes of running a cognitive simulation. The inputs would be psychological facts about the subject (e.g., facts the subject's intelligence, memory, level of curiosity, etc.).¹⁷ The simulation would determine which version of $RULE_p$ (which d value) generates the most epistemic value for a subject of that psychological type. The output of the program would not depend on the subject's actual beliefs nor on the actual evidential standards that the subject tends to use. The aim of the simulation is to determine which evidential rule a subject with a particular type of psychology would need to believe in accordance with in order to maximize the epistemic value of the resulting belief set.¹⁸ The rule could then be used as a standard to evaluate the beliefs of any subject who matches the original psychological description.

In order to get a fully determinate answer, we would need to decide how broad to make the psychological types that serve as the inputs. One possibility is to index rules to individual believers. This is my preferred option, as it seems to be the least arbitrary. However, one might argue that it is better to index rules to a psychological type that averages all humans, or perhaps one that is a representative prototype of humans. The success of the suggestion does not depend on a certain one of these options, but VP response advocates will need to make a theoretical choice about which option to adopt. Despite its practical difficulties, the metaphysical possibility of this method demonstrates the adequacy of the VP response with respect to the determinacy requirement.

The second requirement of ASC is extensional adequacy. The results delivered by the solution must accord with our pretheoretical intuitions. As discussed in Sect. 3, the threshold for justification should be somewhere lower than conclusive support and higher than mere support. And this is indeed the result that the VP response delivers.

We can start with mere support. The version of $RULE_p$ corresponding to mere support does not maximize the promotion of fundamental epistemic value, because

¹⁷ This will also include what sorts of cognitive impulses a subject has toward belief formation to begin with. Some subjects may be more inclined to form many beliefs while others are less so, independently of the level of evidential support that they have for those beliefs.

¹⁸ It would need to be programed into the simulation which things have fundamental epistemic value and how much. A veritist who thinks that true beliefs are the only thing of fundamental epistemic value (and false beliefs the only thing of fundamental epistemic disvalue), for example, would look for the optimal ratio of true to false belief.

it allows for the formation of too many false beliefs, which directly detracts from epistemic value (on most theories).¹⁹ So d must be higher than the value corresponding to mere support.

The version of $RULE_p$ corresponding to conclusive support will also fail to maximize the promotion of fundamental epistemic value. It would set the bar for evidential support so high that we would rarely, if ever, be justified in forming beliefs. To best promote epistemic value generally, we must take some minor epistemic “risks” every now and then. Consider two sets of beliefs. Set A contains 1 true belief and 0 false beliefs and Set B contains 99 true beliefs and 1 false belief. Intuitively, Set B has more fundamental epistemic value than Set A. False beliefs have some epistemic disvalue but not so much as to outweigh any possible positive value possessed by the true beliefs in that set.²⁰ If this is right, then there will be some other version of $RULE_p$ that better promotes fundamental epistemic value. Therefore, the VP response delivers the intuitive result that the threshold falls somewhere between mere support and conclusive support, thereby satisfying the extensional adequacy requirement.

One might worry that the extensional adequacy of the VP response is threatened by the fact that some beliefs have epistemic consequences that may not align with their epistemic integrity. For example, suppose that despite having read earlier that all trains were cancelled due to a strike, Angelo believes that his normal train will be leaving on time and heads to the station, where he finds a book about special relativity abandoned on a bench. Refusing to give up the belief that the train is on its way, he sits down on the bench and reads the entire book and, as a result, finally acquires a good understanding of a topic that always felt ungraspable. In the face of this kind of epistemic luck, how can we be confident that the value-maximizing rule will be one that requires a relatively high degree of evidential support? The answer to this worry lies in the observation that whereas the connection between Angelo’s evidentially unsupported belief and the positive epistemic consequences that follow is entirely accidental, the connection between an evidentially supported belief and its positive epistemic consequences is systematic. As a result, we can expect that the evidentially unsupported beliefs with good consequences will be far less frequent in comparison to the evidentially well-supported beliefs with good consequences. Since we are concerned with the rules that *in general* promote epistemic value, the effect of the former on the overall analysis should be negligible.

The final requirement of ASC is the explanatory requirement. For this, the VP response must be combined with a foundational theory. Suppose that the degree of evidential support required for justification is 0.8. The consequentialist can explain this as follows: the fact that one’s evidence must support p to at least degree 0.8 in order to be justified is explained by the fact that this is the degree of support such

¹⁹ Even most theories of final epistemic value that deny that true belief is the only thing of epistemic value grant an important role for truth. If knowledge or understanding (on factive conceptions of understanding) have final epistemic value, for example, these will also require truth.

²⁰ Parallel points could be made using views that accept, for example, knowledge or understanding, as bearers of final epistemic value.

that when a subject believes in accordance with a rule requiring that much support, they maximize the promotion of epistemic value (as compared to other potential rules). This explanation is congruent with the consequentialist explanation for why there is any evidential requirement on justification at all, namely that forming beliefs that are supported by one's evidence promotes fundamental epistemic value. We can call the combination of the VP response and consequentialism *Consequentialist VP*.

In sum, consequentialism has available a promising answer to the Degree Question, namely Consequentialist VP. The VP response meets conditions (1) and (2) of the ASC, determinacy and extensional adequacy. When combined with a consequentialist explanation, the VP response also meets condition (3), explanatory success. In the next section, we will consider whether the VP response can be combined with brute deontology and/or epistemic Kantianism to yield a similarly satisfying response to ASC.²¹

5 Deontological Explanations of the VP Response

Upon being presented with Consequentialist VP, deontological-leaning readers may find themselves with the impression that there is a possible deontological explanation of the VP response. In this section, I will consider whether the VP response can be combined with brute deontology and/or epistemic Kantianism in order to provide an alternative solution to ASC.

A. Brute Deontology

According to brute deontology, the fundamental explanation of why justified beliefs are justified is that they conform with the norm of belief (for which there is no deeper explanation). The brute deontologist may claim that the norm of belief requires that one form beliefs in accordance with $RULE_p$. They may even agree that the correct version of $RULE_p$ is the one that assigns the value to d that allows the subject who believes in accordance with $RULE_p$ to best promote fundamental epistemic value. And so they can agree that the justified beliefs are the ones that are formed in accordance with $RULE_{pC}$.

An initial worry arises from the fact that there is no direct route from brute deontology to the conclusion that the justified beliefs are the ones that are formed in accordance with $RULE_{pC}$. One cannot reason oneself from the claim that facts about which beliefs are justified are brute facts to the claim that the justified beliefs are the ones that are formed in accordance with $RULE_{pC}$. This ultimately leads to a failure

²¹ It is worth noting that reliabilists face their own threshold problem regarding the level of reliability a belief forming process must have in order to confer justification on the beliefs that are produced by that process. It seems plausible that some version of the VP response would also be a promising way for reliabilists to address this problem. However, since a full treatment of this topic would require discussion of issues that are specific to reliabilism, I will leave its exploration for another time.

to meet condition (3) of ASC, but we must be careful about stating the nature of the problem. I will first dismiss some potential misconstruals.

First, the problem is *not* that brute deontology is *inconsistent* with the VP response. Brute deontology places nearly no restrictions on the content of epistemic norms. As a result, it is *consistent* with myriad factorial theories of justification, including the one proposed by the VP response. Second, one might worry that for the brute deontologist to accept the the VP response would be ad hoc since adopting the VP response seems unmotivated from the brute deontologist's point of view. Given brute deontology alone, there is no reason for thinking that justification has any special relation to the promotion of epistemic value. I am inclined to think that this is not the real problem. When it comes to the factorial level, it is open to the brute deontologist to employ various methods of investigation. For example, they may analyze a variety of cases and construct a factorial theory of justification that is extensionally adequate using something like the method of reflective equilibrium.²² I am dubious about the outlook for such a project, as I find it doubtful that this method would yield a sufficiently determinate factorial theory. However, this is just one example of an alternative methodology. The point stands that the brute deontologist's inability to derive a factorial theory from their foundational theory does not entail that adopting the VP response would be ad hoc on their part.

I grant, then, that there is nothing preventing the brute deontologist from endorsing the VP response. Since the VP response satisfies conditions (1) and (2) of ASC, it follows that they can meet conditions (1) and (2). The crucial question is whether brute deontology can deliver on explanatory adequacy.

This is where we reach the central problem with brute deontology, namely that by its very nature it is unable to *explain why* the required degree of evidential support is d and not some other value. Justified beliefs are justified because they are formed in accordance with the proper epistemic norms. But the brute deontologist cannot explain why those norms have the content that they do. Of course, this comes at no surprise. Brute deontologists are simply not in the business of providing explanations for the content of epistemic norms. In Sect. 6 I will argue that the ability to provide this kind of explanation is indeed a desideratum of a foundational theory, and so brute deontology's inability to do so is a strike against it. For now I merely want to establish that even if brute deontology can accept the VP response, it remains unable to satisfy condition (3) of ASC.

B. Epistemic Kantianism

For epistemic Kantians, the fundamental explanation of why justified beliefs are justified is that they manifest respect for fundamental epistemic value. According to the VP response, the justified beliefs are the ones formed in accordance with $RULE_{PC}$. Combining epistemic Kantianism and the VP response yields the view that the

²² The method of *reflective equilibrium* is one in which one moves back and forth between one's judgments about particular cases or thought experiments on the one hand and one's theory on the other, making whatever revisions are required to maintain coherence. See Rawls (1971) for the classical defense.

justified beliefs are those formed in accordance with $RULE_{PC}$ *because* such beliefs manifest respect for fundamental epistemic value.

A major theme in this section will be the clarification of the notion of respect that is at issue. For Sylvan (2020), fundamental epistemic value (in his view, accuracy) demands something like Darwall's "recognition respect." One treats X with recognition respect when one treats (in the appropriate way) X as a deliberative constraint in reasoning. If Janna treats the fact that trespassing on some property would break the law as a reason not to trespass, she demonstrates recognition respect for the law. (Assume that she believes there to be no chance of suffering negative consequences resulting from her action.) Having recognition respect for accuracy, according to Sylvan, requires treating accuracy with the proper consideration in reasoning. In order to do this, Sylvan claims, one must conform one's beliefs to one's evidence. Already, we can question why this is. Why think that that what recognition respect for truth requires is conforming one's beliefs to the evidence rather than, say, believing the truth or forming beliefs in ways that reliably produce true belief? I will set this issue to the side, as I think that bigger problems arise even if we grant this initial step.

After narrowing the scope of respect to recognition respect, understood in an evidentialist way, questions remain about what respecting accuracy requires. Sylvan (2020) distinguishes between "weak respect" and "strong respect," where weak respect requires believing in accordance with what one's beliefs (and other doxastic attitudes) indicate to be true, whereas strong respect requires believing in accordance with what one's evidence indicates to be true. This characterization is infelicitous given that on some reasonable views of evidence one's evidence is provided by (or even constituted by) one's beliefs (and other doxastic attitudes). The idea seems to be that weak respect requires that the subject believes what she subjectively takes her evidence to support, whereas strong respect requires believing what her evidence objectively supports (from a third person perspective). Sylvan claims that it is the manifestation of strong respect for accuracy that is necessary for justified belief.

This move is even more problematic than the last. Why think that accuracy as such demands strong rather than weak respect? I agree that strong respect for truth, as glossed in the previous paragraph, aligns better with our notion of epistemic justification than does weak respect. But for the epistemic Kantian, it must be the case that the fundamental value of accuracy *demand*s strong (rather than weak) respect, and it is difficult to see what it is about truth that would demand *this* particular kind of respect. Why, for example, does the person who believes what she takes her evidence to support not demonstrate the right kind of respect? The crux of the problem is that the more specific we get about the *kind* of respect that it demands, the less plausible it is that truth itself could make such a demand.

Although the moves from recognition respect for accuracy, to following the evidence, to following the evidence in a particular way are contentious, let us grant all of them in order to evaluate how this epistemic Kantian picture might fit together with the VP response.

First is the question of whether epistemic Kantianism is compatible with the VP response. Recall that the VP response will be part of a factorial theory of justification. It says something about which properties are relevant to justification but says

nothing about *why* these are the justification-relevant properties. In particular, it says that the degree of evidential support required for justification is whatever degree is value-maximizing, but it does not say that this degree of support is required *because* it is value maximizing. A Kantian, then, could hold that the required degree of support is the degree that is value-maximizing, but that this degree of support is required *because* that is what recognition respect for accuracy requires. Accordingly, if the epistemic Kantian adopts the VP response in this way, then they too can satisfy conditions (1) and (2). Once again, the real challenge comes in satisfying condition (3), the explanatory success condition.

Epistemic Kantianism certainly seems better suited than brute deontology to explain the legitimacy of the VP response. The basic explanation is that truth demands a certain kind of respect, strong recognition respect, and that forming beliefs in accordance with $RULE_{PC}$ manifests strong recognition respect for truth. To know whether this explanation succeeds, we must look at some of the details.

We can distinguish between *direct* and *mediated* Kantian explanations of the VP response. A mediated explanation is one that appeals to a requirement to believe in ways that promote epistemic value and then explains this value promotion requirement in terms of respect for fundamental epistemic value. By anchoring the explanatory chain to respect, the mediated explanation retains the commitment to the explanatory fundamentality of respect for epistemic value. A direct explanation, by contrast, is one that is not mediated by appeal to the promotion of epistemic value but rather explains epistemic principles directly in terms of respect.

A direct epistemic Kantian explanation of the VP response is implausible. In order for a direct explanation to succeed, it would need to be the case that respect for fundamental epistemic value requires that one form beliefs in accordance with $RULE_{PC}$ and *not* $RULE_{p^*}$ or $RULE_{p^{**}}$. Otherwise, it would fail to explain why we are required to form beliefs in accordance with $RULE_{PC}$ and not some other version of the rule. For example, when explaining why the standard falls somewhere below conclusive support, it seems that one cannot help but to fall back into the language of value-promotion. One must appeal to something akin to the idea that requiring conclusive support for the formation of a belief would not, as a general policy, be conducive to the promotion of truth, understanding, etc. By contrast, there seems to be nothing in the concept of *respect* that can explain the lenience of the standard. Recall that for the epistemic Kantian, respect for some fundamental epistemic value consists in whatever response is *demand*ed by that value. Even granting that accuracy has fundamental epistemic value and demands strong recognition respect, there is still the question of why the value of accuracy demands some particular degree (or range of degrees) of evidential support. What is it about accuracy that demands this wildly specific kind of respect?

One could start from the other direction: justification does not seem to require conclusive support, and the justified beliefs are the ones that manifest respect for fundamental epistemic value; so, respect for fundamental epistemic value does not require conclusive support. But to adopt this line of reasoning is to reverse-engineer the concept of respect so that it accords with our concept of justification, and a concept cannot explain the phenomenon with which it was contrived to match. So much then for direct Kantian explanations for the VP-response.

The second option is to pursue an indirect explanation of the VP response that is mediated by an appeal to a duty to promote epistemic value. The thought would be that the VP response is explained by a duty to believe in ways that promote fundamental epistemic value, and that this promotional duty is in turn explained by the demand for respect, which is itself generated by fundamental epistemic value. On this view, the epistemic Kantian would grant that factoral theories of justification are (at least partly) explained by facts about value promotion. I will call this view value promotionism.

Value promotionism: Justified beliefs are justified (at least in part) *because* they are formed in value-promoting ways.

The difference between the consequentialist and the epistemic Kantian who endorses the indirect explanation is the “at least in part.” The epistemic Kantian, but not the consequentialist, thinks that there is a further explanation of Value Promotionism itself. We are epistemically obligated to form beliefs in value-promoting ways *because* doing so manifests respect for fundamental epistemic value.

This approach is the epistemic Kantian’s best chance at explaining the VP response. Still, there are major concerns. The first is that since this version of epistemic Kantianism supplements the consequentialist explanation with a further explanatory link, the inclusion of this additional link calls out for validation. If it serves no explanatory purpose, then the link is extraneous. As for explaining the VP response, I see no motivation for including the additional link. But this does not entail that there is no motivation to be found. The capacity to solve ASC is one desideratum of a foundational epistemic theory, but it is not the only one. And there may well be other desiderata that are better satisfied by a theory that appeals to respect for fundamental epistemic value at some level of the explanation.²³

The second concern for the indirect Kantian explanation is that, even supposing that there is some further theoretical need that requires an additional explanatory link, it must also be demonstrated that indirect Kantianism can provide this link. We will want to know what it is about fundamental epistemic value that demands these derived promotional duties, to know how the derived promotional duties are derived.

The third concern for the indirect Kantian explanation is that it is debatable whether this sort of view should be classified as genuinely Kantian. The case in favor of this classification, of course, is that the most fundamental explanation of why we are required to have a certain amount of evidential support appeals to respect for accuracy. There is another case to be made, however, that the primary explanatory work in this sort of account is done by value promotion, and not respect, and thus that the account is really a form of consequentialism in disguise. If this is true, then it seems that there is no version of epistemic Kantianism that is well-suited to providing a successful response to ASC.

²³ For example, Sylvan (2020) claims that one explanandum is the perspective-dependence of epistemic justification. Two people can be justified in believing different things in virtue of their occupying different perspectives. This strikes me as something that might be, at least *prima facie*, better explained by an appeal to respect than by an appeal to value promotion.

I leave it open here that if the epistemic Kantian can respond to these concerns in a satisfying way, then epistemic Kantianism too would be able to provide an explanation for the VP response and, thus, a solution to ASC. This leaves us with one potentially viable version of epistemic deontology.

6 ASC and the Epistemic Consequentialism Versus Deontology Debate

In this section I will argue that ASC highlights a consideration in favor of accepting value promotionism, the claim that justified beliefs are justified (at least in part) because they are formed in value-promoting ways. In Sect. 4, I argued that the VP response provides a promising answer to the Degree Question, one that is capable of satisfying conditions (1) and (2) of ASC, and that epistemic consequentialism can provide a good explanation of *why* the VP response is true. In Sect. 5, I argued that the only version of epistemic deontology that may have an explanation for the VP response is one that endorses value promotionism.

I have not yet argued that there is something unsatisfactory about non-promotional forms of epistemic deontology. In order to establish this stronger claim, we must address two further questions. First, is there some alternative solution to ASC that might be available to the epistemic deontologist? And second, why think that our foundational theory of justification needs to provide the sort of explanation demanded by ASC in the first place? I will respond these two questions in turn.

A. Alternatives to the VP response

I have argued that the only epistemic deontologists who are (potentially) able to provide an explanation for the VP response are those who endorse value promotionism. But many deontologists may object to the VP response itself, hoping instead to find some alternative answer. Unfortunately, as I have mentioned, positive deontological proposals are scant in the epistemological literature. In general though, it is unlikely that this strategy will be profitable. The reason why will be familiar from Sect. 5.

Consider again Sylvan's epistemic Kantianism. The reason why a direct Kantian explanation of the VP response fails is that respect is not determinate enough to explain the specificity of $RULE_{PC}$. There are numerous respect concepts, many of them closely related, which require different things from an epistemic subject. Notice that the problem is not linked to any particular features of the VP response. The lack of determinacy of the respect concept is problematic when it comes to explaining *any* purported answer to the Degree Question. The Kantian explanation will always be that the required degree of evidential support is required *because* this is what respect for fundamental epistemic value demands. But the question remains: *why* would respect for fundamental epistemic value demand *that* degree of evidential support (whatever that degree is)?

Perhaps, then, respect is the source of the problem. Maybe the deontologist could choose some other concept to serve as the basis for explanation. This approach too is

unpromising. The indeterminacy problem seems to generalize beyond one particular brand of deontology. This is because the other sorts of concepts that the deontologist is likely to cite are equally fraught. Concepts such as rationality, reasonableness, or epistemic responsibility seem to fare no better than respect for truth when it comes to delivering a determinate answer to the Degree Question. What grounds could there be, aside from promotional ones, for holding that epistemic responsibility requires one particular threshold of evidential support rather than another? Why is the person who forms a belief on 0.8 degree of evidential support epistemically responsible but the person who forms a belief on 0.79 degree of support is epistemically irresponsible? Similar questions hold for rationality or reasonableness. These are imprecise folk concepts. Indeed, one of the main goals of epistemology as a discipline is to provide more determinate accounts of these very concepts. And so they are just not the sort of concepts that can help us in answering the Degree Question.

Contrast these concepts with that of truth promotion. (Truth is just an example. Other accounts of fundamental epistemic value could be substituted here.) The promotion of truth is a determinate concept. We have a solid independent grasp on what it takes for a rule to be truth promoting and there is an available (in principle) methodology for investigating which rule is the best at promoting truth. It is because the concept of truth promotion (and promotion in general) has these features that it is poised to provide the sort of explanation called for by ASC. Unless there is some similarly determinate deontic concept that can play this same role, which does not seem to be the case, a non-promotional deontological solution will not succeed in responding to ASC.

B. Why ASC matters

I have argued that non-promotional versions of epistemic deontology cannot provide an adequate solution to ASC. Many deontologists will deny that this is a problem. Brute deontologists, for example, claim explicitly that there is *no* explanation for why justified beliefs are justified. I take it that this constitutes a sort of defeatism to which we should retreat only if no good explanations can be found. Other deontologists, perhaps some epistemic Kantians, might agree that epistemic norms do require some kind of explanation but claim that the sort of explanation demanded by ASC falls outside of the scope of what we should expect from a foundational theory of justification.

ASC demands an explanation for why epistemic norms have the particular content that they do. Why must a belief be supported by strong (but not conclusive) evidence in order to be justified? One might doubt that this falls within the domain of *normative* explanations. A normative explanation, one might think, need only explain where the normativity of the norm comes from, not where its content comes from. Accordingly, I would like to say a bit now about why I am inclined to think that content explanations of the sort that ASC is concerned with are relevant to the evaluation of foundational theories of justification.

I suspect that the dismissal of the importance of content explanations arises from one particular way of construing the explanandum. On this construal, we

have in mind some epistemic norm, which has both content and normative force. Since one can separate in one's mind these two different aspects, one is led to think that a normative explanation need only account for the normative force aspect. If so, a foundational theory's failure to explain the content aspect of the norm is no strike against the theory.

But there is another way of construing the explanandum. Suppose that we had in front of us a list of potential epistemic norms. At the factorial level, our question is *which* of these different items on the list has normative force. And at the foundational level, our question is *why* that item on the list is the epistemic principle that has normative force while others lack normative force. The explanation of why *this* principle is the one that has normative force is a normative explanation. And so the ability to provide a solution to ASC does indeed seem to be an important desideratum of a foundational theory of justification.

In summary, I have argued in this section that (a) there is unlikely to be a non-promotional deontological solution to ASC and (b) the ability to respond to ASC is a desideratum of a foundational theory of justification. The outcome of the discussion, then, is that we have one motivation for favoring consequentialism (and potentially promotional versions of deontology, depending on how well such theories are able to address the concerns from Sect. 5) over (non-promotional versions of) deontology.

7 Conclusion

In summary, consequentialism can provide an answer to the Degree Question that is determinate, extensionally adequate, and explainable by appeal to consequentialist principles. This suffices for a successful response to ASC. Epistemic consequentialism, then, has the resources to respond to BonJour's arbitrariness worry. Non-promotional versions of epistemic deontology are not in a position to provide a response that meets these requirements. Finally, forms of deontology that do accept value promotionism may be able to provide a satisfying response to ASC, but this is conditional on a successful response to the concerns for such theories outlined above.

While the foregoing considerations seem to favor consequentialism over non-promotional versions of deontology, I do not claim to have established the all things considered superiority of consequentialism here. For one thing, promotional deontological views deserve further exploration before a final comparison can be made. Furthermore, as noted above, the ability to solve ASC is only one among several desiderata for a foundational theory of epistemic justification. In order to provide a decisive verdict regarding which theory provides the best explanation of epistemic norms, the theories will have to be assessed with respect to each of those desiderata and the desiderata themselves weighted in importance. Accordingly, I will restrict my conclusions to the following claims: (1) the capacity to solve ASC is a

desideratum of a foundational epistemic theory; (2) consequentialists have a promising solution to ASC; and (3) only promotional versions of deontology have a chance of solving ASC, and the success of such a solution is contingent upon an adequate response to important concerns for these views.

Acknowledgements This paper has benefited greatly from discussion with many philosophers including audiences at the Institut Jean Nicod and the University of Edinburgh. I am especially grateful to Hilary Kornblith, Uriah Kriegel, and two anonymous referees from *Erkenntnis* for their extensive comments on previous drafts.

Funding Open access funding provided by Stockholm University.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Ahlstrom-Vij, K., & Dunn, J. (2014). A defence of epistemic consequentialism. *The Philosophical Quarterly*, 64(257), 541–551.
- Berker, S. (2013). The rejection of epistemic consequentialism. *Philosophical Issues*, 23(1), 363–387.
- Bonjour, L. (2009). *Epistemology: Classic problems and contemporary responses*. Rowman & Littlefield Publishers.
- Brown, J. (2014). Impurism, practical reasoning, and the threshold problem. *Noûs*, 48(1), 179–192.
- Byerly, T. R., & Martin, K. (2015). Problems for explanationism on both sides. *Erkenntnis*, 80(4), 773–791.
- Chisholm, R. (1966). *Theory of knowledge*. Prentice Hall. 2nd Edition 1977, 3rd Edition 1989.
- Comesaña, J. (2010). Evidentialist reliabilism. *Noûs*, 44(4), 571–600.
- Conee, E. (1992). The truth connection. *Philosophy and Phenomenological Research*, 52(3), 657–669.
- Conee, E., & Feldman, R. (2004). *Evidentialism: Essays in epistemology*. Clarendon Press.
- Fantl, J., & McGrath, M. (2009). *Knowledge in an uncertain world*. Oxford University Press.
- Feldman, R. (2000). The ethics of belief. *Philosophy and Phenomenological Research*, 60(3), 667–695.
- Feldman, R., & Conee, E. (1985). Evidentialism. *Philosophical Studies*, 48(1), 15–34.
- Goldman, A. I. (1979). What is justified belief? In G. Pappas (Ed.), *Justification and knowledge* (pp. 1–24). D. Reidel.
- Goldman, A. I. (2011). Toward a synthesis of reliabilism and evidentialism? or: Evidentialism's troubles, reliabilism's rescue package. In T. Dougherty (Ed.), *Evidentialism and its discontents* (pp. 254–280). Oxford University Press.
- Joyce, J. M. (1998). A nonpragmatic vindication of probabilism. *Philosophy of Science*, 65(4), 575–603.
- Kagan, S. (1992). The structure of normative ethics. *Philosophical Perspectives*, 6, 223–242.
- Leitgeb, H., & Pettigrew, R. (2010). An objective justification of Bayesianism I: Measuring inaccuracy. *Philosophy of Science*, 77(2), 201–235.
- McCain, K. (2014). *Evidentialism and epistemic justification*. Routledge.
- McCain, K. (2015). Explanationism: Defended on all sides. *Logos and Episteme*, 6(3), 333–349.
- Pettigrew, R. (2013). A new epistemic utility argument for the principal principle. *Episteme*, 10(1), 19–35.

Rawls, J. (1971). *A theory of justice*. Harvard University Press. 2nd Edition 1999.

Sylvan, K. L. (2018). Reliabilism without epistemic consequentialism. *Philosophy and Phenomenological Research*, 3, 525–555.

Sylvan, K. L. (2020). An epistemic nonconsequentialism. *Philosophical Review*, 129(1), 1–51.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.