Certainty in Action

Bob Beddor

Forthcoming in *The Philosophical Quarterly*

Abstract

When is it permissible to rely on a proposition in practical reasoning? Standard answers to this question face serious challenges. This paper uses these challenges to motivate a certainty norm of practical reasoning. This norm holds that one is permitted to rely on p in practical reasoning if and only if p is epistemically certain. After developing and defending this norm, I consider its broader implications. Taking a certainty norm seriously calls into question traditional assumptions about the importance of belief and knowledge. In particular, it raises the possibility that many epistemological jobs that are usually assigned to belief and knowledge should be reallocated to two related but importantly different states: psychological and epistemic certainty.

1 Introduction

We often take various propositions for granted. When leaving my house, I take for granted that my car will start, that I am a certain distance from my office, etc. Say that an agent *relies* on a proposition when they take it for granted in this way.

Whether one should rely on a proposition depends on its epistemic status. If I have good reason to think my car won't start, then I shouldn't take for granted that it will. But what epistemic status, exactly, is required for permissible reliance?

This question has provoked a lively debate. These days, the most popular answer is *knowledge*. Over the last decade, more and more philosophers have signed up for some version of a knowledge norm of practical reasoning:

Knowledge Norm (KN): When p is relevant to an agent A's decision, A is epistemically permitted to rely on p in practical reasoning iff A knows p.[1]

Despite—or perhaps because of—its popularity, KN has generated considerable backlash. A particularly serious challenge concerns its sufficiency direction. A number of epistemologists have offered cases in which an agent knows p, yet they ought not rely on p. And the

[1] While they differ on points of detail, Hawthorne (2004); Williamson (2005); Hawthorne and Stanley (2008); Weatherson (2012); Moss (2018); and Weisberg (2013) all express sympathy for a norm along these lines. For endorsements of its sufficiency direction, see Fantl and McGrath (2009); Ross and Schroeder (2014).

reason they ought not do so is that their epistemic position *vis-à-vis* p, while sufficient for knowledge, is still in some way deficient.[2]

Some epistemologists take these cases to undermine not only KN, but also the larger project of identifying some circumstance-invariant epistemic status that is both necessary and sufficient for permissible reliance. In place of this project, these authors propose what we might call a "variable status norm":

VARIABLE STATUS NORM (VN): When p is relevant to A's decision, A is epistemically permitted to rely on p in practical reasoning iff A's epistemic position vis-à-vis p is of the sort required by A's circumstances.

This paper starts by critically examining the state of play between KN and VN. §2 investigates strategies for overcoming the counterexamples to KN and argues that they prove unsuccessful. §3 considers VN and argues that its proponents fail to provide a plausible story about how an agent's circumstances determine the sort of epistemic status that is required, in those circumstances, for permissible reliance.

This leaves us at an impasse; the rest of the paper charts a path forward. I suggest we can make progress by taking *certainty* to play a central role in practical reasoning. On the view developed here, in order for someone to rely on p in practical reasoning, they must be psychologically certain that p. If their credence in p falls short of certainty, then they are not relying on p, but only on some related proposition—say, that p is extremely likely. It follows that in order for an agent to *permissibly* rely on p, it must be permissible for them to be psychologically certain that p. And this obtains, I argue, if and only if p is *epistemically* certain for them.

The idea that certainty is the norm of practical reasoning will strike many readers as absurd. It's commonly held that virtually none of our knowledge qualifies as epistemically certain. So a certainty norm would seem to entail that we are almost never permitted to rely on everyday premises (e.g., that my car will start) in practical reasoning. However, I argue that the same sort of considerations that cast doubt on skepticism about knowledge also cut against skepticism about certainty. A natural way of defusing the skeptical threat is to adopt a contextualist semantics for certainty claims: while in some contexts precious little qualifies as epistemically certain, in ordinary contexts much more of our knowledge makes the cut.

After developing and defending my account (§§4-7), I explore its consequences for the norms of assertion and belief (§8). In doing so, I aim to cast doubt on a familiar picture, presupposed in much of contemporary epistemology, and to sketch an alternative. The familiar picture goes like this: the primary cognitive state involved in practical reasoning and assertion is *belief*. And the primary epistemic benchmark for this state—its gold standard, if

[2] Brown (2008a,b, 2012b); Gerken (2011); Reed (2010); Lackey (2010); Locke (2015); Roeber (2018).

you will—is *knowledge*. The alternative goes like this: the primary cognitive state involved in practical reasoning and assertion is something more demanding than belief. It's *psychological certainty*. And the primary epistemic benchmark for this state is something more demanding than knowledge. It's *epistemic certainty*.

2 The knowledge norm and the sufficiency challenge

2.1 Motivations for KN

Why have so many been attracted to KN? Mainly because it promises to make sense of our intuitive judgments about when one should—or should not—rely on some premise in practical reasoning. Take the following example from Hawthorne and Stanley (2008: 571):

> **Restaurant** Hannah and Sarah are trying to find a restaurant, at which they have time-limited reservations. They could easily ask for directions. However, instead of asking, Hannah goes on her hunch that the restaurant is down a street on the left. After walking for some amount of time, it becomes quite clear that they went down the wrong street. A natural way for Sarah to point out that Hannah made the wrong decision is to say, "You shouldn't have gone down that street, since you didn't know the restaurant was here."

Sarah's criticism seems perfectly natural. KN explains why. Since Hannah didn't know the restaurant was to the left, she shouldn't have taken for granted that it was.

However, as Brown (2008a) notes, this sort of case only motivates the necessity direction of KN:

K-Nec: When p is relevant to A's decision, A is permitted to rely on p in practical reasoning only if A knows p.

Restaurant provides no direct support for KN's sufficiency direction:

K-Suff: When p is relevant to A's decision, then if A knows p, A is permitted to rely on p in practical reasoning.

Morover, this sufficiency direction gives rise to various challenges, to which I now turn.[3]

[3]Of course, K-Nec also faces challenges. Schiffer (2007) objects that it is at odds with Bayesian decision theory, which recommends acting on the basis of utilities and probabilities. And Brown (2008b), Neta (2009), and Gerken (2011) object that K-Nec conflicts with our intuitions about the permissibility of relying on a rational belief that does not amount to knowledge, either due to falsity or Gettierization.

However, proponents of K-Nec have some promising responses to these objections. In response to the first difficulty, Hawthorne and Stanley (2008) point out that K-Nec permits agents to rely on probabilistic premises

2.2 Challenges to K-Suff

A first difficulty for K-Suff comes from the fact that we sometimes criticize people for acting on the basis of propositions that they don't know with *certainty* (Brown 2008a,b, 2012b; Gerken 2011). To give a variant of one of Hawthorne and Stanley's examples, we might criticize a negligent doctor by saying, "You shouldn't have used that needle; it wasn't certain that the needle was sterilized." Similarly, Brown (2008a: 171) offers the example of a parent remonstrating their teenager who left a party after midnight: "You shouldn't have left so late. You didn't know for sure there would be a bus at that time."

It's plausible that knowing with certainty is a more demanding epistemic state than merely knowing.[4] If this is right, then K-Suff does not explain why these criticisms seem just as legitimate as their knowledge-involving kin.

A related difficulty for K-Suff comes from cases where one knows p, yet intuitively one's epistemic position *vis-à-vis* p is not strong enough to warrant taking p for granted. Here's an example from Reed (2010: 228-229):

> **Jellybean** Hugo is participating in a study where the researcher asks him questions about Roman History, a subject with which Hugo is very familiar. For every correct answer that Hugo gives, Hugo gets a jellybean. For every incorrect answer, he receives a painful shock. He can also remain silent, which will result in neither jellybeans nor shocks. The first question is whether Caesar was born in 100 BC.

As Reed develops the case, Hugo knows that Caesar was born in 100 BC. But he doesn't know it with absolute certainty. By K-Suff, Hugo is permitted to rely on the proposition that Caesar was born in 100 BC. Now, what can we conclude about what Hugo is permitted to *do*? Here's a plausible bridge principle connecting permissible reliance with permissible action:

Reliance-Action Bridge: If A is permitted to rely on p, then A is permitted to act as if p—that is, A is permitted to perform whatever action maximizes expected utility conditional on p.[5]

provided these premises are themselves known. This has invited a charge of over-intellectualization, but see Moss (2013, 2018) and Weisberg (2013), who argue that credences can directly constitute probabilistic knowledge. In response to the second difficulty, advocates of K-Nec observe that agents can be blameless for doing that which is forbidden, provided that they have good reason to think that their conduct is permissible. For versions of this idea, see Sutton (2007); Hawthorne and Stanley (2008). For further discussion of the relation between permissibility and excusability, see Kelp and Simion (2017); Williamson (forthcoming); Greco (forthcoming). There is much more to be said on both of these issues, and a full discussion would require a paper in its own right. The present point is simply that there are potentially viable strategies for defending K-Nec. For this reason, this paper will focus on the challenges to K-Suff, which prove more recalcitrant.

[4] Some might question this. For those who do, I beg for your patience: I tackle this concern at length in §6.1.

[5] For endorsements of Reliance-Action Bridge, Fantl and McGrath (2002, 2009); Ross and Schroeder (2014).

Given this bridge principle, it follows that Hugo is permitted to answer the question ("Yes!"). This conflicts with the intuition that Hugo should keep mum.[6]

The standard response to cases like this is to appeal to *impurism* about knowledge.[7] According to impurism, practical factors affect knowledge: the higher the stakes, the more justification one needs in order to know. According to this diagnosis, the relatively high stakes of Hugo's predicament preclude him from knowing that Caesar was born in 100 BC.

There are, however, two problems with this impurist defense of K-Suff. First, it's not clear that the high stakes do deprive Hugo of his knowledge. Imagine Hugo is later asked, "Why didn't you answer the question? Didn't you know the answer?" We can imagine him replying: "I'll admit I knew the answer. But I wasn't absolutely sure, and I really didn't want to get shocked" (Reed 2010: 229).

Impurists may push back on this way of fleshing out the case. As Fantl and McGrath observe, we can also imagine Hugo explaining his silence by saying, "I was very confident of the answer, but I didn't *know* it" (Fantl and McGrath 2009: 62-63). But this observation alone doesn't solve the problem. After all, knowledge ascriptions are typically vague, both in terms of how much confidence and how much justification is required for knowledge. In underdescribed cases we will find it coherent to resolve the vagueness either way—either by saying the person knows, or by saying they don't. The problem for the impurist defense is that when we resolve the vagueness by saying the person knows, we still find it sensible to deny that they should act on the basis of this knowledge.

Even if defenders of K-Suff remain unmoved by this worry, they face a second—and, in my eyes, much more decisive—problem. The problem is that we can concoct counterexamples to K-Suff that do not involve high stakes. Consider the following scenario, adapted from Roeber (2018: 177-178):

> **Survey** A social scientist approaches you and asks you to participate in an experiment. She explains: "I will give you a survey containing 100 questions. Each contains a pair of propositions, at least one of which is true. Your task will be to select a true proposition from each pair of propositions. You can't select both. If you do, I'll mark the question wrong. If you get at least half of the questions right, I'll give you this keychain." She shows you the keychain, which has a bottle opener, so you want it. As you fill out the survey, you get to the following pairs of propositions:
>
> 26a) Boethius wrote *The Consolations of Philosophy*.
>
> 26b) Either 1=1 or Boethius wrote *The Consolations of Philosophy*.

[6] For a structurally similar counterexample to K-Suff, see Brown's surgeon scenario (Brown 2008a: 176).

[7] AKA "pragmatic encroachment." See, a.o., Fantl and McGrath (2002, 2009); Stanley (2005); Hawthorne and Stanley (2008); Weatherson (2012); Ross and Schroeder (2014).

On a natural way of fleshing out the case, you know both 26a) and 26b). By K-Suff, it would be permissible for you to rely on either proposition. By Reliance-Action Bridge, you are permitted to select either proposition by way of response. But, Roeber argues, this seems wrong: intuitively, you should *not* select 26a). After all, 26a) entails 26b), so any reason you have for selecting the first also counts in favor of selecting the second. Furthermore, you have an additional reason for selecting 26b): it is a necessary truth, hence there's no possibility—however remote—that you could go wrong in selecting it.

Here the impurist defense provides no help. After all, the stakes are very low: a mere bottle opener hangs in the balance. So the impurist must admit that you retain your knowledge of 26a).

3 Variable status

A number of epistemologists take these considerations to cast doubt on the very idea that a single epistemic status is both necessary and sufficient for permissible reliance. In place of this idea, they embrace what I've called the Variable Status Norm (VN). According to this norm, it is permissible for an agent to rely on p just in case their epistemic position *vis-à-vis* p is of the sort required by their circumstances.

At first glance, VN seems to neatly sidestep the counterexamples to K-Suff. In both **Jellybean** and **Survey**, can't we just say that the agent's circumstances require a stronger epistemic position than knowledge? But it's one thing to say this, and another to explain it. What features of the agent's circumstances demand something stronger than knowledge?

More generally, VN faces the following question:

Status Determination Question: How do an agent's circumstances determine which epistemic relation they need to bear to p in order for them to permissibly rely on p in those circumstances?

In the rest of this section, I consider two initially promising answers. Adapting some terminology from Anderson and Hawthorne (2019), I'll call these *the stakes approach* and the *practical adequacy approach*. I argue that neither succeeds. It remains to be seen whether some alternative version of VN fares better.[8]

The stakes approach starts with the familiar idea that we can order epistemic statuses by strength. At the top of the scale we might have *absolute epistemic certainty*. A step down

[8] Anderson and Hawthorne also raise difficulties for both approaches. But their target is somewhat different: they formulate the stakes approach and the practical adequacy approach as impurist theories of knowledge. This makes a difference to the nature of our criticisms. Anderson and Hawthorne object that both the stakes approach and the practical adequacy approach have implausible consequences about whether knowledge is gained or lost in various situations. But these objections do not undercut either approach *qua* answer to the Status Determination Question. After all, VN denies that knowledge is always tied to permissible reliance.

we might have *knowledge*. Another step down we might have some weaker property, such as *being very likely, given the evidence*. Perhaps the stakes of a decision determine which of these statuses is required for permissible reliance. The higher the stakes, the stronger the requisite status.

This handles **Jellybean**. After all, Hugo is in a high stakes situation. According to the stake approach, this means that in order for Hugo to permissibly rely on the proposition, *Caesar was born in 100BC*, it's not enough for him to know it. Rather, he needs to know it with certainty. However, this proposal flounders on **Survey**, and for precisely the same reason that the impurist defense did. In **Survey** the stakes are quite low: a measly keychain is on the line. Yet you still shouldn't rely on 26a).

This problem can be used to motivate an alternative approach to the Status Determination Question. Say that p is *practically adequate* for you if and only if any difference between your actual epistemic position *vis-à-vis* p and the best possible epistemic position *vis-à-vis* p does not make a difference to which actions you should perform. (Equivalently: the action you actually prefer, given your epistemic position, is the same as the action you prefer conditional on p.) We could then propose that whenever your decision depends on p, you are permitted to rely on p if and only if p is practically adequate for you.[9]

This handles **Survey**. If you had the best possible epistemic position *vis-à-vis* 26a), then you would be entitled to choose either 26a) or 26b). So the difference between your epistemic position *vis-à-vis* 26a) and the best possible epistemic position *does* affect which actions you should perform. Hence 26a) is not practically adequate for you.

However, the practical adequacy approach faces a problem of its own. The problem is that a proposition can be practically adequate even though it isn't known, or true, or justifiably believed. Indeed, it can be practically adequate even though the agent has good reason to disbelieve it. Suppose Mindy has carefully consulted the weather forecasts. Given the evidence, it's rational for her to assign a .01 credence to the proposition: *It will snow tomorrow* (SNOW). But now suppose Mindy is offered a bet where she wins $1,000,000 if it snows tomorrow and loses a penny if it doesn't. The action she has most reason to perform is taking the bet, which is also the action she would have most reason to perform if she stood in the best possible epistemic relation to SNOW. And so SNOW is practically adequate for her. But, intuitively, she shouldn't rely on this proposition in practical reasoning. While she should take the bet, she shouldn't treat SNOW as her basis for doing so. Rather, she should take the bet on the basis of the .01 probability that SNOW is true, together with the facts about the payoffs.

As a corollary of this difficulty, the practical adequacy approach has trouble explaining our ordinary patterns of criticism. We saw in §2 that we often criticize people for acting on the basis of propositions that they don't know, or even that they don't know with certainty. The practical adequacy approach does not explain why we find these criticisms so natural.

[9] See Locke (2015) for a defense of this sort of view.

Taking stock: defenders of VN face the Status Determination Question. Two initially promising answers—the stakes approach and the practical adequacy approach—quickly ran into trouble. Of course, this does not show that there is *no* way of answering the Status Determination Question; nothing in this section amounts to an impossibility proof against VN. Still, I think that the considerations put forward here highlight why it is not easy to develop VN into a systematic theory that captures our pretheoretic intuitions about permissible reliance.

Could advocates of VN refuse to answer to the Status Determination Question? Perhaps, some might suggest, they could embrace a "particularist" view, according to which there is no general story to be told about how practical factors select the epistemic status required for permissible reliance.[10] While this is admittedly an option, it strikes me as a rather disappointing one, which should only be embraced as a last resort.

Luckily there is no need to go this route. The rest of this paper develops a new account of permissible reliance. My account agrees with KN that there is a single epistemic status that is both necessary and sufficient for permissible reliance. But it takes this status to be something more demanding than knowledge: it is *epistemic certainty*.

4 Certainty as the norm of practical reasoning

4.1 On certainty

When epistemologists talk about certainty, they typically have in mind psychological or subjective certainty. For someone to be psychologically certain of p is for them to have the highest possible degree of confidence in p. Being certain of p in this sense precludes harboring any doubts about p's truth. But it does not preclude falsity or irrationality.

However, there is also an *epistemic* sense of certainty—a sense that is most naturally conveyed via the impersonal construction, "It is certain that p."[11] How does epistemic certainty relate to psychological certainty? Plausibly, the connection is normative: *epistemic certainty is the epistemic state that warrants psychological certainty.* Here is one way of cashing this out:

Normative Link: p is epistemically certain for A iff A's epistemic position renders it permissible for A to be psychologically certain that p.[12]

[10] Cf. Fantl and McGrath (2019) who raise (without endorsing) the possibility of a particularist version of impurism.

[11] See Moore (1959); Stanley (2008); DeRose (2009); Beddor (forthcoming).

[12] The locution, "p is epistemically certain for A" is admittedly a bit awkward, and does not crop up in ordinary language. However, I think it latches onto a pretheoretic notion—one that we might convey using locutions such as, "In view of A's epistemic position, it is certain that p." On a natural interpretation, this means that A's epistemic position makes it epistemically certain that p. (In this regard, "It is certain that" resembles modals

Normative Link holds considerable appeal. It explains why it would be odd for you to affirm that you are certain of p, while denying that it is certain that p. According to Normative Link, doing so would be tantamount to admitting that your epistemic position does not permit your level of confidence in p.

While there is much more to be said about certainty, this minimal sketch is enough to introduce the norm of practical reasoning I wish to defend.

4.2 Reliance and certainty

Let us return to the question of when it is permissible to rely on a premise in practical reasoning. But let us tackle it from a different angle. What does it mean to *rely* on some premise in the first place?

Start by considering a stock example of practical reasoning:

(P1) If the store is open, then I better stop by.
(P2) The store is open.
(Ccl) I better stop by the store.

As we noted in §1, this reasoning pattern is importantly different from reasoning on the basis of probabilistic premises, e.g.:

(P1) If the store is open, then I better stop by.
(P2′) It's *very likely* that the store is open.
(Ccl) I better stop by the store.

The difference between these two forms of reasoning persists if we replace (P2′) with a premise invoking arbitrarily high probabilities, e.g.:

(P2″) It's 99.99% likely that the store is open.

Here too, the psychological process of relying on (P2) seems to be importantly different from the psychological process of relying on its probabilistic cousin, (P2″).

This suggests the following generalization: if you harbor any doubts about p while engaged in practical reasoning, you are not genuinely relying on p. Instead, you are relying on a closely related premise—say, that p is very likely true, or almost certainly true, etc.

This generalization provides a helpful clue about what sort of psychological state reliance involves. We've already noted that being psychologically certain of p requires being free from doubts about p. Perhaps, then, reliance requires psychological certainty:

such as "must". It is awkward to relativize modals to specific agents via preposition phrases. For example, it sounds odd to say, "It must, for Fred, be raining." But we can relativize modals to agents using "in view of" phrases. For example: "In view of Fred's information, it must be raining." For more on the interaction between "in view of" phrases and modals, see Kratzer (1977).)

RELIANCE REQUIRES CERTAINTY: If A relies on p in practical reasoning, then A is psychologically certain that p.

Reliance Requires Certainty is a descriptive claim. But it has normative implications. If reliance requires psychological certainty, then reliance is epistemically permissible only if psychological certainty is epistemically permissible. Assuming that there are no further epistemic requirements on reliance other than those imposed by the certainty requirement, we can convert this into a biconditional:

PERMISSIBLE RELIANCE-PERMISSIBLE CERTAINTY: If A is facing a decision that depends on p, then it is epistemically permissible for A to rely on p in practical reasoning iff it is epistemically permissible for A to be psychologically certain that p.

According to Normative Link, it is permissible to be psychologically certain that p just in case p is epistemically certain. This gives us:

EPISTEMIC CERTAINTY NORM (ECN): If A is facing a decision that depends on p, then it is epistemically permissible for A to rely on p in practical reasoning iff p is epistemically certain for A.

I have provided an initial argument for an epistemic certainty norm of practical reasoning. This argument was just a warm-up: the main reason to adopt ECN is that it avoids the problems facing rival norms. But before delving into a detailed comparison with its rivals, let me first tackle an obvious objection to ECN. The objection is that ECN is too demanding, since virtually nothing qualifies as epistemically certain.

5 The skeptical objection

5.1 The argument for skepticism about certainty

The idea that virtually none of our knowledge rises to the level of epistemic certainty is often assumed, but rarely defended. The main exception is Unger (1975), who provides a sustained defense of this skeptical conclusion.

Unger's argument starts from the premise that "certain" is an *absolute term*, or, in more contemporary semantic terminology, a *maximum-standard absolute gradable adjective*. Some background: gradable adjectives denote properties that come in degrees. Within the class of gradable adjectives, maximum-standard adjectives are a unified subclass: they only apply to an entity if it has the *maximum* degree of the denoted property.[13]

The main test for a maximum-standard gradable adjective is to see whether we can apply it to some entity x while, in the same breath, insisting that x could have more of the relevant property. Compare:

[13] See Unger (1975); Kennedy (2007).

(1) ? The table is flat, but it could be flatter.

(2) The car is fast, but it could be faster.

(1) seems infelicitous—more so, at any rate, than (2). This is evidence that "flat" is a maximum-standard gradable adjective, whereas "fast" is not. If we apply this test to "certain", we find that it behaves much like "flat":

(3) ? It is certain that the car will start, but it could be more certain.[14]

But if "certain" is a maximum-standard adjective, then a proposition only qualifies as epistemically certain if it has the maximum degree of epistemic certainty. But, as Unger points out, this seems to be a very high bar. And it seems that only a handful of propositions will measure up—perhaps logical truths and the cogito, but little else.

5.2 Resisting the argument

At first glance, Unger's argument seems compelling. But I think we should be wary of its skeptical conclusion. In ordinary contexts, we not only claim to know many things. We also claim that many things are (epistemically) certain:

(4) It's certain that the car will start.

(5) It's certain that China is larger than Singapore.

Certainty skepticism thus has the unwelcome consequence that virtually all of these ordinary epistemic certainty claims are false.

More worrisome still, Unger's argument generalizes beyond certainty. It could be used to show that an ordinary claim such as:

(6) The pavement is flat.

is false, since we can find things—say, a desk—that are even flatter than the pavement. Unger infamously embraced this consequence (1975: 211-212). But at this point it becomes tempting to perform *modus tollens* rather than *modus ponens*. Even those who are willing to stomach skepticism about certainty might hesitate to consign such a large swath of ordinary discourse to error.

But where exactly does Unger's argument go wrong? Lewis offers a plausible diagnosis:

> The right response to Unger, I suggest, is that he is changing the score on you. When he says that the desk is flatter than the pavement, what he says is acceptable only under raised standards of precision. Under the original standards

[14] Psychological uses of "certain" behave much the same: "Lena is certain the car will start, but she could be more certain" seems equally infelicitous.

> the bumps on the pavement were too small to be relevant either to the question whether the pavement is flat or to the question whether the pavement is flatter than the desk… [In the case of certainty] a parallel response is in order. (1979: 353-354)

There are various ways of developing this contextualist response. For our purposes, we need not commit ourselves to any particular implementation. However, I'll briefly sketch one option as a "proof of concept."

5.3 A sketch of a contextualist semantics

According to a rich tradition in epistemic logic, a variety of epistemic vocabulary can be analyzed in modal terms.[15] While this tradition typically focuses on knowledge ascriptions, it can easily be extended to epistemic certainty claims. As a first pass, say that p is maximally epistemically certain for A iff p obtains in all of A's "epistemic possibilities": all of the worlds compatible with A's epistemic state.

To inject a contextualist element into this semantics, we can take our cue once again from Lewis. Lewis (1996) notes that context frequently restricts our domain of quantification: we frequently say things like "All the beer is gone" even when there is some unconsumed beer out yonder. Lewis uses the phenomenon of contextual domain restriction as a model for a contextualist treatment of knowledge ascriptions. I propose doing much the same for epistemic certainty claims:

Contextualist Truth Conditions "It is certain that p" is true, relative to a context c, iff "It is maximally certain that p" is true in c, which in turn obtains iff p holds in all of the speaker's c-relevant epistemic possibilities.[16]

This preserves Unger's insight that "certain" functions as a maximum-standard adjective. In any context, "It is certain that p" entails "It is maximally certain that p", thereby explaining the infelicity of (3). At the same time, we block the skeptical conclusion that virtually all of our everyday epistemic certainty claims are false.

How does context determine which possibilities are relevant? Here too we could turn to contextualists about knowledge ascriptions for guidance. For example, contextualists frequently invoke *salience* when explaining the mechanisms of context domain restriction. Here's one way of cashing this out:

[15] This tradition traces back to Hintikka (1962).

[16] We could also formulate parallel truth conditions for psychological certainty ascriptions. Let A's "psychological possibilities" be the worlds that A takes to be candidates for reality. (Roughly, we can think of these as the worlds to which A assigns some positive credence.) Then an ascription, "A is psychologically certain that p" is true relative to a context c iff p holds in all of A's c-relevant psychological possibilities. For further discussion and development of this sort of contextualist view, see Beddor (forthcoming).

SALIENCE CONSTRAINT *Ceteris paribus*, if a possibility w is made salient in a context c, then w is c-relevant.[17]

While this is surely not the complete story about how contextual domain restriction works, a constraint along these lines has considerable appeal. In particular, it explains why explicitly introducing error possibilities tends to raise the contextual standards. For example, it explains why explicitly thinking about dream scenarios and brains in vats makes us less inclined to assent to (4) and (5).

5.4 Contextualizing the epistemic certainty norm

By going contextualist, we defuse the main objection to ECN. Of course, if "certain" is context-sensitive, our formulation of ECN will need to be tweaked. Which context is relevant for permissible reliance? The most obvious answer is: *the decision-maker's*. That is:

CONTEXTUALIZED ECN: Suppose A is making a decision that depends on p, in a context c_A. Then it is epistemically permissible for A to rely on p if and only if p is epistemically certain for them, relative to the standards of c_A.[18]

To illustrate, consider our earlier store-related reasoning. In an ordinary context, (P2) (*The store is open*) might qualify as maximally epistemically certain for you. By Contextualized ECN, you are permitted to rely on it in your context. But now suppose that as you drive store-ward, you consider the question, "Could the store have closed early?" Attending to this possibility will, *ceteris paribus*, make this possibility relevant (courtesy of the Salience Constraint). Assuming you have no way of ruling out this possibility, (P2) will cease to qualify as maximally epistemically certain, and you will no longer be entitled to rely on this premise.

If context can impact permissible reliance, then it can also impact which actions you are permitted to perform (by Reliance-Action Bridge). Some might find this worrisome: even if (P2) ceases to qualify as certain, relative to your context, aren't you still rationally permitted to perform the same actions you were a moment before? For example, isn't it rational to continue driving store-ward?

In response, we should start by noting that even once a particular proposition fails to qualify as epistemically certain, a number of other propositions will retain this status. Presumably, even once (P2) ceases to be certain, the probabilistic premise (P2′) (*It's very likely*

[17] See the 'Rule of Attention' in Lewis (1996). See also Blome-Tillman (2009).

[18] Cf. DeRose (2009), who advocates a contextualized version of a knowledge norm of assertion, and discusses—though does not endorse—a contextualized version of a knowledge norm of practical reasoning. See also Stanley (2008), who advocates a similarly contextualized certainty norm of assertion. (I discuss Stanley's proposal in §8.)

that the store is open) will remain certain. In many circumstances, this surrogate will rationalize the very same actions.[19]

Of course, we can imagine versions of the case where the two premises do not rationalize the same actions. Suppose that conditional on the store being open, it maximizes expected utility to drive directly there. But conditional on it only being very likely that the store is open, it maximizes expected utility to put in a quick phone call checking the store's hours. Given Reliance-Action Bridge, our view predicts that once (P2) ceases to be certain, you should make the call. This seems like the correct verdict. So Contextualized ECN only predicts that context impacts permissible action in ways that align with our pretheoretic judgments.

Having dealt with the most pressing objection to ECN, we are now in a position to compare it with its main rivals.

6 Comparison with the knowledge norm

In order to compare ECN to KN, we first need to investigate the relationship between epistemic certainty and knowledge. I start (§6.1) by arguing that epistemic certainty is distinct from knowledge. While epistemic certainty entails knowledge, knowledge does not entail epistemic certainty. If this is right, then ECN is well-poised to handle the challenges to K-Suff (§6.2).

6.1 The relation between epistemic certainty and knowledge

How is epistemic certainty related to knowledge? Plausibly, epistemic certainty entails knowledge, or at least being in a position to know. More precisely:

EC⇒K: If "p is epistemically certain for A" is true in a context c, then "A knows (or is in a position to know) p" is also true in c.

In support of this entailment, recall our earlier arguments that "certain" is a maximum-standard adjective: p is only epistemically certain for A if A stands in a maximally strong epistemic relation to p. If A stands in a *maximally* strong relation to p, surely that is good enough for knowledge!

[19]Two points of clarification. First, the surrogate proposition doesn't have to be probabilistic. In some circumstances, the surrogate for p will be a proposition such as p *is approximately true*. For example, Gao (2017) discusses a scientist who knows that Newtonian mechanics is false, but nonetheless uses Newton's laws when making calculations. On the account offered here, the scientist is relying on the approximate truth of Newtonian mechanics. Second, even when the surrogate is probabilistic, we need not assume that the agent consciously contemplates the probabilities. Adapting a maneuver that has been deployed in defense of K-Nec, we could hold that credences themselves can qualify as epistemically certain, assuming they satisfy various epistemic desiderata (e.g., safety, sensitivity, and the like). See the references in fn. 3 for relevant discussion.

One consequence of EC⇒K is that epistemic certainty claims are factive. This consequence seems to be borne out by the data. Suppose Agnes says, "It's certain the show starts at 8." As a matter of fact, the show does not start until 8:30. Even if the assertion was made in a low standards context, and even if Agnes had excellent grounds for thinking the show would start at 8, it would be natural for her to retract her claim.

Some might worry that a factive condition on epistemic certainty is in tension with Normative Link. If falsehoods are never epistemically certain, then, by Normative Link, one should never be psychologically certain of a falsehood. But what if one has excellent (misleading) evidence for a falsehood? In response, one option is to borrow a maneuver from defenses of K-Nec and distinguish between *impermissibility* and *criticizability*. According to this response, it is always impermissible to be psychologically certain of a falsehood. Nonetheless, one is not thereby criticizable, since one might have good reason to think one is complying with Normative Link.

Suppose, then, that we accept EC⇒K. Does the converse entailment hold?

K⇒EC: If "A knows p" is true in a context c, then "p is epistemically certain for A" is also true in c.

At first blush, K⇒EC may seem plausible. An initial point in its favor is that it explains why it is infelicitous to claim:

(7) ? I know the show starts at 8, but it's not certain that it does.[20]

K⇒EC can also be motivated on theoretical grounds. Suppose one accepts that evidence is knowledge (E = K).[21] It's also plausible that there is a close connection between epistemic certainty and evidence: if A's evidence entails p, then p is epistemically certain for A. Put these two theses together, and we get a straightforward argument that knowledge entails epistemic certainty.

However, I think we should resist the initial appeal of K⇒EC. Here I'll mention three pieces of evidence that count against this entailment.[22]

The first piece of evidence comes from cases where it's natural to claim an agent knows some proposition, but it seems unnatural—in fact, downright false—to claim that this proposition is known with certainty. Consider the unconfident examinee (Radford 1966). The time has come for his oral history exam. Remiss student that he is, he has not studied, and he doubts his ability to recall the material. Nonetheless his answers are invariably correct. As a number of authors have noted, it would be natural for his examiner to remark, "How surprising! I guess he knew the answers after all!" However, it would be much less natural to remark, "He knew the answers with certainty" or "He was certain of the answers."[23]

[20] Though stay tuned for §8, where I present an alternative explanation of this incoherence.

[21] The *locus classicus* of this view is Williamson (2000).

[22] See Beddor (forthcoming) for further discussion of these issues.

[23] See, e.g., Armstrong (1969); Stanley (2008); McGlynn (2014).

Some may retort that this just shows that knowledge does not require *psychological* certainty. It may still require epistemic certainty. However, the unconfident examinee might well be *rational* in harboring doubts about his answers. If his epistemic position *vis-à-vis* the answers does not permit psychological certainty, then, by Normative Link, these answers are not epistemically certain for him.

The second piece of evidence comes from the fact that the expression, "knows for certain" is not redundant. Compare:

(8) Lena knows for certain [/with certainty] that the car will start.

(9) Lena knows that the car will start.

A speaker who utters (8) makes a stronger claim than a speaker who utters (9). In particular, there's a natural reading of (8) on which the speaker is communicating that Lena stands in an unusually strong epistemic position *vis-à-vis* the claim that the car will start—a position that is stronger than that which is usually required for knowledge.

Defenders of K⇒EC might try to explain this observation on pragmatic grounds. Perhaps, they might suggest, "knows" does indeed entail "certain" (in at least the epistemic sense), but combining these two expressions raises the contextual standards governing both. However, a problem for this pragmatic explanation is that it does not generalize to other pairs of context-sensitive expressions, one of which entails the other. Consider pairs such as "wealthy" and "affluent", "probable" and "likely", "sad" and "unhappy". Try combining them together—"She's wealthy and affluent", "He's sad and unhappy"—and you'll get a sentence that sounds awfully redundant. This contrasts with (8), which does not generate any similar sense of redundancy.

A further difficulty for this pragmatic explanation—and a third piece of data against K⇒EC—comes from the coherence of claiming someone knows something with *near* or *close to* certainty. Here are some naturally occurring examples:

(10) For the first time since Vietnam, each of Barone's boxers knows with near certainty he is headed off to war.[24]

(11) So I think we know with close to certainty that he'll never be a star again.[25]

Taken together, these three pieces of data make a strong cumulative case against K⇒EC.[26] They thereby motivate the idea that epistemic certainty is more demanding than knowledge.

[24] https://www.nytimes.com/slideshow/2010/04/08/sports/20100409_BOXING_index/s/0409boxing_6.html

[25] https://forums.realgm.com/boards/viewtopic.php?f=6&t=1882971&start=280

[26] It worth noting that these arguments also put pressure on E=K. Suppose we grant that if your evidence entails p, then p is epistemically certain for you (E⇒EC). If K⇏EC, it follows that not all knowledge is evidence. For the purposes of this paper, I will remain noncommittal regarding what lessons about evidence we should draw from this. However, one natural option—developed in Beddor (forthcoming)—would we be to replace E=K with E=EC: an agent's evidence, relative to a context, is whatever propositions qualify as epistemically certain for them, relative to that context.

But it is one thing to establish this; it is another to show that this allows ECN to avoid the objections to K-Suff. I now turn to this task.

6.2 Avoiding the problems for K-Suff

The first problem for K-Suff was that we criticize agents for acting on the basis of uncertain propositions. For example, we noted that it would be natural to complain to a doctor, "You shouldn't have used that needle; it wasn't certain that the needle was sterilized." ECN explains these criticisms. On a natural way of filling out the case, the doctor should only have used the needle if she were in a position to permissibly rely on the premise, *The needle is sterilized*. According to ECN, she can permissibly rely on this premise only if it is epistemically certain for her. By contrast, KN cannot provide a similar explanation, if—as I've just argued—knowledge does not entail epistemic certainty.

Before moving on, let me address two worries about ECN's explanation here. First, some might worry that the explanation breaks down when the standards for certainty in the speaker's context are higher than the standards for certainty in the doctor's context. In such cases, Contextualized ECN would seem to predict that a speaker could say:

(12) ? You acted appropriately in using that needle, even though it wasn't certain that the needle was sterilized.[27]

This is a general problem for any view that ties permissible reliance to some context-sensitive notion. Thankfully, by now some plausible defenses of contextualist norms have been proposed. Perhaps the most detailed is due to DeRose (2009: 244-251). DeRose observes that the contextual standards will often be a function of the conversational purpose. In the case of (12), the conversational purpose is to make a normative evaluation of the doctor's action. By Contextualized ECN, the contextual standards that are most relevant to making this normative evaluation are those in the doctor's context. So it would be natural for the speaker to defer to the doctor's contextual standards when making such normative evaluations. But if the speaker does defer, then (12) is guaranteed to come out false.

The second worry is that ECN will have trouble capturing the full range of ordinary criticisms. In many contexts, it would be equally legitimate to complain, "You shouldn't have used that needle; it wasn't *completely* certain that the needle was sterilized." However, ECN is well-positioned to explain this. Recall that "certain" is a maximum-standard adjective. This predicts that "It is certain that p" is equivalent to "It is completely certain that p." Both require that p possesses the maximum degree of certainty, relative to the context.

Our next challenge to K-Suff came from cases where someone knows a proposition but cannot permissibly act on it. Start with **Jellybean**. As we noted in §2.2, **Jellybean** involves

[27] This worry traces to Hawthorne (2004), who raises it for a contextualized knowledge norm.

high stakes. Plausibly, practical consequences can affect the standards for epistemic certainty in a context: the graver the consequences of being wrong, the higher the standards. And so the proposition, *Caesar was born in 100 BC*, does not qualify as epistemically certain for Hugo, relative to his context.

This diagnosis resembles the impurist defense of K-Suff, which maintained that practical factors affect knowledge. Despite this resemblance, there are three reasons to think ECN has the upper hand.

First, it is independently plausible that practical factors can influence the standards governing maximum-standard gradable adjectives. For example, when deciding whether an antenna counts as "straight", we will use higher standards if we are building a satellite than if we are repairing a television. Why is this? Because small dents in the antenna are more likely to have serious consequences in satellite construction than in television repair. Similarly, whether we count a knife as "clean" depends on whether we are using it for chopping vegetables or for surgery (cf. Rotstein and Winter 2004). Thus the idea that the extension of epistemic uses of "certain" depends on practical factors follows from a more general property of gradable expressions.

No similar justification can be provided for thinking that the extension of "knows" depends on practical factors, since knowledge ascriptions are not gradable (Dretske 1981; Stanley 2004; Pavese 2017). Compare:

(13) It's somewhat/very/completely certain she will win.

(14) # I somewhat/very/completely know she will win.

On to the second reason for preferring ECN to the impurist defense. In §2.2 we noted that it would be coherent for Hugo to say, "I'll admit I knew the answer. But I wasn't absolutely sure, and I really didn't want to get shocked." This causes trouble for the impurist defense, since impurists maintain that the high stakes deprived Hugo of knowledge. By contrast, ECN accounts for the coherence of Hugo's explanation.

Here too, we might wonder if the problem re-emerges for ECN in a different guise. Is it coherent to claim that an agent is rationally certain of p, but they are still not in a position to rely on p? To test this, imagine Hugo had instead explained his behavior as follows:

(15) ? I was certain that Caesar was born in 100BC, and rationally so. But I really didn't want to risk getting shocked.

To my ears, (15) sounds odd—worse, at any rate, than the analogous explanation involving "knows." This suggests that the arguments against K-Suff do *not* generalize to undermine the sufficiency direction of ECN.

The third mark in favor of ECN comes from **Survey**. Recall that you know both propositions on the survey:

26a) Boethius wrote *The Consolations of Philosophy.*

26b) Either 1=1 or Boethius wrote *The Consolations of Philosophy.*

Nonetheless, it seems you are only permitted to select 26b).

ECN explains why. Of course, here we cannot appeal to stakes to explain why 26a) does not qualify as epistemically certain for you (relative to your context). But we can give an independent explanation—an explanation that once again follows from general features of maximum-standard adjectives.

Recall how Unger tried to convince us that the pavement isn't flat. He invited us to compare the pavement with a smooth desk. This led us to notice that the pavement's surface has more bumps than the desk's. And this led us to admit, "The desk is flatter than the pavement", which led us to deny, "The pavement is flat." As we saw in §5.2, Lewis rightly accused Unger of shifting the conversational context. For present purposes, what is important is the way Unger pulled this off. For any maximum-standard adjective *M*, considering the question, "Is x more *M* than y?" will typically lead us consider whether there are any *M*-related differences between x and y. This makes any such differences salient, which tends to raise the contextual standards so that only x or y falls under the extension of *M*.

The setup of **Survey** invites a comparative judgment of precisely this form. Since you can only select either 26a) or 26b), you are confronted with the question, "Is one of these propositions more certain than the other?" This leads you to compare your epistemic position *vis-à-vis* 26a) with your epistemic position *vis-à-vis* 26b). And this comparison serves to raise the contextual standards so that "26a) is certain" ceases to be true (relative to your context of deliberation). By ECN, you cannot rely on 26a) in **Survey**.

Thus ECN offers a plausible treatment of **Survey**. Is there any way for KN to co-opt this treatment? I think not. ECN's treatment of the case exploited the way comparisons affect the interpretation of maximum-standard gradable adjectives. But, as we just saw, "knows" is not gradable. Hence no parallel diagnosis is available to KN.

For this reason, I think the best option for knowledge normers is to retreat to the idea that knowledge entails epistemic certainty. This would enable them to say that even though "knows" is not itself gradable, a knowledge ascription is false whenever the corresponding epistemic certainty claim is false. And they could then help themselves to my explanation of why "26a) is certain" is false in **Survey**. But we've already anticipated the main problem for this move (§6.1): there is good reason to reject K⇒EC. Indeed, I think that **Survey** provides additional reason to reject this principle. After all, it seems implausible that as soon as you encounter 26b) as an option, you no longer count as "knowing" 26a).[28]

[28] Here I agree with Roeber (2018: 179).

7 Comparison with the variable status norm

How does ECN stack up against the Variable Status Norm (VN)? Earlier we saw that VN confronts the Status Determination Question. On the face of it, ECN avoids this challenge. According to ECN, there is a single status that is both necessary and sufficient for permissible reliance: *epistemic certainty*.

But perhaps the problem is not so easily avoided. On my contextualist view, there is no single epistemic status denoted by "epistemic certainty"-talk. And so it would seem that ECN faces a version of the Status Determination Question:

Context-Sensitivity Question: How does a context of utterance c determine which epistemic status is denoted by an epistemic use of "certain" in c?

In fact, we can sharpen this worry.[29] In the previous section, I suggested that practical factors are part of the answer, and I used this to explain our intuitions about **Jellybean**. But just how do practical factors fix the epistemic status denoted by "certainty"-talk? As before, the most natural options invoke either stakes or practical adequacy:

Stakes Constraint: *Ceteris paribus*, the higher the stakes in a context c, the stronger the epistemic position denoted by epistemic uses of "certain" in c.

Practical Adequacy Constraint: *Ceteris paribus*, an epistemic use of "certain" in a context c will denote an epistemic position that is practically adequate in c.

Whichever way we go, won't we run into the same sort of trouble that confronted VN?

While this is a serious concern, I think it can be answered. The answer starts by noting an important difference between VN and ECN. As standardly developed, VN tries to answer the Status Determination Question *entirely* in terms of practical factors. The problems I raised for both versions of VN in §3 exploited this fact. But it would be wildly implausible to try to answer the Context Sensitivity Question entirely in terms of practical factors. Perhaps, then, certainty normers can appeal to independent constraints on the meaning of epistemic certainty claims in order to ward off the problems facing VN.

To flesh this out, let's revisit the problems for VN, starting with the stakes approach. The main problem for the stakes version of VN was that it ran afoul of **Survey**: since **Survey** is low stakes, the stakes approach does not explain why it is impermissible to rely on 26a). But as we saw in §6.2, proponents of ECN have an alternative diagnosis of **Survey**. This alternative diagnosis made no reference to stakes; rather, it appealed to the way explicit comparisons tend to drive up the contextual standards governing gradable adjectives.

Next, turn to the practical adequacy approach. The primary problem for the practical adequacy version of VN was that a proposition can be practically adequate even though the

[29] Thanks to a referee for raising this concern.

agent has no good reason to believe it. Recall Mindy, who only has a .01 credence that it will snow (SNOW), and is offered a bet that pays a million dollars if it snows, and loses a penny if it doesn't. Even though SNOW is practically adequate for her, intuitively she should not rely on this proposition in her reasoning. By contrast, certainty normers who adopt the Practical Adequacy Constraint can avoid this problem by appealing to further constaints on epistemic certainty. We've seen that there is good reason to think epistemic certainty entails knowledge (EC⇒K). On virtually every plausible theory of knowledge, knowledge entails rational belief. So we can use EC⇒K to explain why SNOW is not epistemically certain for Mindy, and hence why she cannot permissibly rely on it.

Let's take stock of the dialectic. Our objector pointed out that ECN faces the challenge of explaining how practical factors affect the truth conditions of epistemic certainty claims. However, the challenge turns out to be relatively benign. Regardless of whether we spell things out in terms of stakes or in terms of practical adequacy, we can appeal to independent constraints on the truth conditions of epistemic certainty claims to avoid the problems facing the analogous versions of VN.

Could defenders of VN make a similar move? The problem with this suggestion is that it is unclear where these independent constraints would come from. In the case of ECN, there is a specific context-sensitive epistemic notion—*epistemic certainty*—that is hypothesized to have a privileged connection with permissible reliance. This allowed us to invoke independently motivated constraints on this epistemic notion to explain the problem cases. By contrast, there is no comparable notion that plays this role for VN: the whole idea behind VN is that permissible reliance is not connected to any specific epistemic notion, but rather to a range of different notions depending on the circumstances.

I've argued that certainty normers can appeal to either the Stakes Constraint or the Practical Adequacy Constraint in order to explain how practical factors influence the extension of "certainty"-talk in a context. However, it's not clear that these are the only two options. Let me close this section by mentioning a third possibility. In §5.3, I suggested that epistemic certainty claims are subject to a Salience Constraint. Perhaps we could use this constraint to explain the pragmatic dimension of certainty claims. Here's the rough idea: when the conversational participants recognize that the cost of being mistaken about p is sufficiently high, this recognition tends to render $\neg p$ possibilities salient. By the Salience Constraint, those possibilities are rendered contextually relevant, thereby making it harder, all else equal, for the corresponding certainty claims to come out true. If a story along these lines can be made to work, it would have the advantage of simplicity (better, all else equal, to make do with fewer constraints).[30] But no matter which constraint we settle on, the

[30] A salience-based explanation may also have further benefits. As noted earlier, Anderson and Hawthorne (2019) raise various problem cases for both stakes-based and practical adequacy-based impurist theories of knowledge. Some might worry that these cases also cause trouble for the corresponding constraints on epistemic certainty claims. In response to this worry, one option is to follow Fantl and McGrath (2019), who point out

important point is that ECN has ample resources for avoiding the challenges facing VN.

8 Beyond Practical Reasoning

Thus far I've focused on practical reasoning. However, many philosophers hold that whatever state serves as the norm of practical reasoning also serves as the norm of assertion and belief.[31] Indeed, one powerful motivation for KN is that it meshes nicely with other knowledge norms:

Knowledge Norm of Assertion (KA): A is epistemically permitted to assert p iff A knows p.

Knowledge Norm of Belief (KB): A is epistemically permitted to believe p iff A knows p.[32]

Should proponents of ECN likewise hold that certainty is the norm of assertion and belief? To answer this question, this section considers the status of KA and KB in light of our argument thus far.

8.1 The norm of assertion

These days, KA is a popular normative constraint on assertion. However, when we examine the arguments in its favor, we see a now-familiar pattern play out: the most compelling arguments motivate its necessity direction, not its sufficiency direction.

One of the main arguments for KA is that it explains the infelicity of Moorean assertions (Moore 1962), e.g.:

(16) ? The show starts at 8, but I don't know it does.

According to KA, anyone who asserted (16) would be guaranteed to violate the norm of assertion, since (16) is unknowable. (If one knows the first conjunct, then the second conjunct is false, hence not known.)

As Stanley (2008) and Brown (2010) note, this argument only supports the necessity direction of KA. And if we want to create trouble for the sufficiency direction, we need look no further than our earlier counterexamples to K-Suff. In **Jellybean**, Hugo knows that Caesar was born in 100 BC, but he doesn't seem to be in a position to assert it. Similarly,

that a salience-based approach may do a better job handling some of these cases.

[31] Brown (2012a) calls this the "commonality assumption." For discussion, see Gerken (2014); Simion (2019).

[32] For endorsements of versions of KA, see, a.o., Unger (1975); Williamson (2000); DeRose (2009), Hawthorne (2004), Sutton (2007); Kelp (2018). For sympathetic discussions of KB, see Sutton (2007); Moss (2018); Williamson (forthcoming).

imagine a variant of **Survey** where you have to assert your answer in order to win. It seems you should not assert 26a), even though you know it.

So both KA and KN are afflicted by similar problems. Similar problems call for similar solutions. In both cases, an epistemic certainty norm offers just what's needed. Here's one way of formulating such a norm (cf. Stanley 2008: 48):

Epistemic Certainty Norm of Assertion (ECA): If A is in a context c_A, then it is epistemically permissible for A to assert p iff p is epistemically certain for A relative to the standards of c_A.

A certainty norm of assertion can be supported on independent grounds. As several philosophers have noted,[33] Moorean assertions remain odd when formulated in terms of certainty:

(17) ? The show starts at 8, but it's not certain that it does.

Given the plausible assumption that knowledge does not entail certainty (§6.1), it is unclear how KA can explain the oddity of (17). By contrast, ECA has no trouble. ECA says that in order for (17) to be felicitous, both conjuncts would need to be epistemically certain. But if the first conjunct were certain, the second conjunct would be false, hence not epistemically certain.

A certainty norm of assertion also helps defuse one of the main arguments for K⇒EC. As we saw in §6, the idea that knowledge entails epistemic certainty gains support from the infelicity of (7) (*I know the show starts at 8, but it's not certain that it does*). ECA provides an alternative explanation of this infelicity: (7) is unassertible for the same reason as (17).

8.2 The norm of belief

Is certainty also the norm of belief?

An affirmative answer will be attractive to recent advocates of a "strong" conception of belief, according to which belief requires credence 1.[34] After all, if belief entails psychological certainty, then, by Normative Link, you should only believe p if p is epistemically certain for you.

However, our everyday "belief"-talk calls this strong conception into question. Some examples "from the wild":

(18) He believes, but is not certain, that the M'Gees knew and could have shown the land.[35]

[33] Unger (1975); Stanley (2008); Petersen (2019); Beddor (forthcoming).
[34] See e.g. Clarke (2013); Greco (2015); Dodd (2017); Moss (2018, 2019); Kauss (forthcoming).
[35] Hunt et al. v. Wickliffe, *Cases Argued to the Supreme Court*, January term, 1829.

(19) Tony Blankley, press secretary to House Speaker Newt Gingrich (R-Ga.), said he believes, but is not certain, that House conferees will accept the legislation.[36]

These examples suggest that we ordinarily conceive of belief as a weaker state than psychological certainty (cf. Hawthorne et al. 2016; Beddor and Goldstein 2018). Moreover, it seems that the belief states reported in (18) and (19) are not only psychologically possible, they are also *rationally* possible. Given Normative Link, it follows that one can permissibly hold a belief that isn't epistemically certain.

However, this should not be regarded as a problem for ECN and ECA, since there is independent evidence that the norm of belief is weaker than the norms of practical reasoning and assertion.

For evidence that the norm of belief is weaker than the norm of practical reasoning, note that it sounds perfectly natural to say things like:

(20) I believe the flight leaves from Terminal A, but we should check.

Assume the speaker is epistemically permitted to hold the belief reported in the first conjunct: *the flight leaves from Terminal A* (TERMINAL). If belief and practical reasoning are governed by the same norm, then the speaker is permitted to rely on TERMINAL. But if the speaker is permitted to rely on TERMINAL, then, by Reliance-Action Bridge, they are permitted to perform whatever action has the highest expected utility conditional on TERMINAL. So there should be no need to check.

This line of reasoning dovetails nicely with recent arguments by Hawthorne et al. (2016) that the norm of belief is weaker than the norm of assertion. Hawthorne et al. note a difference in felicity between:

(21) I believe the show starts at 8, but I know/realize it might start later.

(22) ? The show starts at 8, but I know/realize it might start later.

This contrast is unexpected if belief and assertion are governed by the same norm. But it is readily explained by the view offered here. According to ECA, felicitiously asserting (22) requires that both conjuncts are epistemically certain. But if the first conjunct were epistemically certain, then there would be no possibility of being mistaken, which is contradicted by the second conjunct. By contrast, in order to permissibly believe the show starts at 8, a weaker epistemic position suffices.

Is there any way to defend a certainty norm of belief against these arguments? Some proponents of a strong conception of belief acknowledge that their view stands in tension with our ordinary "belief"-talk. In response to this tension, they sometimes insist that we can only make sense of the roles that belief plays in assertion and practical reasoning on the assumption that belief requires credence 1.[37] Forced to choose between respecting ordinary

[36] https://www.latimes.com/archives/la-xpm-1996-09-13-mn-43572-story.html

[37] E.g. Greco (2015: 180).

"belief"-talk and making sense of the functional role of belief, they go the latter route.

The framework developed in this paper allows us to avoid this dilemma. We can respect the ordinary conception of belief as weak, while agreeing that a strong cognitive state is needed to explain practical reasoning and assertion. The solution is to give up the presupposition that this strong cognitive state is belief. Fortunately, ordinary language already furnishes us a perfectly good name for this strong state: "certainty"![38]

9 Conclusion

In this paper, I've made a case for a certainty norm of practical reasoning. This norm avoids the problems facing KN and VN. It also gains support from independent arguments for a certainty norm of assertion.

These arguments also motivate a more general perspective on the role of certainty in psychological and epistemological explanation. According to this outlook, the central psychological state involved in explaining action—or at least in explaining reliance—is *psychological certainty*. And this state is governed by its own normative standard: *epistemic certainty*. Rather than trying to assimilate either form of certainty to more familiar categories—belief or knowledge—we should recognize them for what they are, and put them into action.[39]

References

Charity Anderson and John Hawthorne. Knowledge, Practical Adequacy, and Stakes. In Gendler and Hawthorne, editors, *Oxford Studies in Epistemology*, pages 234–257. Oxford University Press, Oxford, 2019.

D.M. Armstrong. Does Knowledge Entail Belief? *Proceedings of the Aristotelian Society*, 70(1):21–36, 1969.

Bob Beddor. *Reduction in Epistemology*. PhD thesis, Rutgers University, 2016.

Bob Beddor. New Work for Certainty. *Philosophers' Imprint*, forthcoming.

[38] We should be careful not overstate the difference between the view defended in this paper and the strong conception of belief. Both views agree that a strong cognitive state is needed to perform various theoretical roles. Where we disagree is whether this state deserves the name "belief." That said, this difference is not merely terminological, since distinguishing between belief and psychological certainty opens up new theoretical possibilities. In particular, it opens up the possibility that there may be no single state that does all of the work traditionally assigned to belief. Some of the theoretical roles are played by a weak cognitive state—what we ordinarily call "belief." Other roles are played by a distinct state—psychological certainty.

[39] An early version of some of these ideas appeared in chapter three of my dissertation (Beddor 2016). I am grateful to Simon Goldstein, Blake Roeber, the NUS philosophy reading group, and three referees at *The Philosophical Quarterly* for detailed comments on an earlier draft. Thanks also to Andy Egan, Alvin Goldman, Jonathan Schaffer, and participants in Ernie Sosa's dissertation group for helpful discussions about certainty and its role in epistemology.

Bob Beddor and Simon Goldstein. Believing Epistemic Contradictions. *The Review of Symbolic Logic*, 11(1):87–114, 2018.

Michael Blome-Tillman. Knowledge and Presuppositions. *Mind*, 118(470):241–294, 2009.

Jessica Brown. Subject-Sensitive Invariantism and the Knowledge Norm for Practical Reasoning. *Noûs*, 42(2):167–189, 2008a.

Jessica Brown. Knowledge and Practical Reason. *Philosophy Compass*, 3(6):1135–1152, 2008b.

Jessica Brown. Knowledge and Assertion. *Philosophy and Phenomenological Research*, 81(3):549–566, 2010.

Jessica Brown. Assertion and Practical Reasoning: Common or Divergent Standards? *Philosophy and Phenomenological Research*, 84(1):123–157, 2012a.

Jessica Brown. Practical Reasoning, Decision Theory and Anti-Intellectualism. *Episteme*, 9(1):43–62, 2012b.

Roger Clarke. Belief Is Credence One (In Context). *Philosophers' Imprint*, 13(11):1–18, 2013.

Keith DeRose. *The Case for Contextualism: Knowledge, Skepticism, and Context.* Clarendon Press, Oxford, 2009.

Dylan Dodd. Belief and Certainty. *Synthese*, 194(11):4597–4621, 2017.

Fred Dretske. The Pragmatic Dimension of Knowledge. *Philosophical Studies*, 40:363–379, 1981.

Jeremy Fantl and Matthew McGrath. Evidence, Pragmatics, and Justification. *Philosophical Review*, 111(1):67–94, 2002.

Jeremy Fantl and Matthew McGrath. *Knowledge in an Uncertain World.* Oxford University Press, Oxford, 2009.

Jeremy Fantl and Matthew McGrath. Clarifying pragmatic encroachment. In Szabó Gendler and Hawthorne, editors, *Oxford Studies in Epistemology*, volume 6. Oxford University Press, Oxford, 2019.

Jie Gao. Rational Action Without Knowledge. *Synthese*, 194(6):1901–1917, 2017.

Mikkel Gerken. Warrant and Action. *Synthese*, 178(3):529–547, 2011.

Mikkel Gerken. Same, Same but Different. *Philosophical Studies*, 168:725–744, 2014.

Daniel Greco. How I Learned to Stop Worrying and Love Probability 1. *Philosophical Perspectives*, 29(1):179–201, 2015.

Daniel Greco. Justification and Excuses in Epistemology. *Noûs*, forthcoming.

John Hawthorne. *Knowlege and Lotteries.* Oxford University Press, New York, 2004.

John Hawthorne and Jason Stanley. Knowledge and Action. *Journal of Philosophy*, 105(10):571–590, 2008.

John Hawthorne, Daniel Rothschild, and Levi Spectre. Belief Is Weak. *Philosophical Studies*, 173(5): 1393–1404, 2016.

Jaakko Hintikka. *Knowledge and Belief: An Introduction to the Logic of Two Notions.* Cornell University Press, Ithaca, 1962.

Dominic Kauss. Credence as Doxastic Tendency. *Synthese*, forthcoming.

Christoph Kelp. Assertion: A Function First Account. *Noûs*, 52(2):411–442, 2018.

Christoph Kelp and Mona Simion. Criticism and Blame in Action and Assertion. *Journal of Philosophy*, 114(2):76–93, 2017.

Christopher Kennedy. Vagueness and Grammar: The Semantics of Relative and Absolute Gradable Adjectives. *Linguistics and Philosophy*, 30:1–45, 2007.

Angelika Kratzer. What 'Must' and 'Can' Must and Can Mean. *Linguistics and Philosophy*, 1:337–355, 1977.

Jennifer Lackey. Acting on Knowledge. *Philosophical Perspectives*, 24:361–382, 2010.

David Lewis. Scorekeeping in a Language Game. *Journal of Philosophical Logic*, 8(3):339–359, 1979.

David Lewis. Elusive Knowledge. *Australasian Journal of Philosophy*, 74(4):549–567, 1996.

Dustin Locke. Practical Certainty. *Philosophy and Phenomenological Research*, 90(1):72–95, 2015.

Aidan McGlynn. *Knowledge First?* Palgrave Macmillian, New York, 2014.

G.E. Moore. Certainty. In *Philosophical Papers*. Allen and Unwin, London, 1959.

G.E. Moore. *Commonplace Book: 1919-1953*. Allen and Unwin, London, 1962.

Sarah Moss. Epistemology Formalized. *Philosophical Review*, 122(1):1–43, 2013.

Sarah Moss. *Probabilistic Knowledge*. Oxford University Press, Oxford, 2018.

Sarah Moss. Full Belief and Loose Speech. *Philosophical Review*, 128(3):255–291, 2019.

Ram Neta. Treating Something as a Reason for Action. *Noûs*, 43:684–699, 2009.

Carlotta Pavese. Know How and Gradability. *The Philosophical Review*, 126(3):345–383, 2017.

Esben Nedenskov Petersen. A Case for a Certainty Norm of Assertion. *Synthese*, 196(11):4691–4710, 2019.

Colin Radford. Knowledge: By Examples. *Analysis*, 27(1):1–11, 1966.

Baron Reed. A Defense of Stable Invariantism. *Noûs*, 44(2):224–244, 2010.

Blake Roeber. The Pragmatic Encroachment Debate. *Noûs*, 52(1):171–195, 2018.

Jacob Ross and Mark Schroeder. Belief, Credence, and Pragmatic Encroachment. *Philosophy and Phenomenological Research*, 88(2):259–288, 2014.

Carmen Rotstein and Yoad Winter. Total Adjectives vs. Partial Adjectives: Scale Structure and Higher-Order Modifiers. *Natural Language Semantics*, 12:259–288, 2004.

Stephen Schiffer. Interest-Relative Invariantism. *Philosophy and Phenomenological Research*, 75(1): 188–195, 2007.

Mona Simion. Epistemic Norm Correspondence and the Belief-Assertion Parallel. *Analysis*, 79(2): 265–275, 2019.

Jason Stanley. On the Linguistic Basis for Contextualism. *Philosophical Studies*, 119(1/2):119–146, 2004.

Jason Stanley. *Knowlege and Practical Interests*. Oxford University Press, New York, 2005.

Jason Stanley. Knowledge and Certainty. *Philosophical Issues*, 18(1):35–57, 2008.

Jonathan Sutton. *Without Justification*. The MIT Press, Cambridge, MA, 2007.

Peter Unger. *Ignorance: The Case for Skepticism*. Clarendon Press, Oxford, 1975.

Brian Weatherson. Knowledge, Bets, and Interests. In Brown and Gerken, editors, *Knowledge Ascriptions*, pages 75–103. Oxford University Press, New York, 2012.

Jonathan Weisberg. Knowledge in Action. *Philosophers' Imprint*, 13(22):1–23, 2013.

Timothy Williamson. *Knowledge and Its Limits*. Oxford University Press, Oxford, 2000.

Timothy Williamson. Contextualism, Subject-Sensitive Invariantism and Knowledge of Knowledge. *The Philosophical Quarterly*, 55(219):213–235, 2005.

Timothy Williamson. Justifications, Excuses, and Sceptical Scenarios. In Dorsch and Dutant, editors, *The New Evil Demon Problem*. Oxford University Press, Oxford, forthcoming.