

CERTAINTY FIRST

Abstract

This paper argues that we should assign certainty a central place in epistemology. While epistemic certainty played an important role in the history of epistemology, recent epistemology has tended to dismiss certainty as an unattainable ideal, focusing its attention on knowledge instead. I argue that this is a mistake. Attending to certainty attributions in the wild suggests that much of our everyday knowledge qualifies, in appropriate contexts, as certain. After developing a semantics for certainty ascriptions, I put certainty to explanatory work. Specifically, I argue that by taking certainty as our central epistemic notion, we can shed light on a variety of important topics, including evidence and evidential probability, epistemic modals, and the normative constraints on credence.

1 Introduction

For much of its history, epistemology focused on certainty. In the medieval and early modern periods, epistemological discussions centered on *scientia*, conceived as the highest grade of knowledge. Most authors in this tradition explicitly associated *scientia* with certainty. For example, we find Aquinas, Scotus, and Descartes all characterizing *scientia* in terms of “certain [*certa*] cognition.”¹

For these authors, *scientia*—and more specifically, certainty—is the optimal epistemic status: it is “perfect cognition.” And while certainty was conceived as the epistemic ideal, it was considered an attainable ideal. This epistemic optimism is perhaps clearest in Descartes, who urged that by properly directing our reason, we can elevate much of our everyday knowledge to the status of *scientia*.

But as time went on, this optimism waned. According to an emerging consensus, most of our knowledge only approaches the optimal epistemic status; precious little reaches this ideal. Thus we find Russell (1912) concluding that only our knowledge of our sensory states makes the cut—all else is merely probable. In a similar vein, Reichenbach concludes that all attempts to achieve certainty have failed: “The search for certainty had to burn itself out” (1963: 49). And Unger (1971, 1975) develops a linguistically sophisticated version of this line of thought, arguing that *certain* functions as an absolute term, which only applies to a proposition if that proposition has the maximal degree of certainty. This is then

¹Aquinas 1970; Scotus 1987; Descartes 1983, AT X: 362. See Pasnau 2014, 2017 for discussion of the connection between *scientia* and certainty in the medieval and early modern tradition.

conjoined with a general argument that absolute terms almost never apply to anything in the world.²

If certainty is so rarely attained, it is natural to wonder whether it should matter much for epistemology. For many contemporary epistemologists, the answer seems to be ‘No.’ According to a common view, while certainty is largely beyond our reach, knowledge is much more easily attained; consequently, it’s better-suited to play a central role in epistemology. This shift in focus from certainty to knowledge is perhaps most apparent in the ‘Knowledge First’ program in epistemology. For Knowledge Firsters, knowledge is the central epistemic notion: it can be used to illuminate other epistemic notions (e.g., evidence, evidential probability, justification) and to provide normative constraints on assertion and action.³ Typically, Knowledge Firsters devote scant space to certainty: peruse the major contributions to Knowledge First epistemology and you’ll find certainty mentioned in passing, if at all.

In this paper, I seek to restore certainty to its former centrality. I open by inquiring into the nature of certainty. I argue that attempts to reduce certainty to knowledge flounder: certainty is a *sui generis* epistemic notion. At the same time, I contest the perception of certainty as an unattainable ideal. I argue that much—though by no means all—of our ordinary knowledge can rise to the level of certainty.⁴

Having laid the groundwork, the rest of the paper explores how much explanatory mileage we can get out of certainty: if we were to take certainty as our primitive, how far could this take us? Surprisingly far, it turns out. We can use certainty to illuminate evidence, evidential probability, and epistemic modals, as well as the normative constraints on credence and assertion. In view of these results, I urge that certainty deserves a central place among our epistemic concepts.

2 Towards an account of certainty

While the analysis of knowledge has generated a voluminous literature, the analysis of certainty has received comparatively little attention. In this section, I try to remedy this state of neglect. I propose an account of certainty that has two main virtues: it makes sense of the semantic properties of everyday *certainty*-talk, and it sheds light on the connections between certainty, knowledge, and belief. Once our account is in place, we will be in a better position to evaluate the charge that certainty is an unattainable ideal—a topic that I tackle in §3.

2.1 Subjective vs. epistemic certainty

It’s common to distinguish between subjective and epistemic certainty.⁵ Subjective certainty is a matter of strength of conviction. A belief can be subjectively certain even if it is

²Other prominent 20th century endorsements of the view that certainty is seldom, if ever, attainable include Dewey 1929, C.I. Lewis 1929: 309, and Ayer 1936: 127.

³See Williamson 1996, 1997, 2000; Hawthorne 2004; Hawthorne and Stanley 2008; Sutton 2007; Weatherson 2012; Littlejohn forthcoming, among many others.

⁴On this front I agree with Miller (1978); Klein (1981); and Stanley (2008).

⁵See e.g. Moore 1959; Klein 1981; Stanley 2008; DeRose 2009; Reed 2011.

held for no good reason. By contrast, if a belief is epistemically certain, the believer must stand in a strong epistemic relation to its content.

While *certain* and its cognates are ambiguous between these two senses, some constructions favor one reading over the other (Moore 1959; Stanley 2008; DeRose 2009). Claiming that a person is certain of something usually conveys subjective certainty:

(1) I'm certain/sure that the butler did it.

(1) can be true even if the speaker irrationally believes the butler did it.

Claiming that a proposition is certain usually conveys epistemic certainty:

(2) It's certain that the butler did it.

(2) seems to entail that the speaker stands in a strong epistemic position with regards to the proposition that the butler did it.

What is the relation between these two species of certainty? A natural thought is that the link is normative: p is epistemically certain for A iff A *ought* to be subjectively certain that p . This explains the oddity of conjunctions of the form:

(3) ? $\left\{ \begin{array}{l} \text{It's certain} \\ \text{I'm certain} \end{array} \right\}$ that the butler did it, but $\left\{ \begin{array}{l} \text{I'm not certain} \\ \text{it's not certain} \end{array} \right\}$ he did it.

According to this proposal, (3) is infelicitous because anyone who uttered it would be committed to violating a basic rational requirement.

2.2 Certainty vs. knowledge

Can more be said about either form of certainty? For many—particularly those sympathetic to the Knowledge First program—it will be tempting to understand certainty in terms of knowledge. Subjective certainty, some may suggest, is the level of confidence required for knowledge. Epistemic certainty is the epistemic position required for knowledge: it is being in a position to know.⁶

However, I think there is reason to doubt that knowledge requires either species of certainty. First, *knows for certain* is not redundant. To see this, imagine that it's the first day of Epistemology 101, and you're trying to get your students to feel the pull of Descartes' project. Most likely, you'd ask (4a) rather than (4b):

- (4) a. What can we know for certain/with certainty?
b. What can we know?⁷

⁶The idea that knowledge entails either subjective or epistemic certainty (or both) can be found in Ayer 1956; Moore 1959; Unger 1975, among others.

⁷Arguably, asking (4a) rather than (4b) fits better with Descartes' own views on knowledge. While Descartes is widely interpreted as holding that knowledge requires certainty, Descartes' discussion of the atheist mathematician in the Second Replies casts doubt on this interpretation. In his discussion, Descartes draws a distinction between *cognitio* and *scientia*: the atheist's belief that a triangle's three angles are equal to two right angles amounts to *cognitio*, but not *scientia* (AT VII 141). On a natural reading, *cognitio* still amounts to a species of knowledge, it is simply a lower grade than *scientia*. For further discussion, see Sosa 1997; Wykstra 2008; Pasnau 2014.

More generally, if I say that someone knows something with certainty, I seem to be making a stronger claim than if I merely say that they know it.

Could we explain the difference between (4a) and (4b) on pragmatic grounds? Perhaps, some may suggest, both *knows* and *certain* are context-sensitive expressions governed by the same standards. And so in any context in which *A knows p* is true, the corresponding subjective and epistemic certainty ascriptions are also true. However, perhaps coupling *knows* and *certain* together in the complex phrase *knows for certain* drives up the standards for both knowledge and certainty. As a result, we interpret (4b) and its ilk as invoking unusually stringent standards for certainty.

In order for this pragmatic explanation to be plausible, it would need to follow from a more general principle governing the interpretation of context-sensitive expressions. According to this more general principle, whenever two context-sensitive expressions are governed by the same standards, combining them in a complex phrase drives up the standards associated with each. But if we consider other context-sensitive expressions we find that things don't work this way. For example, *certain* and *sure* are presumably governed by the same standards: if I say that the detective is certain the butler did it, this seems to entail that the detective is sure the butler did it, and *vice versa*. But claiming that the detective is *certain and sure* that the butler did it sounds odd: it is not naturally interpreted as saying that the detective is certain to an unusually high degree. Similarly, *likely* and *probable* are presumably governed by the same standards. But claiming that an event is *likely and probable* smacks of redundancy; it's not naturally interpreted as saying that the event is extremely likely.

A further difficulty for a pragmatic explanation of the difference between (4a) and (4b)—and a further reason to doubt the knowledge-certainty entailment—comes from cases where it's natural to ascribe knowledge while denying certainty. Consider the unconfident examinee (Woozley 1952; Radford 1966). Throughout his oral history exam, his answers are fumbling and hesitant, yet invariably correct. The exam concluded, it would be natural for his surprised examiner to remark, 'Turns out he knew the answers all along'. Yet it would also be natural to deny that he was certain of the answers (Armstrong 1969; Stanley 2008; McGlynn 2014).

Ascriptions of knowledge without certainty are not confined to the pages of philosophical journals. Some examples 'from the wild':

- (5) [W]e know without certainty, but with a high degree of probability, that returns over the next 10 years or so will be very poor.⁸
- (6) When [a false ID] is handed to a cop, he knows with near certainty the guy before him is not the guy identified on the flimsy piece of paper.⁹
- (7) We now know with near-certainty that Russia did this with the goal of electing Trump.¹⁰

⁸http://www.smithers.co.uk/news_article.php?id=16&o=50.

⁹Geeting 2005: 96

¹⁰<http://nymag.com/daily/intelligencer/2016/12/trump-mcconnell-putin-and-the-triumph-of-the-will-to-power.html>

- (8) The only way to know with close to certainty the gender of your baby is to have medical tests performed.¹¹

In (5), the speaker allows for knowledge without certainty. In (6)-(8), the speaker allows for knowledge that is nearly certain, implying that it is not actually certain.¹²

Faced with these examples, some may concede that knowledge does not require subjective certainty, while still maintaining that it requires epistemic certainty. However, this position stands in tension with the normative connection between subjective and epistemic certainty. As we've seen, it's natural to hold that one should be subjectively certain of p iff p is epistemically certain. If knowledge entails epistemic certainty, then anyone who knows p is rationally required to be subjectively certain that p . But this seems wrong. Consider again the unconfident examinee. While the examinee's memory is highly reliable, it could be rational for him to harbor doubts about its reliability. As a result, it could be rational for him to be less than certain that, say, Elizabeth I died in 1603.

While this is hardly the last word on the matter, I think these considerations give reason to doubt that knowledge entails either subjective or epistemic certainty.¹³ On the picture that emerges, subjective certainty involves a particularly high degree of confidence—higher than that required for knowledge. And in order for such a high degree of confidence to be warranted, one must be in a particularly strong epistemic position—stronger than that usually required for knowledge.

How should we understand these differences in strength of epistemic position? According to a common view, knowledge involves eliminating possibilities of error: to know p is to be in a state that rules out possibilities in which p is false. However, it need not rule out *all* possibilities of error, only those that are sufficiently plausible, or sufficiently nearby. Perhaps epistemic certainty likewise eliminates possibilities of error, just a wider range thereof.

To illustrate, take one of our ascriptions of knowledge without certainty, (6) (*When [a false ID] is handed to a cop, he knows with near certainty the guy before him is not the guy identified on the flimsy piece of paper*). Here the speaker is claiming that when a cop receives a false ID, the cop's epistemic state eliminates all plausible scenarios in which the person in front of him is the person whose name is on the ID. But his epistemic state leaves

¹¹<http://www.pregnancy-info.net/fetal-development/baby-boy.html>

¹²These examples notwithstanding, some people find it odd to claim someone *knows without certainty*. I suspect this is because, without further elaboration, it may be unclear what the speaker is trying to convey. Is the speaker claiming that the subject knows p even though p is *not at all* certain? If so, then the claim is indeed absurd. Or is the speaker claiming that the subject knows p with a relatively high degree of certainty, just not high enough to qualify as *certain* (full-stop)? Note that in all of the naturally occurring examples I've quoted, the speaker clarifies that the second interpretation was intended—by adding *with a high degree of probability* in (5), by using *near* or *close* to rather than *without* in (6)-(8). Indeed, we may expect speakers to make some clarification along these lines, since we expect speakers to be as clear and informative as possible. Once such clarifications are made, the relevant claims sound much more natural.

¹³While I'll be operating with this picture in what follows, people who think that knowledge entails certainty can still accept much of what I say in the remainder of this paper. In particular, they can accept many of my claims about certainty, as well as many of my proposals for understanding various epistemological phenomena in terms of certainty. However, such readers will extract a different lesson from the paper as a whole: for them, a Certainty First framework should not be regarded as a *rival* to Knowledge First epistemology, instead, it's the best way of developing the latter.

open various far-fetched possibilities in which this isn't the case—for example, scenarios in which someone created a fake ID for themselves in order to sow confusion.

2.3 *Certain as a quantifier over worlds*

One way to develop this proposal with greater precision draws on the resources of epistemic and doxastic logic. The standard approach to epistemic and doxastic logic, due to Hintikka (1962), treats *knows* and *believes* akin to modal operators. For someone to know p is for p to hold in all worlds consistent with what they know—call these the '*K*-alternatives'.¹⁴ For someone to believe p is for p to hold in all worlds consistent with what they believe—call these the '*B*-alternatives'. One attractive feature of this framework is that it allows us to model properties of knowledge and belief in terms of constraints on the underlying accessibility relations. For example, to capture the factivity of knowledge, it's standard to take the *K*-alternatives at w to include w . To capture the idea that knowledge asymmetrically entails belief, it's standard to take the *K*-alternatives at w to include the *B*-alternatives, but not *vice versa*.

This framework extends naturally to certainty. We can propose that for p to be epistemically certain is for p to hold in every world consistent with what is epistemically certain—call these the '*E*-alternatives'. To capture the idea that epistemic certainty requires a stronger epistemic position than knowledge, we require that the *K*-alternatives are always a subset of the *E*-alternatives, but not *vice versa*.¹⁵ (See Fig. 1.)




Figure 1: Knowledge & Epistemic Certainty

Formulated thus, the account is non-reductive: it does not try to explain epistemic certainty in more basic terms. It could, however, be supplemented with a reductive account of the *E*-alternatives. For example, internalists could take the *E*-alternatives to be the worlds consistent with the agent's phenomenal states. Reliabilists could take them to be

¹⁴I use '*K*-alternatives' rather than the more familiar label, 'epistemic alternatives' because I want to distinguish the possibilities consistent with what's known from the possibilities consistent with what's epistemically certain. Talk of 'epistemic alternatives' can be read either way, and hence is liable to cause confusion.

¹⁵This predicts that epistemic certainty entails knowledge. Some readers may question this, on the grounds that p could be epistemically certain without being believed—for example, if someone's evidence entails p , but they have not yet drawn this inference. Such readers can still accept my framework by interpreting the *K*-alternatives slightly differently, as the worlds consistent with what one is in a position to know, rather than the worlds consistent with what one knows.

the worlds consistent with whichever of the agent's beliefs are produced by a maximally reliable process. I suspect that the truth is more complicated than either of these simple pictures, and that it may not be possible to characterize the *E*-alternatives in more basic terms. However, there is no need to take a stand on this matter here.

Even if we lack a reductive characterization of the *E*-alternatives, we could perhaps use the *E*-alternatives in service of a reductive account of knowledge. For example, suppose we help ourselves to a notion of comparative *closeness* between worlds. We could then define the *K*-alternatives at *w* as the *E*-alternatives that are sufficiently close to *w*. If *closeness* can be understood without recourse to knowledge, then this would amount to a definition of knowledge in terms of closeness and epistemic certainty. Here too I do not want to take a stand on whether such a reduction is defensible; I only aim to highlight the explanatory resources of the present framework.

This framework can also be used to model subjective certainty. Intuitively, subjective certainty is a stronger state than belief. Someone can believe that it will rain later without being certain of it, but not *vice versa*. This suggests the following picture: *subjective certainty is to belief as epistemic certainty is to knowledge*. To model this, we can hold that someone is subjectively certain of *p* iff *p* holds in every world consistent with their subjective certainties—call these the '*S*-alternatives'. To capture the asymmetric entailment between subjective certainty and belief, we require that the *S*-alternatives are always a superset of the *B*-alternatives.¹⁶

A quantifier-over-worlds model of certainty yields a number of further predictions, two of which are worth mentioning. First, since the *E*-alternatives include the *K*-alternatives, which include the actual world, our model predicts that epistemic certainty entails truth. This seems plausible. Suppose the detective has good but not conclusive evidence that the butler is guilty, and consequently exclaims (2) (*It's certain that the butler did it*). Suppose that further investigation reveals that the butler was framed. It would be natural for the detective to retract her claim:

(9) Ok, I guess I was wrong when I said that it was certain the butler did it.

By contrast, it would be far less natural for the detective to 'stick to her guns' and defend the truth of her earlier claim:

(10) ? What I said was perfectly true. After all, I didn't say he did it. Only that it's *certain* that he did it.¹⁷

A second prediction of the quantifier-over-worlds approach is that both epistemic and subjective certainty are closed under entailment: if $p_1 \dots p_n$ are epistemically or subjectively certain, so is any consequence of $p_1 \dots p_n$. This seems plausible as well. Indeed, a closure requirement on certainty seems more plausible than a closure requirement on rational belief, or even knowledge. After all, some authors contend that one can rationally

¹⁶Here too there is a question as to whether this can be leveraged into a reductive account of belief. One promising option here would be to identify the *B*-alternatives with the *S*-alternatives that the believer takes to be the most plausible candidates for actuality.

¹⁷Cf. von Fintel and Gillies (2010), who offer a similar argument for the conclusion that epistemic *must* is factive.

believe of each ticket in a lottery that it will lose, without being in a position to believe that all the tickets will lose (Kyburg 1961). Whatever we wish to say about the rationality of *believing* lottery propositions, presumably one cannot rationally be *certain* that a given ticket will lose. And so lotteries do not create even a *prima facie* difficulty for closure requirements on certainty.¹⁸

2.4 *Certain* as a gradable adjective

While an account along these lines strikes me as promising, I don't think it can be the complete story. As it stands, this account leaves out an important aspect of certainty: its *gradability*.

Certainty comes in degrees. This is reflected in the fact that *certain*—on both its subjective and epistemic uses—combines with degree morphology:

- (11) It's fairly/very/95% certain the Mets will win.
- (12) Sal is fairly/very/95% certain the Mets will win.

How should we analyze these 'graded' uses of *certain*? The quantifier-over-worlds approach offers no answer. It tells us how to analyze ungraded or 'pos form' certainty ascriptions such as (1) and (2), but not their graded cousins.¹⁹

If we turn to the semantics literature for guidance, we find a well-developed framework for analyzing gradable adjectives. The core idea is that gradable adjectives are associated with scales. In the case of *tall*, the scale will be degrees of height; in the case of *expensive*, it will be units of cost. The semantic value of a gradable adjective is taken to be a function from entities to degrees on the associated scale.

In order to apply this scalar semantics to *certain*, we can associate subjective uses of *certain* with a *subjective certainty function* (SC) from propositions and agents to degrees on a *subjective certainty scale*, which measures degrees of subjective confidence. Likewise, we associate epistemic uses of *certain* with an *epistemic certainty function* (EC) from propositions and agents to degrees on an *epistemic certainty scale*, which measures strength of epistemic position.²⁰ This yields a simple analysis of graded certainty ascriptions such as (11) and (12). On this analysis, degree modifiers combine with *certain* in the usual ways

¹⁸Some may question the plausibility of the closure requirement on subjective certainty. Are ordinary agents subjectively certain of every complex tautology? In response, I think we should grant that this is indeed a worrisome consequence, but point out that this is a special instance of a more general problem that arises for standard epistemic and doxastic logics: the problem of logical omniscience. For present purposes, I do not take a stand on how best to resolve the problem. Perhaps quantifying over 'impossible worlds' will help (Hintikka 1975), perhaps not. Whatever one's preferred strategy for handling the problem of logical omniscience, one can use that strategy to address the special case of this problem involving subjective certainty.

¹⁹A gradable adjective occurs in the 'pos' (short for 'positive') form if it lacks overt degree morphology—e.g., *x is full* (pos form) vs. *x is fairly/very/95% full* (graded).

²⁰For our purposes, we can refrain from taking a stand on whether it's possible to give a reductive account of EC. Those of an internalist persuasion might try to understand a proposition's degree of epistemic certainty in terms of the degree to which it coheres with the agent's doxastic and phenomenal states. Those of a reliabilist bent might understand it in terms of the degrees of reliability of the agent's belief-forming processes. But this is a point on which the semantics can remain neutral.

to deliver particular degrees on the subjective and epistemic certainty scales: *fairly* (subjectively) *certain* will deliver a fairly high degree of subjective certainty; *95%* (epistemically) *certain* will deliver a .95 degree of epistemic certainty, etc.

However, an important question remains. How exactly do graded certainty ascriptions relate to their ungraded, pos form counterparts? We'd like our analysis to shed light on this. For example, we'd like to predict that (13a) entails (13b), but not *vice versa*:

- (13) a. Sal is certain that the Mets will win.
b. Sal is fairly certain that the Mets will win.

Happily, the standard scalar semantics also comes with a story about this. The standard strategy is to take pos form constructions to contain a null morpheme (pos) that combines with a gradable adjective to deliver some threshold on the associated scale. Thus the underlying form of (13a) is:

- (14) Sal is POS certain that the Mets will win.

In the case of 'relative' gradable adjectives such as *long*, *tall*, and *expensive*, the threshold will be settled by context, and is often vague. In the case of 'maximum-standard' gradable adjectives (*max adjectives*, hereafter) such as *clean*, *straight*, and *full*, the threshold is always the maximal element of the associated scale.²¹

The differences between relative and max adjectives show up on a range of diagnostics. First, sentences of the form, *x is A, but x could be A-er* are fine when *A* is a relative adjective, but anomalous when it is a max adjective:

- (15) This line is long, but it could be longer.
(16) ? The line is straight, but it could be straighter.²²

A second diagnostic looks at interactions with degree modifiers. Max adjectives tolerate the modifiers *almost* and *completely* to a much greater degree than their relative brethren (Rotstein and Winter 2004; Kennedy 2007):

- (17) This line is almost straight/#long.
(18) This line is completely straight/#long.

As a number of authors have noted,²³ *certain*—on both its subjective and epistemic uses—seems to pass the tests for a max adjective with flying colors:

- (19) ? It's certain to rain, but it could be more certain.

²¹For discussion, see Unger 1975; Kennedy and McNally 2005; Kennedy 2007. Some authors also posit a class of minimum-standard adjectives (*bent*, *dirty*, *open*) whose threshold is always the minimal element of the associated scale. For our purposes, we can afford to ignore minimum-standard adjectives.

²²Not everyone shares the intuition that (16) is infelicitous. This may be due to the fact that speakers are willing to use max adjectives imprecisely (Kennedy 2007). However, most speakers I've consulted agree with the comparative claim that (16) is *more* anomalous than (15). Note also that (16) becomes worse if we change the second conjunct to *it could be much straighter*. By contrast, the corresponding adjustment to (15) does not diminish its acceptability.

²³See e.g., Unger 1975; Lassiter 2010, 2011, 2017; Klecha 2012.

- (20) We're almost certain to lose.
 (21) I'm/it's completely certain she'll be there.

If we help ourselves to an off-the-shelf scalar semantics for max adjectives, we get the following picture. A pos form subjective certainty ascription, *A is pos certain that p*, is true iff *A* has the maximal degree of subjective certainty in *p*. Likewise, a pos form epistemic certainty ascription, *p is pos certain*, is true iff *p* has the maximal degree of epistemic certainty (for the contextually supplied agent).²⁴

A scalar semantics along these lines provides a way of relating pos form certainty ascriptions to their graded counterparts. In doing so, it validates entailments between the two—for example, that (13a) asymmetrically entails (13b). And while the primary motivation for such an account is semantic, the core idea meshes well with the traditional thought that certainty constitutes a particularly exalted ideal—that it is the highest form of cognition.²⁵

2.5 Integrating the two approaches

We have, then, two analyses of certainty. One uses the tools of epistemic logic, analyzing *certain* as a quantifier over worlds. The other uses the tools of scalar semantics, analyzing *certain* as a measure function. Both have their advantages. The quantifier-over-worlds approach captures the connections between epistemic certainty and knowledge, and between subjective certainty and belief. The scalar semantics captures the gradability of certainty, as well as the relations between pos form and graded certainty ascriptions. It would be nice if we could integrate the two approaches in a way that preserves their advantages.

Luckily we can. In developing the scalar approach, we said little about the *structure* of the subjective and epistemic certainty functions. Let us now venture the following hypothesis: both are probability functions, defined over algebras generated from the *S*-alternatives and the *E*-alternatives respectively. *SC* is a subjective probability function; *EC* an epistemic probability function.²⁶ Let us also assume that these probability functions are regular, in that they only assign the maximal degree of certainty to a proposition if it holds at every accessible world. This allows us to synthesize our two approaches: *A is pos certain that p* is true iff *A* assigns *p* the maximal degree of subjective certainty, which

²⁴A candidate formal implementation: let Θ_{EC} be the epistemic certainty scale, and let EC_x be *x*'s epistemic certainty function. The semantic values of *certain* and *pos certain* can be formulated as follows:

$$[[\text{certain}_E]] = [\lambda p_{\langle s,t \rangle} [\lambda x_e. EC_x(p)]]; \quad [[\text{pos certain}_E]] = [\lambda p_{\langle s,t \rangle} [\lambda x_e. EC_x(p) = \mathbf{max}(\Theta_{EC})]].$$

Similar lexical entries can be provided, *mutatis mutandis*, for subjective uses of *certain*.

²⁵For a very different scalar treatment of certainty ascriptions, see Stanley 2008: 54. On Stanley's approach, a pos form certainty ascription is true iff the relevant proposition's degree of certainty exceeds some contextually determined threshold. One concern for Stanley's approach is that it in effect amounts to analyzing *certain* as a relative gradable adjective. It thus has trouble explaining why *certain* behaves differently from relative gradable adjectives on the various diagnostics canvassed here.

²⁶Cf. Lassiter (2010, 2011), who also argues that the semantic value of *certain* is best modeled by a probability function. However, Lassiter says relatively little about the nature of this probability function or the set of worlds it is defined over. One contribution of my framework is to show how a probabilistic approach to certainty can be integrated with a Hintikka semantics, delivering the quantifier-over-worlds approach in the special case where the degree of certainty is 1.

will in turn obtain iff p holds in all of A 's S -alternatives. Likewise, *It is pos certain that p is true* iff p has the maximal degree of epistemic certainty (for some contextually supplied agent), which will obtain iff p holds in all of the agent's E -alternatives. Lesser degrees of certainty will correspond to lower probabilities. *It's 95% certain that the Mets will win* means that the epistemic probability that the Mets will win is .95. *Sal is fairly certain that the Mets will win* means that Sal has a fairly high credence that the Mets will win.²⁷

This integration retains the advantages of both approaches. It also yields downstream benefits. For example, it explains why multi-premise closure is valid for pos form certainty ascriptions but not for graded certainty ascriptions. Similarly, it explains why graded epistemic certainty ascriptions are not factive, unlike their pos form counterparts: *It is 99% certain that the Mets will win* $\not\Rightarrow$ *the Mets will win*. After all, no probability shy of 1 guarantees truth.

3 Is certainty scarce?

3.1 Scarcity: for and against

Now that we have an account of certainty on the table, let us turn to the worry that led many epistemologists to renounce the quest for certainty. The worry is that certainty is *scarce*: precious little of our knowledge ever rises to the level of certainty. On the face of it, our account seems to feed directly into this worry. After all, our account says something is only certain (full-stop) if it has the maximal degree of certainty. But this is a high bar, and it would seem that hardly any of our knowledge measures up. Take, for instance, my knowledge that Marseille is in France. Does this knowledge rise to the *maximal* degree of certainty, subjective or epistemic? It is natural to think the answer is 'No'. After all, I can imagine scenarios in which this belief is mistaken—for example, scenarios in which I am the victim of an elaborate geographic hoax. But this seems to entail that this belief isn't as certain as, say, the cogito or basic logical truths.

But perhaps we shouldn't be so quick. After all, if asked, I'd readily assert both:

- (22) I'm certain that Marseille is in France.
- (23) It's certain that Marseille is in France.

More generally, people are fairly liberal in their certainty ascriptions: they don't reserve *certain* for a tiny sliver of their knowledge.

Thus, our everyday certainty ascriptions count against the idea that certainty is scarce. Of course, some might simply insist that most of these ascriptions are false—a line taken

²⁷A technical point: I require that the probability functions are regular in order to derive the quantifier-over-worlds analysis of pos form certainty ascriptions as a special case of the scalar semantics. Some may object that this leads to implausible consequences when the E and S -alternatives include uncountably many possibilities. For example, it would appear to entail that it is impossible for a infinitely fine dart to fall on a particular point on the number line (Hájek 2003). The issues here are complex, and fall beyond the scope of this paper. (See Lewis 1980; Easwaran 2014 for discussion.) However, note that it sounds quite odd to claim, *p is 100% likely, but not completely certain*. This suggests that our ordinary concept of certainty—the concept reflected in our everyday linguistic intuitions—takes the underlying probability function to be regular, even if this requirement leads to difficulties when it comes to infinitely fine darts and the like.

by Unger (1971, 1975). But this seems like a rather desperate and undesirable maneuver. *Ceteris paribus*, it would be preferable to find a way to make sentences such as (22) and (23) come out true.

3.2 Maximality without scarcity

We can reconcile the thesis that *certain* is a max adjective with the truth of (22) and (23) by relativizing gradable expressions to contextually determined standards of precision.²⁸ To illustrate with a different max adjective, consider again *straight*. In any context, something only qualifies as *straight* if it has the maximal degree of straightness, as revealed by the oddity of (16) (*This line is straight, but it could be straighter*). Still, it seems there is considerable contextual variability in what we regard as *straight*. Some contexts call for strict standards. If we are building a satellite, a microscopic dent might preclude an antenna from qualifying as *straight*. Other contexts are more lax. If we are repairing my television, I may be happy to call an antenna *straight* provided it is not noticeably bent.

One way to develop this idea is to allow the measure function denoted by a gradable adjective to vary with context. In a context with lax standards, *straight* denotes a coarse-grained measure function—one that maps x to the maximal degree of straightness as long as x is free from any noticeable bends. In a context with strict standards, *straight* denotes a fine-grained measure function—one that only maps x to the maximal degree of straightness if x is free from the tiniest dent.

This contextualist maneuver extends smoothly to *certain*. In a context with lax standards, *certain* denotes a function that maps much of an agent's knowledge to the maximal degree of certainty (subjective or epistemic). In stricter contexts, *certain* denotes a function that allows far fewer propositions to qualify as maximally certain.²⁹

If the subjective and epistemic certainty functions vary with context, then so too do the sets of worlds they are defined over. How does this work? One option is to suppose that a contextual standard of precision determines a set of *relevant alternatives*: a set of possibilities that are worth taking seriously, for the purposes of the conversation. These are the worlds that are not *too* distant or far-fetched, where what counts as *too* distant or far-fetched is a function of context.³⁰ We could then use the relevant alternatives to restrict the E and S -alternatives: the contextually restricted E -alternatives are the contextually relevant alternatives that are consistent with what's epistemically certain, and similarly, *mutatis mutandis*, for the S -alternatives. The context-relative epistemic and subjective certainty functions are probability functions defined over the contextually restricted E and S -alternatives. (Equivalently, a context-relative certainty function is what you get from conditionalizing a context-independent certainty function on the proposition that none of the contextually irrelevant alternatives obtain.)

To illustrate, take (22) and (23). In ordinary contexts, far-fetched scenarios in which I'm the dupe of an elaborate Marseillan deception are irrelevant. In such contexts, the

²⁸Lewis (1979) sketches a response to Unger along these lines. For recent work on standards of precision, see Sauerland and Stateva 2007; van Rooij 2011; Sassoon and Zevakhina 2012.

²⁹It matters little, for my purposes, *which* context is relevant. It could be the context of utterance, or the context of assessment, or that of the subject of the certainty ascription.

³⁰For relevant alternatives accounts of knowledge, see Dretske 1970, Goldman 1976, and, esp., Lewis 1996.

proposition *Marseille is in France* does have the maximal degree of epistemic and subjective certainty: it obtains in all of contextually relevant worlds compatible with what is epistemically and subjectively certain. However, when we contemplate various deception scenarios, we expand the sphere of relevant alternatives.³¹ Relative to this new context, *Marseille is in France* does not hold throughout all of the contextually restricted *E* and *S*-alternatives. And so, relative to this new context, it does not qualify as maximally certain.

By going contextualist, we block the conclusion that all of our ordinary certainty attributions are false. At the same time, we preserve the advantages of the semantic framework developed in §2. First, we still explain the data that led us to classify *certain* as a max adjective. After all, pos form certainty ascriptions still require the maximal degree of certainty, it's just that now whether something qualifies as maximally certain depends on context. And so in any context, an utterance of e.g., (19) (*It's certain to rain, but it could be more certain*) is predicted to be infelicitous. Second, we still capture the connections between certainty, knowledge, and belief that motivated the quantifier-over-worlds aspect of our approach. To do so, we need only maintain that in any context, the *E*-alternatives include the *K*-alternatives, and the *S*-alternatives include the *B*-alternatives.³²

3.3 Taking stock

I've argued that we should resist two impulses: the impulse to analyze certainty in terms of knowledge, and the impulse to dismiss certainty as unattainable. According to the treatment of certainty offered here, certainty comes in two forms: subjective and epistemic. The former consists in a strong conviction; the latter consists in a strong epistemic position, not reducible to knowledge. And while *certain* is a max adjective, this does not entail that certainty is scarce: in many contexts, a non-negligible subset of our everyday knowledge qualifies as both subjectively and epistemically certain.

That said, the worry that certainty is unattainable contains a kernel of truth. After all, typically philosophers became convinced that certainty is unattainable only after reflecting on far-fetched error possibilities involving demonic deception and the like. As they engaged in these reflections, they expanded the sphere of relevantly close worlds to include many that would normally be excluded. In doing so, they created a context where their exclamations of *Little to nothing is certain!* may well have been true. But from this it doesn't follow that when ordinary agents, occupying ordinary contexts, pursue certainty, they are embarking on a fool's errand.

By showing that certainty is often attainable (at least by ordinary standards), we've paved the way for putting certainty to work in epistemological theorizing. In what follows, I explore two specific applications: evidential probability (§4) and epistemic modals (§5). My main contention will be that a Certainty First approach to these topics allows us to

³¹In doing so, we exploit some version of Lewis' 'Rule of Attention' (1996: 559): attending to some possibility tends to render it relevant. For discussion and refinement of this rule, see Blome-Tillman 2009.

³²This contextualist picture of certainty ascriptions pairs naturally with a contextualist picture of knowledge and belief ascriptions: just as the *E* and *S*-alternatives expand or shrink with context, so too do the *K* and *B*-alternatives. But while a natural partnership, this is not forced. In theory, one could hold that *K* and *B*-alternatives remain fixed, while the *E* and *S*-alternatives sway. (Perhaps the *K* and *B*-alternatives are limiting cases of the *E* and *S*-alternatives; they coincide only in the laxest contexts.)

account for a range of data that would otherwise be left unexplained.

4 Evidence and evidential probability

The notion of evidence plays a vital role in both traditional and formal epistemology. But what does it take for an agent to have some proposition as part of their evidence? Epistemologists in the Bayesian tradition typically don't say much on this point. One of the central contributions of Knowledge First epistemology is to try to fill this lacuna. According to Williamson, someone possesses a proposition as evidence iff they know it:

E = K: For any agent A and time t , A 's total evidence at $t = \{p \mid A \text{ knows } p \text{ at } t\}$.³³

Williamson goes on to use this proposal as the backbone of a theory of evidential probability (2000: chp.10). On the resulting theory, the evidential probability of a proposition is its probability given what one knows.

In this section, I argue that evidence and evidential probability are intimately connected with epistemic certainty. These connections are difficult to explain on a Knowledge First account, but are readily explained by a Certainty First alternative.

4.1 Evidence

Judgments about evidence possession are closely bound up with judgments about certainty. Plausibly, if p is epistemically certain, then one's evidence entails p . Note how odd it would be to claim:

- (24) ?? It's certain that smoking causes cancer. But the evidence leaves open the possibility that smoking doesn't cause cancer.

By itself, this is no trouble for $E = K$, given the plausible assumption that epistemic certainty entails knowledge. The trouble begins when we note that the converse seems equally plausible: if p is entailed by the evidence, then p is epistemically certain. To motivate this, note that the following sound equally odd:

- (25) ?? The medical evidence entails that smoking causes cancer. But it isn't quite certain that smoking causes cancer.

Earlier we found reason to doubt that knowledge entails epistemic certainty. If it doesn't, then $E = K$ has trouble accounting for these data. After all, we should expect (25) to describe a perfectly coherent situation, one where the medical community's knowledge falls short of certainty.

Defenders of $E = K$ might seek to explain the data by appealing to Williamson's suggestion that we are reluctant to let the "contextually set standards for knowledge and certainty diverge" (2000: 204). On this view, while knowledge does not entail either epistemic or subjective certainty, a knowledge ascription will typically be true in a context c only if the corresponding certainty ascriptions are also true in c .

³³See Williamson 1997, 2000: chp.9.

However, the same considerations that suggest knowledge does not entail certainty cast doubt on the idea that we're reluctant to let the standards for knowledge and certainty diverge. As we saw in §2.2, claiming that someone knows something with certainty is not redundant; rather, it's naturally interpreted as claiming that they know it with a particularly high degree of certainty. Moreover, we saw that ordinary speakers are often happy to speak of knowledge that falls short of certainty, as revealed by (5)-(8). These considerations suggest that the standards for certainty ascriptions are typically higher than those for knowledge ascriptions.

From the perspective of the present essay, there is a natural remedy for this difficulty. The remedy is to identify one's evidence with one's epistemic certainties rather than one's knowledge. Of course, if epistemic certainty ascriptions are context-sensitive, then this leads to a contextualist account of evidence possession ascriptions:

E = C: In any context, the expression *A's evidence* is co-extensive with the expression *A's epistemic certainties*.

Some might balk at the idea that evidence possession claims are context-sensitive in this way. However, I think our ordinary patterns of *evidence*-talk actually fit quite nicely with a contextualist treatment. Suppose someone asks the detective, *What's your evidence that the butler did it?* In many contexts, it would be natural to cite the fact:

(26) The cook saw the butler fleeing the scene, weapon in hand.

Suppose, however, the questioner raises the possibility that the cook's eyesight is unreliable. If the detective is willing to take this possibility seriously, it would be natural for the detective to admit that, strictly speaking, she doesn't have (26) as a part of her evidence. Rather, she has:

(27) The cook thought he saw the butler fleeing the scene, weapon in hand.

And we can imagine continuations of the conversation in which the detective begrudgingly admits that not even (27) is part of her evidence. For example, we can imagine a context in which she seriously entertains the possibility that the cook is lying, or in which she starts to worry about whether all her experiences are a demon-induced deception. This is precisely the sort of contextual variability in our judgments about evidence possession that we should expect if $E = C$ is correct.³⁴

³⁴In addition to jiving with everyday *evidence*-talk, contextualism about evidence possession comes with epistemological benefits. As Greco (2017) observes, such a view allows us to reconcile a Bayesian account of evidential support with the view that all evidence is defeasible. The idea is that in any particular context, one's evidence cannot be undermined by new evidence (since if $Pr(p | E) = 1$, then for any further evidence E' , $Pr(p | E \wedge E') = 1$). However, receipt of new evidence often functions to change the context, which changes what counts as evidence. Relative to the new context, what was formerly one's evidence is no longer immune to defeat. For the purposes of this paper, we need not take a stand on whether all evidence is defeasible, or whether this is the best way of modeling defeat. However, those who are attracted to a 'contextualist foundationalism' along these lines will find a natural ally in $E = C$.

4.2 Evidential probability

The evidential probability of a proposition p is standardly defined as the probability of p given the evidence. If $E = C$, then this will be the probability of p given what is certain. More precisely: for any context c , the evidential probability of p is the probability of p conditional on whatever propositions qualify as epistemically certain in c .

This allows us to unify evidential probabilities and degrees of certainty. Recall that our semantics for *certain* appealed to an epistemic certainty function (EC), which we took to be a probability function. Given $E = C$, we can now venture a further hypothesis: the evidential probability function simply *is* the epistemic certainty function. Call this the ‘Certainty Account of Evidential Probability’:

Certainty Account of Evidential Probability: The evidential probability of p (relative to a contextual standard s) is p ’s degree of epistemic certainty (relative to s).

In what follows, I highlight two considerations in favor of the Certainty Account of Evidential Probability. The account explains linguistic data suggesting a close connection between evidential probability ascriptions and epistemic certainty ascriptions. It also explains the normative connections between evidential probabilities and credences.

4.2.1 Certainty and probability: linguistic data

While *evidential probability* is something of an epistemologist’s term of art, it maps onto an intuitive notion. This intuitive notion is reflected in our everyday use of probability operators, e.g.:

(28) It’s likely/probable that the Mets will win.

In everyday discourse it may not always be clear what sort of probability is at issue. However, at least some uses of (28) seem to convey a distinctly evidential notion of probability. Such evidential readings can be made explicit using *in view of*-phrases, e.g.:

(29) In view of the evidence, it’s likely/probable that the Mets will win.

These evidential probability ascriptions are closely connected to epistemic certainty ascriptions. Both accept percentage modifiers (e.g., 99%). And when both are embedded under the same percentage modifier, they seem to be equivalent:

- (30) a. It’s 99% likely that the Mets will win.
b. It’s 99% certain that the Mets will win.

(30a) and (30b) seem interchangeable, at least when (30a) is interpreted in terms of evidential probability. Indeed, it would be quite odd to affirm one while denying the other:

(31) # It’s 99% certain/likely the Mets will win. But it’s only 98% likely/certain that they’ll win.

To my ears, this would only be coherent if we impose some non-evidential interpretation on the probability operators.³⁵

This cries out for explanation. The Certainty Account of Evidential Probability provides one.³⁶ By contrast, it is far less clear how to account for this data if evidential probabilities are probabilities conditional on knowledge, unless we take knowledge and epistemic certainty to be co-extensive.

Some might worry that I've cherry-picked my data. According to this objection, the equivalence between graded epistemic certainty ascriptions and graded evidential probability ascriptions only holds for percentage modifiers that denote very high degrees on the corresponding scale. But when we look at mid-scale percentage modifiers (e.g., 60%) the equivalence breaks down:

- (32) a. It's 60% likely the Mets will win.
b. ? It's 60% certain the Mets will win.

In a situation where there's a 60% chance the Mets will win, (32a) seems to be in perfectly good order, while (32b) does not.

But I think it would be too hasty to abandon the Certainty Account of Evidential Probability on these grounds alone. First, observe that the intuition here is not really one of *inequivalence*: it's not that there's a situation in which (32a) is true, whereas (32b) is false. Rather, it's that claiming something is *60% certain* just sounds odd. Moreover, there is a natural explanation for this oddity. We've already seen that a max adjective targets the upper end of its scale. (32b) shows that our tendency to reserve *certain* for the upper end of its scale persists even under degree modifiers: we are happier with *99% certain* than *70% certain*, and less happy still with *60% certain*.

By contrast, *probable* and *likely* are relative gradable adjectives: they do not target the upper end of the probability scale, but rather some contextually determined point along it (Yalcin 2010: 930; Lassiter 2010: 204). After all, if there's a 70% chance that it will rain today and an 80% chance that it will rain tomorrow, one can say:

- (33) It's likely to rain today, but it's more likely that it will rain tomorrow.

Since *likely* and *probable* have no tendency to target the upper end of their scale, they happily combine with mid-scale percentage modifiers. This explains the contrast in (32)—an explanation that's fully compatible with the Certainty Account of Evidential Probability.³⁷

³⁵Lassiter (2017: chp. 5) makes a similar observation. He gives the example of a CNN broadcast on the 2014 AirAsia crash, where multiple speakers seem to use *95% certain* and *95% likely* interchangeably when discussing whether some debris was part of the airliner.

³⁶Of course, the explanation assumes that the meaning of our ordinary probability talk is best modeled by a probability function—an assumption that is by no means trivial. However, much recent work on the semantics of probability operators seems to vindicate this assumption; see, among others, Yalcin 2010; Lassiter 2010, 2017; Moss 2015.

³⁷There are two ways of developing this explanation in more detail. One possibility is that *certain*'s distaste for mid-scale percentage modifiers is an instance of a very general phenomenon involving max adjectives: namely, that they do not tolerate modifiers that lower the degree on the associated scale significantly below the maximum. Another possibility is that this is an instance of a more specific phenomenon involving max adjectives that co-habit a scale with a relative gradable adjective. According to this possibility, we prefer,

4.2.2 Constraints on credences

A second argument in favor of the Certainty Account of Evidential Probability comes from considering the normative relation between evidential probability and credence. At least as far back as Locke's *Essay*, philosophers have been attracted to the view that rationality requires one to proportion one's degree of belief to the evidence.³⁸ This suggests what we can call the 'Credal Constraint': an agent's credence in p should equal the probability of p given her evidence.

The Credal Constraint fits very naturally with the Certainty Account of Evidential Probability. Earlier, we suggested that there is a normative connection between subjective and epistemic certainty. Plausibly, this normative connection also extends to degrees of certainty. That is:

Matching Requirement Relative to any context, your degree of subjective certainty in p should equal the degree to which p is epistemically certain.

This requirement has considerable appeal. It seems quite odd to claim that p is $n\%$ epistemically certain while denying that one is $n\%$ subjectively certain of p :

(34) ? It's 99% certain the Mets will win. But I'm $\left\{ \begin{array}{l} 100\% \\ \text{only } 98\% \end{array} \right\}$ certain that they'll win.

The Matching Requirement explains this oddity. (34) is infelicitous for the same reason as (3): no one could truly assert it while adhering to the requirements of rationality.

By contrast, the Knowledge First account of evidential probability proves harder to integrate with the Credal Constraint. On the Knowledge First account, evidential probabilities are probabilities conditional on what's known, and so everything one knows gets assigned probability 1. Given the Credal Constraint, it follows that one should have credence 1 in everything one knows, which is tantamount to the claim that one should be subjectively certain of everything one knows. But we have already seen reason to reject this claim. Recall our unconfident examinee, who knows that Queen Elizabeth I died in 1603, without being certain of it. As we saw in §2.2, it seems we can develop the case in such a way that his lack of subjective certainty is perfectly rational.

More generally: cases of knowledge without subjective certainty are not *ipso facto* cases of irrationality. This observation is difficult to reconcile with the Knowledge Account, whereas it is predicted by the Certainty Account.³⁹

when possible, to combine a degree modifier M with an adjective A when MA denotes a degree that is close to $\text{pos } A$. Since *certain* and *probable* are scale-mates, this leads us to prefer 60% *likely* to 60% *certain*. I am inclined to favor the latter hypothesis, since it seems that some max adjectives, such as *full*, do tolerate mid-scale percentage modifiers. (Claiming the tank is 60% *full* is in perfectly good order.) However, for present purposes we need not choose between these two possibilities.

³⁸Williamson also endorses this requirement (2000: 223).

³⁹Kaplan (2003, 2009) also objects to the consequence that one should have credence 1 in everything one knows. However, Kaplan's objection assumes that if one has credence 1 in p , one is rationally required to accept a bet wherein one gains a penny if p is true, and loses one's life otherwise. This leaves open a potential escape route, which is to simply deny this assumption (Williamson 2009). By contrast, my formulation of the difficulty does not assume any connection between credence 1 and life-in-the-balance bets.

5 Epistemic modals

The language of probability is part of a richer fragment of the language: modal discourse. We not only talk about whether something is *likely* to be the case, we also talk about whether something *might* or *must* be the case. If certainty is closely connected with evidential probability, we should also expect certainty to be closely connected with other modal expressions. In this section, I'll argue that this is precisely what we find.

5.1 Two analyses of epistemic modals

The standard analysis of epistemic modals takes them to be quantifiers over the possibilities compatible with some epistemic state. Possibility modals (denoted ' \diamond ') such as *might* and *possibly* are analyzed as existential quantifiers. Necessity modals (denoted ' \square ') such as *must*, *has to*, and *necessarily* are analyzed as universal quantifiers:

Classical Analysis of Epistemic Modals:

$\lceil \diamond p \rceil$ is true at a point of evaluation i iff p is compatible with the relevant epistemic state.

$\lceil \square p \rceil$ is true at i iff the relevant epistemic state entails p .

Much of the recent literature on epistemic modals focuses on *whose* epistemic state is relevant. Is it some agent(s) determined by the context of utterance, as contextualists maintain? Or is it some agent(s) determined by the context of assessment, as relativists insist? For the purposes of this paper, I want to set this question aside and focus on what *sort* of epistemic state is relevant.

The most common view is that the relevant state is knowledge—call this the 'Knowledge Analysis.'⁴⁰ I propose instead that the relevant state is epistemic certainty:

Certainty Analysis of Epistemic Modals:

$\lceil \diamond p \rceil$ is true relative to a contextual standard s iff p is compatible with what's epistemically certain (relative to s).

$\lceil \square p \rceil$ is true relative to s iff p is entailed by what's epistemically certain (relative to s).⁴¹

In what follows, I offer novel linguistic data demonstrating a close connection between epistemic certainty and epistemic modals. The Certainty Analysis explains these data; the Knowledge Analysis does not.⁴²

⁴⁰Versions of the Knowledge Analysis are defended by Hacking (1967); Kratzer (1981); DeRose (1991); Egan et al. (2005); Egan and Weatherson (2011); Stanley (2005); Stephenson (2007); Dorr and Hawthorne (2013).

⁴¹Here's one way of fleshing this out formally. For any contextual standard of precision s , let $E_s(w)$ be the s -restricted E -alternatives—that is, the s -relevantly close worlds compatible with what is epistemically certain for the relevant agent(s). Then:

$$\lceil \diamond p \rceil^s = \{w \mid \exists w' \in E_s(w) \ \& \ w' \in p\}; \quad \lceil \square p \rceil^s = \{w \mid \forall w' (w' \in E_s(w) \rightarrow w' \in p)\}.$$

⁴²Only a couple of authors have entertained something like the Certainty Analysis. DeRose suggests that *might* is the dual of *it is certain that* (1998; 2009: 20). However, he seems to think certainty should be analyzed in terms of knowledge, indicating that he doesn't take this approach to be an alternative to the Knowledge Analysis, which he explicitly endorses in DeRose 1991. As far as I'm aware, the only author who defends a Certainty Analysis as an alternative to the Knowledge Analysis is Littlejohn (2011). According to Littlejohn, p

5.2 In favor of the Certainty Analysis

Suppose a detective asserts:

(35) The butler must/has to have done it.

We'd expect her to also be willing to assert:

(36) It's certain that the butler did it.

Indeed, it sounds very odd to follow an assertion of (35) with a denial of (36):

(37) ?? The butler must/has to have done it. But it's not certain that the butler did it.⁴³

We find a similarly close connection between epistemic possibility modals and certainty ascriptions. In particular, $\neg\Diamond p$ seems to entail that it is certain that $\neg p$, as suggested by the oddity of saying:

(38) ?? There's no possibility that the cook was involved. But it isn't certain that the cook wasn't involved.

The Certainty Analysis explains these observations. According to the Certainty Analysis, (35) says that it's epistemically certain that the butler did it. And so (37) is predicted to be contradictory. Similarly, the first conjunct of (38) entails that it's epistemically certain that the cook was not involved, which contradicts the second conjunct.

By contrast, the Knowledge Analysis leaves these data unexplained. According to the Knowledge Analysis, (35) says that the relevant agents' knowledge entails the butler did it. But if knowledge does not entail certainty, the assertability of (35) provides no guarantee that (36) is assertable. Similarly, we'd expect (37) and (38) to be coherent: they will be true whenever the relevant knowledge falls shy of certainty.

Could proponents of the Knowledge Analysis explain these data on pragmatic rather than semantic grounds? One way of doing so would be to appeal to the idea that certainty is the norm of assertion. According to this explanation, while (37) and (38) could

is epistemically possible for S iff $\neg p$ is not obviously entailed by something S knows with certainty. While my proposal differs from Littlejohn's in points of detail, in large part this section can be seen as providing new data in support of Littlejohn's thesis, and embedding this thesis within a broader Certainty First framework.

⁴³Some might question this judgment. After all, people do say things like:

(i) No one heard from him after 1995. He must be dead, but we can't know for sure whether he is.

However, it's natural to hear the final conjunct (*We can't know for sure . . .*) as hedging the *must*-claim. Thus a natural diagnosis is that the speaker is raising the standards for certainty mid-utterance. On this diagnosis, any apparently felicitous instances of (i)/(37) are due to mid-discourse context shifts.

For those who are unpersuaded by this response, two further points are worth noting. First, insofar as (i) causes a problem for the Certainty Analysis, close variants of it cause a problem for the Knowledge Analysis. (If we drop the 'for sure' in the final conjunct the discourse still sounds felicitous.) Second, the problem seems to be specific to the word, *must*. If we replace *must* in (i) with some other necessity modal, such as *has to*, *necessarily*, or *There's no possibility that not*, then the discourse becomes much worse. Thus another possible conclusion is that *must* is not a genuine necessity modal after all, but instead means something like, *highly probable* (Lassiter 2016). Note that this is still compatible with a broadly Certainty First framework, as long as probabilities are themselves understood in terms of degrees of certainty.

be true, anyone who uttered them would be violating the norm of assertion. However, I suspect this explanation would be unwelcome to most Knowledge Firsters, since it would involve replacing a central tenet of the Knowledge First program (the knowledge norm of assertion) with a certainty norm. While I'll be arguing shortly that such a replacement is independently motivated, it's a concession that most Knowledge Firsters would be reluctant to make.

Moreover, even if proponents of the Knowledge Analysis are willing to make this concession, trouble is still in store. This is because the connection between certainty and modals persists in embedded contexts:

- (39) # Suppose both that there's no possibility that the cook was involved and it's not certain that the cook wasn't involved.
- (40) # If the butler must have done it and it's not certain whether he did it ...⁴⁴

A merely pragmatic explanation of the infelicity of (37)-(38) does not generalize to explain the oddity of (39)-(40). By contrast, the Certainty Analysis has no trouble here. According to the Certainty Analysis, (39) and (40) invite the addressee to entertain an incoherent state of affairs, thereby accounting for their infelicity.

Thus the infelicity of (37)-(40) provide compelling evidence for the Certainty Analysis. An independent source of evidence comes from the phenomenon of *modal concord*. Modal concord arises when two modals occur next to each other, but only seem to contribute the force of a single modal (Halliday 1970; Geurts and Huitink 2006; Huitink 2012). Compare:

- (41) a. You may possibly have read my little monograph on the subject.⁴⁵
 b. You may have read my little monograph on the subject.

The most natural reading of (41a) is a 'concord reading', on which it's simply equivalent to (41b) (or a slightly hedged version thereof). It is less natural to give (41a) a 'cumulative' reading, according to which it's possible that there is a possibility that the addressee has read the speaker's monograph.

It is a matter of debate how best to analyze modal concord. However, it's widely agreed that in order for a concord reading to be possible, the two modals must be equivalent. This explains why (41a) allows for a concord reading, but (42) and (43) do not:

- (42) ? You must possibly have read my monograph.
 (43) ? You may certainly have read my monograph.⁴⁶

To see why this supports the Certainty Analysis, note that *must certainly* allows for a concord reading (Huitink 2012). Here are some naturally occurring examples retrieved through the Corpus of Contemporary American English (Davies 2008-):

⁴⁴Yalcin (2007) argues that the behavior of epistemic modals in embedded contexts requires revising the classical semantics for epistemic modals. I discuss this issue in §5.3, arguing that even if Yalcin is right on this front, this is perfectly compatible with a revised version of the Certainty Analysis.

⁴⁵Geurts and Huitink (2006) take (41a) from Sir Arthur Conan Doyle's *The Hound of the Baskervilles*.

⁴⁶This is not to say that there are no coherent readings of these sentences. There may be a coherent reading of (43) on which it says that it is possible that it's certain that the addressee has read the speaker's monograph. However, this would be a cumulative reading rather than a concord reading.

- (44) Something about her told him that she must certainly be noble.
- (45) She simply observed the process, as though it were not her hand at all, but someone else's, despite the fact that the pain must certainly have been very great indeed.
- (46) And the plates on my Subaru station wagon back in New England must certainly be among the billions contained in private databases.

In each of these examples, it's natural to give *must certainly* a concord reading. For example, the most natural reading of (46) claims that it's epistemically necessary that the plates are contained in private databases, not that it's epistemically necessary that it's epistemically certain that the plates are in private databases. Those who reject the Certainty Analysis will thus be forced to reject the well-supported generalization that modal concord is only possible when both modals are equivalent.

5.3 An Objection to the Certainty Analysis

Some may object that the Certainty Analysis has trouble explaining the infelicity of *epistemic contradictions* (e.g., (47)) and *concessive knowledge attributions* (e.g., (48)):

- (47) # It's raining but it might not be raining.
- (48) # I know that it's raining. But it might not be.⁴⁷

According to the Certainty Analysis, *It might not be raining* is true as long as it's not certain that it's raining. Since knowledge does not require certainty, it seems that a speaker could both know that it's raining and also know that it's not certain (for her) that it's raining. And so it is unclear why (47) and (48) are infelicitous.

By contrast, the Knowledge Analysis seems to be well-positioned to explain these data (Stanley 2005). On the Knowledge Analysis, $\diamond\neg p$ entails that the relevant agents don't know p . If we assume the 'Speaker Inclusion Constraint', according to which the speaker is always one of the relevant agents (Egan et al. 2005), concessive knowledge attributions such as (48) are guaranteed to be self-contradictory. And while epistemic contradictions such as (47) are not predicted to be self-contradictory, proponents of the Knowledge Analysis arguably have a plausible pragmatic explanation of their infelicity. After all, many Knowledge Firsters hold that assertion is governed by a knowledge norm:

Knowledge Norm of Assertion (KA): Assert p only if you know p .⁴⁸

Combining the Knowledge Analysis with KA predicts that (47) is never assertable. Knowing the first conjunct (*It's raining*) precludes knowing the second (*It might not be*), so no one could assert (47) while abiding by KA. On the face of it, this is an elegant result, since a major argument for KA is that it explains the infelicity of Moorean assertions, e.g.:

- (49) # It's raining but I don't believe/know it's raining.

⁴⁷The label 'epistemic contradictions' is due to Yalcin (2007). The label 'concessive knowledge attributions' is due to Rysiew (2001).

⁴⁸See, in particular, Williamson 1996, 2000. Unger (1975), Slote (1979), DeRose (2002), Sutton (2007), and Kelp (2018) all endorse versions of KA.

Hence by appealing to KA, proponents of the Knowledge Analysis promise a unified explanation of the infelicity of epistemic contradictions and Moorean assertions.

There are two ways that proponents of the Certainty Analysis could respond to this objection. The first is to replace KA with a certainty norm of assertion:

Certainty Norm of Assertion (CA): Assert p only if p is epistemically certain for you.^{49,50}

Armed with CA, advocates of the Certainty Analysis can explain the infelicity of epistemic contradictions on pragmatic grounds. If the first conjunct of (47) is epistemically certain, then *It is not raining* is incompatible with what's epistemically certain. And so the second conjunct is false, hence not epistemically certain (since pos form epistemic certainty ascriptions are factive). Thus no one could assert (47) without violating CA. Moreover, since knowledge is factive, this explanation generalizes to explain the infelicity of concessive knowledge attributions.

Some may worry that CA is *ad hoc*. But it can be motivated on independent grounds. Stanley (2008), following Unger (1975), notes that it sounds odd to say:

(50) ?? It's raining but it's not certain that it's raining.

Given that knowledge doesn't entail epistemic certainty, KA doesn't explain the oddity of (50). CA does: it would be impossible to utter (50) while obeying CA. Hence (50) suffers from the same ailment as (47). Assuming that epistemic certainty entails knowledge, CA also accounts for the original Moorean assertions (e.g., (49)) that motivated KA.

Thus a certainty norm of assertion offers one promising way of explaining the infelicity of (47) and (48). An alternative strategy is to modify the Certainty Analysis to predict that such sentences are semantically defective. For example, we could recast the Certainty Analysis as a version of *update semantics* (Veltman 1996). According to update semantics, the meaning of a sentence is its ability to change an information state (a set of worlds representing some body of information). An atomic sentence such as *It's raining* updates an information state s by removing any not-raining worlds from s . By contrast, modals are tests on information states. *It might be raining* tests to see whether s contains at least one world where it isn't raining. If so, s passes the test, and is returned unscathed. If not, s crashes, returning the absurd information state. Similarly, *It must be raining* tests whether s contains only worlds where it's raining.

Where does certainty come in? We can reframe the Certainty Analysis as an account of information states. According to this proposal, an information state is simply a set of contextually restricted E -alternatives: it's the set of contextually relevant worlds compatible with what's epistemically certain for the relevant folks. This 'Updated Certainty Analysis' still accounts for the connection between epistemic modals and certainty that motivated our original analysis (§5.2). At the same time, it provides a *semantic* explanation of the

⁴⁹The idea for such a norm can be traced to Moore (1959), who claims that when I assert p , I imply that p is certain. However, Moore thought that knowledge entailed certainty. As far as I know, the only author who explicitly advocates a certainty norm as an alternative to KA is Stanley (2008).

⁵⁰As Stanley (2008) observes, a norm along these lines seems most plausible if the standards of certainty are taken to be set by the asserter's context. On this construal, the norm says that you should only assert p in a context c if p counts as epistemically certain in c .

infelicity of (47) and (48). This is because update semantics predicts that epistemic contradictions are semantically defective, in the sense that they are guaranteed to crash any information state (Veltman 1996; Gillies 2001). Consider (47) (*It's raining and it might not be*). Updating an information state with the first conjunct results in an information state that contains only worlds where it's raining. And so this information state is bound to fail the test imposed by the second conjunct. This explanation generalizes to explain the infelicity of concessive knowledge attributions: given the factivity of knowledge, (48) will also crash any information state.

Thus while the infelicity of epistemic contradictions and concessive knowledge attributions poses a *prima facie* hurdle for the Certainty Analysis, there are two natural strategies for explaining the data: one pragmatic, one semantic. Which of these strategies is preferable? On the one hand, we saw that a certainty norm of assertion can be motivated on independent grounds (specifically, its ability to explain the full range of Moorean assertions). On the other hand, Yalcin (2007) argues that epistemic contradictions are infelicitous in embedded contexts, unlike Moorean assertions. According to Yalcin, this creates trouble for purely pragmatic explanations of the infelicity of epistemic contradictions. This observation, if correct, speaks in favor of a semantic explanation of the data.

For present purposes, we need not choose between the two strategies. (It may even turn out that both strategies are required to account for the full range of data.) The important point is that they offer ample resources for warding off the main objection to analyzing epistemic modals in terms of certainty.

5.4 Taking stock

Our ordinary uses of epistemic modals suggest that they're closely tied to certainty. This motivates a Certainty Analysis, according to which epistemic modals quantify over the possibilities compatible with what's epistemically certain.

The Certainty Analysis also fits naturally with the treatment of evidential probability in §4. Both necessity modals and pos form epistemic certainty ascriptions are logically stronger than *probably* claims, which are in turn stronger than *might* claims:

- (51) a. It's certain the butler did it/The butler must have done it. \Rightarrow
 b. It's likely/probable the butler did it. \Rightarrow
 c. The butler might have done it.

On the picture that emerges, epistemic certainty ascriptions, epistemic modals, and expressions of evidential probability all reside on the same scale. Pos form epistemic certainty ascriptions and necessity modals target the top of the scale: both are used to indicate that a proposition is maximally certain (relative to the context). Probability operators live lower on the scale: they indicate that a proposition has a fairly high degree of epistemic certainty. Finally, epistemic *might* inhabits the bottom of scale: it indicates that a proposition isn't ruled out by what's epistemically certain.

6 Conclusion

Recent epistemology has tended to give short shrift to certainty. In this essay, I've mounted a rehabilitation campaign. The notion of certainty is worthy of attention in its own right; moreover, it can be enlisted into epistemological service. By analyzing evidential probability and epistemic modals in terms of certainty we can account for a wide range of data—linguistic and otherwise—that are left unexplained by rival approaches.

If I'm right, then it's natural to wonder whether we can use certainty to make progress on other issues in epistemology. Consider, for example, the normative constraints on action. We've already floated the idea that certainty is the norm of assertion. Is it the norm of action more generally? That is, should people only act on those propositions that are epistemically certain for them? Or consider the goal of inquiry: are there some contexts in which we take inquiry to have concluded not merely when we arrive at knowledge, but only when we arrive at certainty? I hope to pursue these questions in future research.

It is too early to tell whether a Certainty First program—or indeed any 'X First' program—should be embraced in full generality. However, I take the arguments in this paper to have provided grounds for cautious optimism. At the very least, I aim to have shown that certainty enjoys a high degree of explanatory priority: we can make considerable progress by taking certainty as our lodestar, and explaining other epistemological phenomena in terms of it.

References

- Thomas Aquinas. *Commentary on the Posterior Analytics of Aristotle*. Magi Books Inc., New York, 1970.
- D.M. Armstrong. Does Knowledge Entail Belief? *Proceedings of the Aristotelian Society*, 70:21–36, 1969.
- A.J. Ayer. *Language, Truth, and Logic*. Dover, New York, 1936.
- A.J. Ayer. *The Problem of Knowledge*. Macmillan, London, 1956.
- Michael Blome-Tillman. Knowledge and Presuppositions. *Mind*, 118:241–294, 2009.
- Mark Davies. The Corpus of Contemporary American English (COCA): 520 million words, 1990-present, 2008-. URL <https://corpus.byu.edu/coca/>.
- Keith DeRose. Epistemic Possibilities. *Philosophical Review*, 100:581–605, 1991.
- Keith DeRose. Simple 'Might's, Indicative Possibilities, and the Open Future. *Philosophical Quarterly*, 48:67–82, 1998.
- Keith DeRose. Assertion, Knowledge, and Context. *Philosophical Review*, 111(2):167–203, 2002.
- Keith DeRose. *The Case for Contextualism: Knowledge, Skepticism, and Context*. Clarendon Press, Oxford, 2009.
- Rene Descartes. Rules for the Direction of the Mind. In Adam and Tannery, editors, *Oeuvres De Descartes*. Librairie Philosophique, Paris, 1983.
- John Dewey. *The Quest for Certainty*. Minton, Balch, New York, 1929.
- Cian Dorr and John Hawthorne. Embedding Epistemic Modals. *Mind*, 122(488):867–913, 2013.
- Fred Dretske. Epistemic operators. *Journal of Philosophy*, 67:1007–1023, 1970.
- Kenny Easwaran. Regularity and Hyperreal Credences. *The Philosophical Review*, 123(1):1–41, 2014.
- Andy Egan and Brian Weatherson. Epistemic Modals and Epistemic Modality. In Egan and Weatherson, editors, *Epistemic Modality*. Oxford University Press, Oxford, 2011.

- Andy Egan, John Hawthorne, and Brian Weatherson. Epistemic Modals in Context. In Preyer and Peter, editors, *Contextualism in Philosophy*. Oxford University Press, Oxford, 2005.
- Jim Geeting. *Truckers and Troopers*. Tall Cop Productions, 2005.
- Bart Geurts and Janneke Huitink. Modal Concord. In *Concord Phenomena and the Syntax-Semantics Interface*. ESSLLI, 2006.
- Thony Gillies. A New Solution to Moore's Paradox. *Philosophical Studies*, 105:237–250, 2001.
- Alvin Goldman. Discrimination and Perceptual Knowledge. *Journal of Philosophy*, 73:771–791, 1976.
- Daniel Greco. Cognitive Mobile Homes. *Mind*, 126(501):93–121, 2017.
- Ian Hacking. Possibility. *Philosophical Review*, 76:143–168, 1967.
- Alan Hájek. What Conditional Probability Could Not Be. *Synthese*, 13(3):273–323, 2003.
- Michael Alexander Kirkwood Halliday. Functional Diversity in Language as Seen From a Consideration of Modality and Mood in English. *Foundations of Language*, 6:322–261, 1970.
- John Hawthorne. *Knowledge and Lotteries*. Oxford University Press, New York, 2004.
- John Hawthorne and Jason Stanley. Knowledge and Action. *Journal of Philosophy*, 105(10):571–590, 2008.
- Jaakko Hintikka. *Knowledge and Belief*. Cornell University Press, 1962.
- Jaakko Hintikka. Impossible Worlds Vindicated. *Journal of Philosophical Logic*, 4:475–484, 1975.
- Janneke Huitink. Modal Concord: A Case Study in Dutch. *Journal of Semantics*, 29(3):403–437, 2012.
- Mark Kaplan. Who Cares What You Know? *Philosophical Quarterly*, 53:105–116, 2003.
- Mark Kaplan. Williamson's Casual Approach to Probabilism. In Greenough and Pritchard, editors, *Williamson on Knowledge*. Oxford University Press, New York, 2009.
- Christoph Kelp. Assertion: A Function First Account. *Noûs*, 52(2):411–442, 2018.
- Chris Kennedy. The Semantics of Relative and Absolute Gradable Adjectives. *Linguistics and Philosophy*, 30:1–45, 2007.
- Chris Kennedy and Louise McNally. Scale Structure, Degree Modification, and the Semantics of Gradable Predicates. *Language*, 81:345–381, 2005.
- Peter Klecha. Positive and Conditional Semantics for Gradable Modals. In Guevara et al, editor, *Proceedings of Sinn und Bedeutung 16*, volume 2, pages 363–376. MIT Working Papers in Linguistics, 2012.
- Peter Klein. *Certainty: A Refutation of Scepticism*. University of Minnesota Press, Minneapolis, 1981.
- Angelika Kratzer. The Notional Category of Modality. In Eikmeyer and Rieser, editors, *Words, Worlds, and Contexts: New Approaches in Word Semantics*. W. de Gruyter, Berlin, 1981.
- Henry Kyburg. *Probability and the Logic of Rational Belief*. Wesleyan University Press, Middletown, CT, 1961.
- Daniel Lassiter. Gradable Epistemic Modals, Probability, and Scale Structure. In Li and Lutz, editors, *Semantics and Linguistic Theory (SALT)*, volume 20, pages 197–215. CLC Publications, Ithaca, New York, 2010.
- Daniel Lassiter. *Measurement and Modality: The Scalar Basis of Modal Semantics*. PhD thesis, NYU, 2011.
- Daniel Lassiter. *Must*, Knowledge, and (In)directness. *Natural Language Semantics*, 24(2):117–163, 2016.
- Daniel Lassiter. *Graded Modality*. Oxford University Press, Oxford, 2017.
- C.I. Lewis. *Mind and the World Order: Outline of a Theory of Knowledge*. Charles Scribners, New York, 1929.
- David Lewis. Scorekeeping in a Language Game. *Journal of Philosophical Logic*, 8(1):339–359, 1979.
- David Lewis. A Subjectivist's Guide to Objective Chance. In Jeffrey, editor, *Studies in Inductive Logic and Probability*, volume 2, pages 263–293. University of California Press, 1980.
- David Lewis. Elusive Knowledge. *Australasian Journal of Philosophy*, 74(4):549–567, 1996.
- Clayton Littlejohn. Concessive Knowledge Attributions and Fallibilism. *Philosophy and Phenomenological Research*, 83(3):603–619, 2011.

- Clayton Littlejohn. How and Why Knowledge is First. In Carter and Gordon, editors, *Knowledge First Epistemology*. Oxford University Press, Oxford, forthcoming.
- Aidan McGlynn. *Knowledge First?* Palgrave MacMillan, New York, 2014.
- Richard Miller. Absolute Certainty. *Mind*, 87:46–65, 1978.
- G.E. Moore. Certainty. In *Philosophical Papers*. Allen and Unwin, London, 1959.
- Sarah Moss. On the semantics and pragmatics of epistemic vocabulary. *Semantics and Pragmatics*, 8(5):1–81, 2015.
- Robert Pasnau. Epistemology Idealized. *Mind*, 122:987–1021, 2014.
- Robert Pasnau. *After Certainty: A History of Our Epistemic Ideals and Illusions*. Oxford University Press, Oxford, 2017.
- Colin Radford. Knowledge: By Examples. *Analysis*, 27:1–11, 1966.
- Baron Reed. Certainty. *The Stanford Encyclopedia of Philosophy*, winter 2011.
- Hans Reichenbach. *The Rise of Scientific Philosophy*. University of California Press, 1963.
- Carmen Rotstein and Yoad Winter. Total Adjectives vs. Partial Adjectives: Scale Structure and Higher Order Modifiers. *Natural Language Semantics*, 12:259–288, 2004.
- Bertrand Russell. *The Problems of Philosophy*. Henry Holt and Company, New York, 1912.
- Patrick Rysiew. The Context-Sensitivity of Knowledge Attributions. *Noûs*, 35(4):477–514, 2001.
- Galit Sassoon and Natalia Zevakhina. Granularity Shifting: Experimental Evidence from Degree Modifiers. *Proceedings of SALT*, 22:226–246, 2012.
- Uli Sauerland and Penka Stateva. Scalar vs. Epistemic Vagueness: Evidence from Approximators. *Proceedings of SALT*, 17, 2007.
- Duns Scotus. *Philosophical Writings*. Hackett, Indianapolis, 1987.
- M.A. Slote. Assertion and Belief. In Dancy, editor, *Papers on Language and Logic*. Keele University Library, Keele, 1979.
- Ernest Sosa. How to Resolve the Pyrrhonian Dialectic: A Lesson from Descartes. *Philosophical Studies*, 85:229–249, 1997.
- Jason Stanley. Fallibilism and Concessive Knowledge Attributions. *Analysis*, 65(2):126–131, 2005.
- Jason Stanley. Knowledge and Certainty. *Philosophical Issues*, 18(1):35–57, 2008.
- Tamina Stephenson. Judge Dependence, Epistemic Modals, and Predicates of Personal Taste. *Linguistics and Philosophy*, 30(4):487–525, 2007.
- Jonathan Sutton. *Without Justification*. MIT Press, Cambridge, MA, 2007.
- Peter Unger. A Defense of Skepticism. *The Philosophical Review*, 80(2):198–219, 1971.
- Peter Unger. *Ignorance*. Clarendon Press, Oxford, 1975.
- Robert van Rooij. Vagueness and Linguistics. In Ronzitti, editor, *The Vagueness Handbook*, pages 123–170. Springer, Dordrecht, 2011.
- Frank Veltman. Defaults in Update Semantics. *Journal of Philosophical Logic*, 25(3):221–261, 1996.
- Kai von Fintel and Thony Gillies. *Must... Stay... Strong!* *Natural Language Semantics*, 18(4):351–383, 2010.
- Brian Weatherson. Knowledge, Bets, and Interests. In Brown and Gerken, editors, *Knowledge Ascriptions*, pages 75–103. Oxford University Press, New York, 2012.
- Timothy Williamson. Knowing and Asserting. *Philosophical Review*, 105(4):489–523, 1996.
- Timothy Williamson. Knowledge as Evidence. *Mind*, 106(424):1–25, 1997.
- Timothy Williamson. *Knowledge and its Limits*. Oxford University Press, Oxford, 2000.
- Timothy Williamson. Replies to Critics. In Greenough and Pritchard, editors, *Williamson on Knowledge*. Oxford University Press, New York, 2009.
- A.D. Woozley. Knowing and Not Knowing. *Proceedings of the Aristotelian Society*, 53:151–172, 1952.
- Stephanie Wykstra. *A Defense of Cartesian Certainty*. Rutgers PhD dissertation, 2008.
- Seth Yalcin. Epistemic Modals. *Mind*, 116(464):983–1026, 2007.
- Seth Yalcin. Probability Operators. *Philosophy Compass*, 5(11):916–937, 2010.