New Work for Certainty

Bob Beddor

Forthcoming in *Philosophers' Imprint*

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 and assertion.

1 Introduction

For much of its history, epistemology focused on certainty. Philosophers such as Aquinas, Scotus, and Descartes all conceived of certainty—or *scientia*—as the epistemic ideal. Moreover, there was broad agreement on what this involved. In order for a belief to qualify as certain, it needed to be immune to both doubt and rational revision. For these philosophers, a central task of epistemology was to determine which of our beliefs could attain this exalted status.[1]

But these days, epistemologists have little time for certainty. Knowledge has stolen the spotlight.

To hear most epistemologists tell it, this shift from certainty to knowledge is an improvement. In our Cartesian youth we thought that certainty was attainable, but time has taught us better. After all, the reasoning goes, precious few of our beliefs are so secure that they cannot be doubted or rationally revised. Perhaps some logical truths and the cogito meet this high bar; perhaps not even these.[2] Knowledge, by contrast, is more abundant. Consequently, it's a better candidate to serve as the foundation for a useful epistemology.

This paper questions received wisdom on this front; I argue that certainty should occupy a central place in epistemology. In doing so, I do not advocate a return to some

[1] Aquinas 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.

[2] For prominent 20th century endorsements of the view that certainty is seldom, if ever, attained, see a.o. Russell 1912; Dewey 1929, C.I. Lewis 1929; Ayer 1936; Reichenbach 1963; Unger 1971, 1975.

sort of Cartesian naiveté. Contemporary epistemologists are right that very few beliefs are in principle immune to doubt or rational revision. However, I'll argue that the right response to this insight is not to reject certainty, but rather to reject the ultra-demanding conception of certainty that is so often assumed.

The first half of this paper develops a more plausible account of certainty (§§2-3). Attending to our everyday certainty attributions suggests that many humdrum facts can qualify as certain. A natural way of accommodating this is to opt for a contextualist treatment of certainty ascriptions. In some contexts, certainty is extremely difficult to attain; in others, the bar is lower. At the same time, I argue that we should resist assimilating certainty to knowledge. While certainty is often attainable, it is still a more demanding state than knowledge.

Equipped with a new account of certainty, the second half of the paper (§§4-5) puts this notion to explanatory work. Suppose we were to take certainty as our central notion and try to understand other epistemic phenomena in terms of it. How far could we get? Surprisingly far, it turns out. I'll argue that we can use certainty to illuminate evidence, evidential probability, and epistemic modals, as well as the normative constraints on credence and assertion. Moreover, the explanations that emerge have important advantages over more familiar 'Knowledge First' treatments of these topics.[3] The upshot: many of the epistemological jobs traditionally assigned to knowledge are better performed by certainty.

2 Towards an account of certainty

While the analysis of knowledge has spawned a massive literature, the analysis of certainty has received comparatively little contemporary attention. This section tries 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 take up in §3.

2.1 Psychological vs. epistemic certainty

It's common to distinguish between psychological and epistemic certainty.[4] Psychological certainty is a matter of strength of conviction. A belief can be psychologically certain even if it is 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 psychological certainty:

[3]The *locus classicus* of 'Knowledge First' epistemology is Williamson 2000. For important elaborations of the Knowledge First approach, see a.o. Hawthorne 2004; Hawthorne and Stanley 2008; Sutton 2007; Weatherson 2012; Littlejohn forthcoming.

[4]See e.g. Moore 1959; Klein 1981; Stanley 2008; DeRose 2009; Reed 2011.

(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 psychologically 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. Psychological 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.[5]

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?[6]

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

The difference between *knows* and *knows for certain* is not a quirk of English. A wide variety of languages carve out the same distinction. Here are some examples from Italian, Romanian, Bahasa Indonesian, Malayalam, Korean, and Japanese, respectively:

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

[6]Arguably, 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.

(5) So per certo che Ronaldo non giochera' la prossima partita
I know for sure that Ronaldo not will play the next game
'I know for sure that Ronaldo will not play the next game'

(6) Bine, dar stii tu sigur ca vine maine?
OK, but know you sure that she's coming tomorrow?
'OK, but do you know for sure she's coming tomorrow?'

(7) Tetapi anda tidak tahu dengan pasti.
But you do not know with certainty.
'But you do not know for certain.'

(8) enik:jə at̪ə orap:a:jiʈ:e arija:m.
to.me that solidly know.
'I know that for sure'

(9) na.nun pi-ga o.go-it'a-nun.kos-ul hwakʃɾ-i an-da.
I rain falling certain know.
'I know for certain that it's raining'

(10) [doko-ni iru-noka], [dou shi-teiru-noka]-o kakujitsu-ni shiru-tame-no houhou.
where be how do certain know-for method.
'methods for knowing with certainty where [they] are and how [they] are doing'[7]

In each of these languages, the counterpart of *knows for certain* picks out a stronger state than the counterpart of *knows*.[8]

Could we explain the difference in strength between *knows* and *knows for certain* 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 psychological and epistemic certainty ascriptions are also true. However, perhaps yoking *knows* and *certain* together in the complex phrase *knows for certain* drives up the standards for both knowledge and 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, *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.[9]

[7] Example taken from a blog post on how to get in touch with friends and family after a natural disaster: `http://shinsairegain.jp/2016/03/20/communicationline/`

[8] For these examples, and discussion of their interpretation, I am grateful to Carlotta Pavese, Mona Simion, Qu Hsueh, Savithry Namboodiripad, Jiseung Kim, Andrew Moon, and Mitcho Erlewine.

[9] As a referee points out, there may be some cases where combining two synonyms in a complex expression serves to drive up the standards—perhaps *full and complete answer* works this way. However, such examples tend to be fairly isolated; they also tend to have a conventional flavor. By contrast, *knows for certain* crops up in a wide variety of languages hailing from different language families, suggesting that it is not similarly conventional.

A further difficulty for a pragmatic explanation of the non-redundancy of *knows for certain*—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 (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':

(11) [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.[10]

(12) 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.[11]

(13) We now know with near-certainty that Russia did this with the goal of electing Trump.[12]

(11) explicitly ascribes knowledge without certainty. (12)-(13) ascribe knowledge that is nearly certain, implying that it is not actually certain.[13]

At this point, some may concede that knowledge does not require psychological certainty, while still maintaining that it requires epistemic certainty. However, this position stands in tension with the normative connection between psychological and epistemic certainty. As we've seen, it's natural to hold that one should be psychologically certain of p iff p is epistemically certain. If knowledge entails epistemic certainty, then anyone who knows p is rationally required to be psychologically certain that p. But this seems wrong.

[10] `http://www.smithers.co.uk/news_article.php?id=16&o=50`.

[11] Geeting 2005: 96

[12] `http://nymag.com/daily/intelligencer/2016/12/trump-mcconnell-putin-and-the-triumph-of-the-will-to-power.html`

[13] A referee raises the possibility that these expressions of uncertainty modify the content of the knowledge rather than the manner of knowing. For example, perhaps (11) and (12) should be paraphrased as:

(i) We know that it's not certain, but highly probable, that returns...

(ii) When [a false ID] is handed to a cop, he knows that it's nearly certain the guy before him is not...

However, compositional considerations count against this 'content-modifying' interpretation. The most straightforward compositional semantics for the prepositional phrases *without certainty* and *with near certainty* treats them as restrictive modifiers. In (11)-(13), these modifiers combine with the verb *knows*, suggesting that they characterize the manner of knowing. One advantage of this interpretation is that it allows us to analyze (11)-(13) in the same way we would analyze structurally similar sentences where the content-modifying interpretation gives the wrong results, e.g.:

(iii) The detective stated without hesitation that the butler did it.

Here, the prepositional phrase *without hesitation* does not modify the content of the detective's statement. (The detective is not claiming that without hesitation the butler did it.) Rather, it characterizes the manner in which the detective made the statement.

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 psychological or epistemic certainty.[14] On the picture that emerges, psychological 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, (12) (*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 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* as modal operators. For someone to know p is for p to hold in all worlds consistent with what they know—call these the '$\mathcal{K}$-alternatives'. For someone to believe p is for p to hold in all worlds consistent with what they believe—call these the '$\mathcal{B}$-alternatives'. Here the 'worlds' in question are not assumed to be metaphysically possible; instead, they can be viewed as maximally complete states of information.

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 $\mathcal{K}$-alternatives at w to include w. To capture the idea that knowledge asymmetrically entails belief, it's standard to take the $\mathcal{K}$-alternatives at w to include the $\mathcal{B}$-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 all the '$\mathcal{E}$-alternatives'—that is, all the worlds consistent with what is epistemically certain. To capture the idea that epistemic certainty requires a

[14] A residual worry: if knowledge doesn't require certainty, why does it sound odd to claim, *I don't know for certain p, but I do know p*? I'll defer this issue to §5, where I'll argue that a natural explanation is pragmatic. Epistemic certainty is the norm of assertion; since knowledge is factive, no one could make such a claim while abiding by the norm of assertion.

stronger epistemic position than knowledge, we require that the $\mathcal{K}$-alternatives are always a subset of the $\mathcal{E}$-alternatives, but not *vice versa*. (See Fig. 1.)

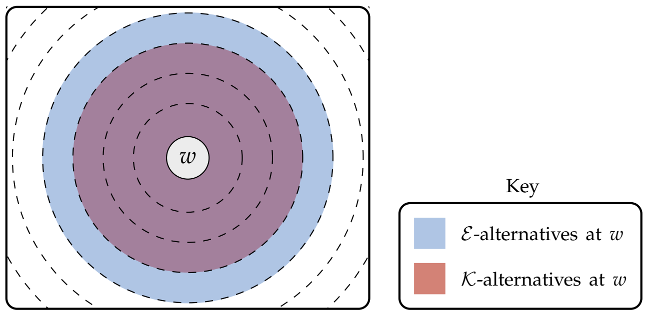

Figure 1: Knowledge & Epistemic Certainty

Formulated thus, the account is non-reductive: it does not try to explain epistemic certainy in more basic terms. It could, however, be supplemented with a reductive account of the $\mathcal{E}$-alternatives. For example, internalists could take the $\mathcal{E}$-alternatives to be the worlds consistent with the agent's phenomenal states. Reliabilists could take them to be 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; however, there is no need to take a stand on this matter here.

Even if we lack a reductive characterization of the $\mathcal{E}$-alternatives, we could perhaps use the $\mathcal{E}$-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 $\mathcal{K}$-alternatives at w as the $\mathcal{E}$-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.[15]

This framework can also be used to model the relation between belief and psychological certainty. Belief seems to be a weaker state than psychological certainty. Someone can believe that the butler did it without being psychologically certain of the butler's guilt, but not *vice versa*. This suggests the following picture: *psychological certainty is to belief as epistemic certainty is to knowledge.* Just as epistemic certainty asymmetrically entails knowledge, so psychological certainty asymmetrically entails belief. To model this, we can hold that someone is psychologically certain of p iff p holds in every world consistent with their psychological certainties—call these the '$\mathcal{P}$-alternatives'. To capture the asymmetric entailment between psychological certainty and belief, we require that the $\mathcal{P}$-alternatives are always a superset of the $\mathcal{B}$-alternatives.

A quantifier-over-worlds model of certainty yields a number of further predictions, two of which are worth mentioning. First, since the $\mathcal{E}$-alternatives include the $\mathcal{K}$-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

[15]Another option in the ballpark would be to help ourselves to a notion of comparative *normality* across worlds. We could then define the $\mathcal{K}$-alternatives at w as the $\mathcal{E}$-alternatives that are at least as normal as w (cf. Goodman and Salow 2018; Beddor and Pavese 2018). This would amount to a definition of knowledge in terms of normality and epistemic certainty.

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:

(14) 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:

(15) ? What I said was perfectly true. After all, I didn't say he did it. Only that it's *certain* that he did it.[16]

A second prediction of the quantifier-over-worlds approach is more controversial: both epistemic and psychological certainty are closed under logical entailment. This may seem implausible, particularly when it comes to psychological certainty. Every tautology is trivially entailed by one's certainties. But is everyone psychologically certain of every tautology?

However, we should note that this is a special instance of a more general problem: the problem of logical omniscience. Notoriously, epistemic and doxastic logics in the tradition of Hintikka (1962) predict that every agent knows every logical truth. By now much ink has been spilled over this problem; for our purposes, we need not take a stand on how best to deal with it. (Perhaps quantifying over impossible worlds will help (Hintikka 1975), perhaps not.) The important point is that using this framework to model epistemic and psychological certainty does not incur any new costs that were not already implicit in the framework.

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*.

Both psychological and epistemic certainty comes in degrees:

(16) It's fairly/very/95% certain the Mets will win.

(17) 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.[17]

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*,

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

[17] A 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).

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 psychological uses of *certain* with a *psychological certainty function* (*PC*) from propositions and agents to degrees on a *psychological certainty scale*, which measures degrees of 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. This yields a simple analysis of graded certainty ascriptions such as (16) and (17). On this analysis, degree modifiers combine with *certain* in the usual ways to deliver particular degrees on the psychological and epistemic certainty scales: *fairly (psychologically) certain* will deliver a fairly high degree of psychological 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 (18a) entails (18b), but not *vice versa*:

(18) 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 (18a) is:

(19) 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.[18]

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

(20) This line is long, but it could be longer.

(21) ? The line is straight, but it could be straighter.

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):

(22) This line is almost straight/#long.

(23) This line is completely straight/#long.

[18] For 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.

As a number of authors have noted,[19] *certain*—on both its psychological and epistemic uses—seems to pass the tests for a max adjective with flying colors:

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

(25) We're almost certain to lose.

(26) I'm/it's completely certain she'll be there.

So if we apply the standard scalar semantics to certainty ascriptions, we get the following picture. A pos form psychological certainty ascription, *A is* POS *certain that p*, is true iff *A* has the maximal degree of psychological 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).

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 (18a) asymmetrically entails (18b). 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.[20]

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 psychological 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 psychological and epistemic certainty functions. Let us now venture the following hypothesis: both are probability functions, defined over algebras generated from the $\mathcal{P}$-alternatives and the $\mathcal{E}$-alternatives, respectively. *PC* is a psychological probability function; *EC* an epistemic probability function. 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 psychological certainty, which will in turn obtain iff *p* holds in all of *A*'s $\mathcal{P}$-alternatives. Likewise, *It is* POS *certain that p* is true iff *p* has the maximal degree of epistemic certainty (for some contextually supplied

[19] See e.g., Unger 1975; Lassiter 2010, 2011, 2017; Klecha 2012.

[20] For 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.

agent), which will obtain iff p holds in all of the agent's $\mathcal{E}$-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.[21]

This integration retains the advantages of both approaches. It also yields downstream benefits. For example, 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.[22]

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, psychological 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. In ordinary contexts, I'd readily assert both:

(27) I'm certain that Marseille is in France.

(28) 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

[21] A 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 $\mathcal{E}$- and $\mathcal{P}$-alternatives include uncountably many possibilities. For example, it would appear to entail that it is impossible for an 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.

[22] A further benefit: this integration explains why *certain* is gradable, whereas *knows* is not (or only marginally so). The reason is that *certain* denotes a function from agents and propositions to degrees. By contrast, *knows* does not relate agents and propositions to degrees. Instead, it is a universal quantifier over a particular subset of contextually relevant worlds. In general, quantifiers over worlds are only marginally gradable. (Consider the oddity of claiming, *It very much must be raining*, or, *You somewhat ought to study*.)

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 (27) and (28) come out true.

3.2 Maximality without scarcity

We can reconcile the thesis that *certain* is a max adjective with the truth of (27) and (28) by relativizing gradable expressions to contextually determined standards of precision.[23] 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 (21) (*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 thought is to allow the function denoted by a gradable adjective to vary with context. In a context with lax standards, *straight* denotes a coarse-grained 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 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 (psychological or epistemic). In stricter contexts, *certain* denotes a function that allows far fewer propositions to qualify as maximally certain.[24]

If the psychological 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.[25] We could then use the relevant alternatives to restrict the $\mathcal{E}$- and $\mathcal{P}$-alternatives: the contextually restricted $\mathcal{E}$-alternatives are the contextually relevant alternatives that are consistent with what's epistemically certain, and similarly for the $\mathcal{P}$-alternatives. The context-relative epistemic and psychological certainty functions are probability functions defined over the contextually restricted $\mathcal{E}$- and $\mathcal{P}$-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

[23]Lewis (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.

[24]One interesting consequence of this approach is that credences are themselves context-sensitive. For independent arguments for this conclusion, see Clarke 2013. While our approaches are largely congenial, an important difference between our views is that Clarke takes this approach to support the idea that belief requires credence 1. One of the upshots of my paper is that what Clarke says about belief is much more plausible as a claim about psychological certainty.

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

irrelevant alternatives obtain.)

To illustrate, take (27) and (28). In ordinary contexts, far-fetched scenarios in which I'm the dupe of an elaborate Marseillan deception are irrelevant. In such contexts, the proposition *Marseille is in France* does have the maximal degree of epistemic and psychological certainty: it obtains in all of contextually relevant worlds compatible with what is epistemically and psychologically certain. However, when we contemplate various deception scenarios, we expand the sphere of relevant alternatives.[26] Relative to this new context, *Marseille is in France* does not hold throughout all of the contextually restricted $\mathcal{E}$- and $\mathcal{P}$-alternatives. And so it not longer qualifies 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., (24) (*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 $\mathcal{E}$-alternatives include the $\mathcal{K}$-alternatives, and the $\mathcal{P}$-alternatives include the $\mathcal{B}$-alternatives.[27]

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: psychological 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 psychologically and epistemically certain.

This contextualist conception of certainty differs in many respects from the traditional conception of *scientia* championed by philosophers in the medieval and early modern tradition. Still, there is an important thread of commonality. Both conceptions take certainty to require freedom from doubt. Philosophers such as Aquinas, Scotus, and Descartes imposed a particularly stringent version of this requirement: they took certainty to involve indubitability. While the contextualist does not go this far, the contextualist still maintains that in any context where p qualifies as certain, p cannot be seriously doubted. After all,

[26] In 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.

[27] Of course, contextualism is controversial in its own right. One source of difficulty comes from cross-contextual assessments: cases where a speaker makes a certainty attribution in one context, and an assessor inhabiting a different context evaluates this claim for truth or falsity. There is at least some temptation for the assessor to use the standards in their own context, rather than the speaker's. This is a complex issue, and one that I will bracket for the purposes of this paper. For those moved by this objection, one option is to move from contextualism to relativism; we could take the relevant 'contexts' to be contexts of assessment rather than contexts of utterance. For an overview and defense of relativism, see e.g., MacFarlane (2014).

entertaining doubts about p's truth will expand the sphere of relevant alternatives, thereby shifting the context. Relative to this new context, p will no longer qualify as certain.[28]

Equipped with a contextualist treatment of certainty, 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 by explaining these notions in terms of certainty, we can 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 $t = \{p \mid A \text{ knows } p \text{ at } t\}$.[29]

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-based analysis.

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:

(29) ?? 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 sounds equally odd:

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

[28] Cf. Greco (2017), who makes a similar point in defending a contextualist version of foundationalism. Greco distinguishes between a classic version of foundationalism that requires the epistemic foundations to be indubitable and a contextualist version that only requires the foundations to be *undoubted*.

[29] See Williamson 1997, 2000: chp.9.

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 (30) 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 psychological certainty, a knowledge ascription will typically be true in a context c only if the corresponding certainty ascriptions are also true in c.

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 (11)-(13). 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:

(31) 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 (31) as a part of her evidence. Rather, she has:

(32) 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 (32) 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 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.[30]

[30]See Greco 2017 for independent considerations in favor of a contextualist account of evidence possession.

4.2 Evidential probability

In addition to evaluating whether some hypothesis is consistent with—or entailed by—the evidence, we can also evaluate the probability of a hypothesis given a body of evidence. These evidential probabilities frequently have practical import. Suppose we've winnowed the suspect list down to two: either the butler or the gardener did it. We don't know for sure which is the culprit. But it's much more likely, given the evidence, that it was the butler. This probabilistic difference might well make a practical difference. For example, given our limited resources, it might be rational to focus on investigating the butler first.

Epistemologists in the Bayesian tradition have developed a rich formal apparatus for investigating evidential probabilities. In this tradition, the evidential probability of some 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 first is epistemological: the Certainty Account explains the normative connections between evidential probabilities and credences. The second is linguistic: the account explains linguistic data suggesting a close connection between evidential probability ascriptions and epistemic certainty ascriptions.

4.2.1 Certainty and probability: the normative link

What is point of positing evidential probabilities? What epistemological work do they serve? 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. This idea is taken for granted within much of the Bayesian tradition, where it's frequently assumed evidential probabilities constrain a rational agent's credences. Call this the 'Credal Constraint':

Credal Constraint Your credence in p should equal the probability of p given your evidence.

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

Matching Requirement Relative to any context, your degree of psychological 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*% psychologically certain of *p*:

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

The Matching Requirement explains this oddity. (33) 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 psychologically 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 psychological certainty is perfectly rational.

More generally, cases of knowledge without psychological 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.[31]

4.2.2 Certainty and probability: linguistic data

Our second argument for the Certainty Account is that it explains a range of 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.:

(34) 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 (34) convey a distinctly evidential notion of probability. Such evidential readings can be made explicit using *in view of*-phrases, e.g.:

(35) 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:

[31] Kaplan (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.

(36) a. It's 99% likely that the Mets will win.
b. It's 99% certain that the Mets will win.

(36a) and (36b) seem interchangeable, at least when (36a) is interpreted in terms of evidential probability. Indeed, it would be quite odd to affirm one while denying the other:

(37) ?? 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.[32]

This cries out for explanation. The Certainty Account of Evidential Probability provides one. 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:

(38) 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, (38a) seems to be in perfectly good order, while (38b) 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 (38a) is true, whereas (38b) is false. Rather, it's that claiming something is *60% certain* sounds odd. Second, we should be careful not to overstate its oddity. We can find naturally occurring examples of this construction, e.g.:

(39) FSU President Bernard Sliger said it is 60 percent certain the school will join a conference.[33]

What all of this suggests is that we should hold on to the idea that evidential probability ascriptions are equivalent to the corresponding epistemic certainty ascriptions. It's just that we prefer to avoid combining *certain* with mid-scale percentage modifiers. Why is this? Here's a natural thought. We've already seen that a max adjective targets the upper end of its scale. (38) 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

[32] Lassiter (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.

[33] `https://www.sun-sentinel.com/news/fl-xpm-1990-08-09-9002080020-story.html`

70% certain, and less happy still with *60% certain*. (Though, as (39) shows, even this is sometimes tolerated.)

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:

(40) 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, which is why (38a) is preferrable to (38b).

Some might question this explanation on the grounds that other max adjectives happily combine with mid-scale percentage modifiers. Take for example, *full*:

(41) The glass is 60% full.

But at least some speakers judge that (41) sounds fine—better, at any rate, than (38b).

However, there is an important difference between *certain* and *full*. Whereas *certain* shares its scale with the relative gradable adjectives *probable* and *likely*, no relative gradable adjective shares a scale with *full*. To leverage this observation into an explanation of our data, we can appeal to a principle along the following lines:

Competition Principle (CP) A combination of an absolute adjective α with a modifier m is dispreferred if both:

1. $m\alpha$ denotes a point very far from the value of POS α,
2. α has some scale-mate α', and the point denoted by $m\alpha$ could be denoted by $m'\alpha'$ (where m' is some modifier that may or may not be the same as m),

unless the combination $m'\alpha'$ is itself dispreferred by CP.

This explains why *60% certain* is dispreferred, whereas *60% full* is not. The point denoted by *60% certain* could equally well be denoted by *60% probable*, and the latter is not dispreferred by CP (since *probable* is a relative gradable adjective). However, *60% full* could not equally well be denoted by a combination of a modifier with a relative gradable adjective, so it is not dispreferred by CP.[34]

CP has independent explanatory appeal. First, it explains why it seems odd to describe a glass as *1% full*, or *only the slightest bit full*. After all, while *full* does not share a scale with a relative adjective, it does share a scale with another max adjective—namely, *empty*, whose maximal element is the minimal element of *full*. According to CP, the point denoted by *almost empty* is close to the value of POS *empty*, and hence should be preferred.

[34] While to my knowledge no other researchers have advanced CP, the basic idea behind such competition-based explanations is familiar. For example, competition explains why *Some students passed* implicates *Not all did* (since the more informative competitor *All passed* should have been preferred, if it were true). For competition-based explanations of other linguistic phenomena, see Aronoff (1976).

Second, CP explains similar contrasts between other absolute and relative adjectives that co-habit a scale. Consider the max adjective *filthy* and the relative adjective *dirty*. While neither of these adjectives accepts percentage modifiers, they pattern differently with minimizing modifiers:

(42) a. The floor is slightly/a little bit dirty.
b. ? The floor is slightly/a little bit filthy.

CP correctly predicts that combining *filthy* with a minimizing modifier such as *slightly* or *a little bit* is dispreferred, since the same point could be equally well denoted by combining some other modifier with its relative scalemate *dirty*, for example:

(43) The floor is rather/fairly/very dirty.

Summing up: epistemic uses of *certain* seem to be equivalent to evidential uses of *probable*, as illustrated by (36). This provides a strong argument for the Certainty Account of Evidential Probability. While differences in the acceptability of *certain* and *probable* under mid-scale percentage modifiers might seem to challenge this argument, closer inspection suggests that these differences do not reveal an inequivalence between the relevant expressions, and hence do not constitute counterexamples to the Certainty Account. Rather, these differences are better explained by a general dispreference for combining mid-scale percentage modifiers with a max adjective when a relative gradable adjective will serve just as well.

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 quantifers. Necessity modals (denoted '$\Box$') such as *must*, *has to*, and *necessarily* are analyzed as universal quantifiers:

Classical Analysis of Epistemic Modals:
$\ulcorner \Diamond p \urcorner$ is true at a point of evaluation i iff p is compatible with the relevant epistemic state.
$\ulcorner \Box p \urcorner$ is true at i iff the relevant epistemic state entails p.

What *sort* of epistemic state is relevant here? The most common view is that the relevant state is knowledge—call this the 'Knowledge Analysis.'[35] I propose instead that the relevant state is epistemic certainty:

Certainty Analysis of Epistemic Modals:
$\ulcorner \Diamond p \urcorner$ is true relative to a contextual standard s iff p is compatible with what's epistemically certain (relative to s).
$\ulcorner \Box p \urcorner$ is true relative to s iff p is entailed by what's epistemically certain (relative to s).

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.[36]

5.2 In favor of the Certainty Analysis

Suppose a detective asserts:

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

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

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

Indeed, it sounds very odd to follow an assertion of (44) with a denial of (45):

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

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:

(47) ?? 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, (44) says that it's epistemically certain that the butler did it. And so (46) is predicted to be contradictory. Similarly, the first conjunct of (47) entails that it's epistemically certain that the cook was not involved, which contradicts the second conjunct.

[35] Versions 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).

[36] Only 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 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-centric framework.

By contrast, the Knowledge Analysis leaves these data unexplained. According to the Knowledge Analysis, (44) says that the relevant agents' knowledge entails the butler did it. But if knowledge does not entail certainty, the assertability of (44) provides no guarantee that (45) is assertable. Similarly, we'd expect (46) and (47) 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 (46) and (47) could 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:

(48) ?? Suppose both that there's no possibility that the cook was involved and it's not certain that the cook wasn't involved.

(49) ?? If the butler must have done it and it's not certain whether he did it . . . [37]

A merely pragmatic explanation of the infelicity of (46)-(47) does not generalize to explain the oddity of (48)-(49). By contrast, the Certainty Analysis has no trouble here. According to the Certainty Analysis, (48) and (49) invite the addressee to entertain an incoherent state of affairs, thereby accounting for their infelicity.

Thus the infelicity of (46)-(49) provides 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:

(50) a. You may possibly have read my little monograph on the subject.[38]
b. You may have read my little monograph on the subject.

The most natural reading of (50a) is a 'concord reading', on which it's simply equivalent to (50b) (or a slightly hedged version thereof). It is less natural to give (50a) a 'cumulative' reading, according to which it's possible that there is a possibility that the addressee has read the speaker's monograph.[39]

[37] See also Beddor and Goldstein 2018 for data on the connection between epistemic modals and certainty in belief reports.

[38] Geurts and Huitink (2006) take (50a) from Sir Arthur Conan Doyle's *The Hound of the Baskervilles.*

[39] This not to deny that there's a cumulative reading available for such constructions, or that certain contexts might make the cumulative reading more easily accessible. There may also be dialectical variation in how easily accessible cumulative readings are for double modals. For example, dialects of English spoken in the southern United States allow for two modal auxiliaries to occur in the same clause; these occurrences are often given a cumulative rather than a concord reading.

It's widely agreed that in order for a concord reading to be possible, the two modals must be equivalent. This explains why (50a) allows for a concord reading, but (51) and (52) do not:

(51) ? You must possibly have read my monograph.

(52) ? You may certainly have read my monograph.

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-):

(53) Something about her told him that she must certainly be noble.

(54) Vanguard keeps costs low, but people must certainly be making financial services industry salaries.

(55) 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. 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., (56)) and *concessive knowledge attributions* (e.g., (57)):

(56) ?? It's raining but it might not be raining.

(57) ?? I know that it's raining. But it might not be.[40]

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 (56) and (57) are infelicitous.

By contrast, the Knowledge Analysis seems 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 (57) are guaranteed to be inconsistent. And while epistemic contradictions such as (56) are not predicted to be inconsistent, proponents of the Knowledge Analysis have a plausible pragmatic explanation of their infelicity. After all, many Knowledge Firsters hold that assertion is governed by a knowledge norm:

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

Knowledge Norm of Assertion (KA): Assert *p* only if you know *p*.[41]

Combining the Knowledge Analysis with KA predicts that (56) is never assertable. Knowing the first conjunct (*It's raining*) precludes knowing the second (*It might not be*), so no one could assert (56) 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.:

(58) ?? It's raining but I don't believe/know it's raining.

So by appealing to KA, proponents of the Knowledge Analysis offer a unified account 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. There are a couple of ways of formulating such a norm, depending on whether one thinks permissible assertion requires epistemic certainty, psychological certainty, or both. For my purposes, I will focus on a simple version of a certainty norm, according to which epistemic certainty is the requisite state:

Certainty Norm of Assertion (CA): Assert *p* only if *p* is epistemically certain for you.[42,43]

Armed with CA, advocates of the Certainty Analysis can explain the infelicity of epistemic contradictions on pragmatic grounds. If the first conjunct of (56) 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. Thus no one could assert (56) 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:

(59) ?? 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 (59). CA does; it would be impossible to utter this sentence while obeying CA. Hence (59) suffers from the same ailment as (56). Assuming that epistemic certainty entails knowledge, CA also accounts for the original Moorean assertions (e.g., (58)) that motivated KA.[44]

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

[42] The 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. For a defense of CA as an alternative to KA, see Stanley 2008.

[43] As 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*.

[44] Moorean assertions involving psychological certainty also sound odd, e.g.,

(iv) ?? It's raining but I'm not certain that it's raining.

A further advantage of CA is that it addresses a residual challenge for the idea that knowledge does not entail certainty. The challenge arises from the fact that it sounds odd to say things like:

(60) ?? I know it's raining, but I don't know for certain that it's raining.

According to CA, (60) is infelicitous for the exact same reason as (59). No one who is obeying CA could utter (60) unless both conjuncts were epistemically certain, but—given the factivity of knowledge—that is impossible.[45]

Thus a certainty norm of assertion offers one promising way of explaining the recalcitrant data. 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 $\mathcal{E}$-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 infelicity of (56) and (57). 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 (56) (*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

To explain this, we can appeal to the normative connection between psychological and epistemic certainty (§2.1). Suppose the first conjunct of (iv) is epistemically certain. Then the speaker ought to be psychologically certain that it's raining. So either the second conjunct is false (and hence not epistemically certain) or the speaker is being irrational.

[45] These advantages notwithstanding, some might worry that a certainty norm of assertion is too demanding. Suppose I stand in a strong epistemic relation to p, but my relation does not quite rise to the level of certainty. If I am asked whether p is true, wouldn't it be unduly hesitant to refrain from asserting p? (Thanks to a referee for raising this concern.) While I acknowledge the force of this objection, two points are worth noting. First, we are not usually restricted to just two options: asserting p or falling silent. Another option is to use probabilistic language (§4.2.2). For example, we can say *Probably p*, or even *It's 99% likely that p*. In cases where we clearly fall short the contextual standards for certainty, arguably a hedged assertion along these lines is the most appropriate response. Second, it is telling that we often *opt out* of providing an answer by citing our lack of certainty, e.g., Q: *When is the bus coming?*, A: *I'm not certain/sure*. For more on 'opting out', see Dorst (2014).

infelicity of concessive knowledge attributions: given the factivity of knowledge, (57) 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:

(61) 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 Wherefore Knowledge?

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.

By way of conclusion, I want to address a residual question that may be bothering some readers. I've argued that many of the roles traditionally assigned to knowledge are better served by certainty. If I'm right about this, then what work is left for knowledge to do?

This may seem particularly puzzling given the analysis of knowledge and certainty defended here. On our analysis, knowledge and epistemic certainty have much in common. Both involve ruling out error possibilities. It's just that epistemic certainty involves ruling out a wider range thereof. There is a distinction here, to be sure, but why is this a distinction worth drawing? What is the point of having both notions?[46]

In response, we should start by noting that natural language frequently makes similar distinctions. Consider the distinction between *good* and *excellent*. Both terms have similar meanings. But the latter is stronger than the former. An excellent paper is a good paper, but a good paper need not be excellent. Moreover, this difference can matter a great deal (e.g., when it comes to publishing in certain journals).

For an even closer parallel with knowledge and certainty, consider the distinction between weak necessity modals (*should*, *ought*) and strong necessity modals (*must*, *necessarily*, *have to*). While these two classes of expressions have similar meanings, the latter are logically stronger than the former, as shown by the contrast in (62):

(62) a. You should give all of your extra income to charity, but you don't have to.
b. ?? You must give all of your extra income to charity, but you don't have to.

It is a matter of debate how best to analyze the distinction between weak and strong necessity modals. However, one natural approach is to analyze both as universal quantifiers over worlds. It's just that the domain of the former is a proper subset of the domain of the latter. To illustrate this idea with deontic modals (e.g. (62)), we might hold that deontic *should* universally quantifies over all of the very best worlds in some contextually relevant domain. By contrast, deontic *must* and *have to* universally quantify over all of the acceptable worlds in the contextually relevant domain—that is, all the worlds that are good enough for the purposes of the context.[47] If we accept an analysis in this vein, the structural relationship between weak and strong necessity modals provides a particularly close analogue for the relationship between knowledge and certainty.

Thus my analysis of the relation between knowledge and certainty is not without precedent; we carve out similar distinctions elsewhere in the normative landscape. Still, the question remains: what is the point of having a notion of knowledge in addition to certainty?

Faced with this question, one option would be to go *pluralist*: while many important epistemological roles are best served by certainty, others are best served by knowledge. For example, one *prima facie* attractive hypothesis is that knowledge serves as the normative standard for *belief*.[48]

[46] Thanks to a referee for raising this question.

[47] For an analysis along these lines, see von Fintel and Iatridou 2008; Beddor 2017.

[48] For sympathetic discussion of the idea that knowledge is the norm of belief, see Williamson 2000, forthcoming; Sutton 2007; Moss 2018.

This suggestion can be made to fit rather nicely into the theoretical framework developed here. On our framework, belief is a weaker state than psychological certainty: one can believe something without being psychologically certain of it. Precisely because belief is weaker than psychological certainty, an epistemic certainty norm of belief is highly implausible. Suppose you believe the butler did it. But you are not quite certain of it, and rationally so—your evidence leaves open the possibility, albeit a remote one, that the gardener was the culprit.

So we need a less demanding epistemic state than certainty to serve as the standard for belief. Knowledge is a natural candidate. This would offer one way of further developing the analogy floated in §2: *knowledge is to belief as psychological certainty is to epistemic certainty*. On the view that emerges, knowledge imposes the same sort of normative constraint on belief that epistemic certainty imposes on psychological certainty.

While this hypothesis is appealing, we should recognize that it faces some important challenges. In particular, some of the linguistic data indicating that one can permissibly believe something in the absence of certainty also seem to indicate that one can permissibly believe something in the absence of knowledge. Consider:

(63) I believe she's going to accept the job, but I don't know one way or the other.

While intuitions may differ across speakers, to my ears (63) sounds perfectly coherent. This creates at least a *prima facie* challenge for a knowledge norm of belief.[49]

Those who reject a knowledge norm of belief could go one of two ways. One would be to stick with the pluralist route and find some other function for knowledge to serve. But a more radical option is also worth considering. According to the *monist* route, we should reject the assumption that knowledge has any interesting work to do. Perhaps once we look into the matter, we'll find that certainty is better qualified to fulfill *all* of the jobs traditionally associated with knowledge. I will leave it to future research to investigate which of these routes proves more promising.[50,51]

[49]For further linguistic data suggesting that belief is a relatively weak state, see Hawthorne et al. 2016; Beddor and Goldstein 2018. Another potential challenge for a knowledge norm of belief is more theoretical. The knowledge norm of belief is sometimes motivated by the idea that the norm of belief parallels the norm of assertion. Belief, on this view, is a sort of 'inner assertion,' and hence should be held to the same normative standard as 'outer assertion.' But if we accept my earlier suggestion that epistemic certainty is the norm of assertion, then we cannot use this argument to motivate the idea that knowledge is the norm of belief. (Thanks to a referee for helpful comments here.)

[50]For further exploration of the applications of certainty in epistemology, see Beddor forthcoming, where I argue for a certainty norm of practical reasoning.

[51]A very early version of some of these ideas appeared in chapter three of my dissertation (Beddor 2016). This material has benefited enormously from the feedback of many different people over the years, including D Black, Andy Egan, Branden Fitelson, Georgi Gardiner, Alvin Goldman, Simon Goldstein, Dan Greco, Dan Lassiter, Ben Levinstein, Ricardo Mena, Aidan McGlynn, Andrew Moon, Jonathan Schaffer, and Holly Smith. I am also grateful to audiences at the Bled Epistemology Conference, the University of Valencia, NUS, the University of British Columbia, and Ernie Sosa's dissertation group at Rutgers. Finally, thanks to two referees for *Philosophers' Imprint* for their careful and insightful comments.

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.
Mark Aronoff. *Word Formation in Generative Grammar*. MIT Press, Cambridge, 1976.
A.J. Ayer. *Language, Truth, and Logic*. Dover, New York, 1936.
A.J. Ayer. *The Problem of Knowledge*. Macmillan, London, 1956.
Bob Beddor. *Reduction in Epistemology*. PhD thesis, Rutgers University, 2016.
Bob Beddor. Justification as Faultlessness. *Philosophical Studies*, 174:901–926, 2017.
Bob Beddor. Certainty in Action. *The Philosophical Quarterly*, forthcoming.
Bob Beddor and Simon Goldstein. Believing Epistemic Contradictions. *Review of Symbolic Logic*, 11:87–114, 2018.
Bob Beddor and Carlotta Pavese. Modal Virtue Epistemology. *Philosophy and Phenomenological Research*, Advance Publication Online, 2018. doi: https://doi.org/10.1111/phpr.12562.
Michael Blome-Tillman. Knowledge and Presuppositions. *Mind*, 118:241–294, 2009.
Roger Clarke. Belief is Credence One (In Context). *Philosophers' Imprint*, 13(11):1–18, 2013.
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. *The 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*. Libraire 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.
Kevin Dorst. Can the Knowledge Norm Co-Opt the Opt-Out? *Thought*, 3(4):273–282, 2014.
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.
Jeremy Goodman and Bernhard Salow. Taking a Chance on KK. *Philosophical Studies*, 175:183–196, 2018.
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. *Knowlege and Lotteries*. Oxford University Press, New York, 2004.

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

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

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.

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. *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.

John MacFarlane. *Assessment-Sensitivity: Relative Truth and its Applications*. Oxford University Press, Oxford, 2014.

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

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

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

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.
Kai von Fintel and Sabine Iatridou. How to Say 'Ought' in Foreign: the Composition of Weak Necessity Modals. In Guéoron and Lecarme, editors, *Time and Modality (Studies in Natural Language and Linguistic Theory 75)*. Springer, 2008.
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.
Timothy Williamson. Justifications, Excuses, and Sceptical Scenarios. In Dorsch and Dutant, editors, *The New Evil Demon Problem*. Oxford University Press, Oxford, forthcoming.
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.