Expressing Uncertainty

Anhui Huang

(Draft Oct 2026; please do not cite or circulate)

Abstract

There is an asymmetry between denials of knowledge and denials of certainty. The latter, but not the former, generates an *opinion inference*: "S is not sure" communicates that S has *some* opinion bearing on the question under discussion, whereas "S does not know" does not. This simple observation has been overlooked in the literature on certainty. I argue that it is not an isolated puzzle, but a revealing symptom of what is wrong with the dominant view of certainty, which identifies certainty with maximal credence. In this paper, I propose an alternative account of certainty that can explain the observation: being certain *just is* taking inquiry to be closed. My view of certainty also has a range of independent advantages over the credence account. For example, it explains why we can be sure or unsure about both what is the case and what to do, why we shouldn't be sure when we lack knowledge or the right *kind* of evidence, and why knowledge does not entail certainty but is normatively in tension with uncertainty.

1 Introduction

Suppose someone asks you: Who won the last Dutch election? And suppose that you know absolutely nothing about Dutch politics—you have no idea when the election occurred, who the candidates were, or what the election procedure involves. In fact, before the question was asked, you were not even aware that there *was* a Dutch election. Now two responses are available to you:

(1) I don't know.

(2) I am not sure.

Intuitively, you should answer (1) rather than (2), because (2) seems to communicate that you have *something* in mind that bears on the question, *Who won the last Dutch election?*. Since you have none, (2) is misleading.

A similar observation can be made about third-person reports. Consider:

(3) Jane doesn't know who won the last Dutch election.

(4) Jane is not sure who won the last Dutch election.

Intuitively, (4) communicates that Jane has *something* in mind that bears on who won the Dutch election, whereas (3) does not.

This is puzzling, though. Why is it that, simply by denying that one is sure, we communicate that one has *something* in mind that bears on the question under discussion (QUD)?[1] In other words, why would denials of certainty generate such an *opinion inference*?[2]

In §2, I provide additional linguistic data in support of the opinion inference and examine more closely how it behaves. In particular, I show that it exhibits two striking features. First, the opinion inference arises both when the QUD is theoretical, (e.g., who won the last Dutch election) and when the QUD is practical (e.g., what to make for the department potluck). Secondly, while certainty talk behaves differently from knowledge talk or belief talk with respect to the opinion inference, it patterns well with settling talk.

In §3, I make the content of the opinion inference more precise by spelling out the conditions on which one can felicitously report a subject as not sure. I argue that believing some answer to the QUD is too strong, but merely comprehending the QUD is too weak, and actively inquiring into the QUD is sometimes too strong and sometimes too weak. I conclude that the content of the opinion inference is that one has made some *progress in inquiry*.

These observations are hard to explain for the dominant view of certainty in contemporary epistemology, which identifies certainty with maximal credence.[3] It is standard to accept, either explicitly or implicitly, the following:[4]

> (Credence account of certainty) To be sure of a proposition is to have maximal credence in that proposition.

---

[1] I am using the QUD framework developed by Roberts (2012), according to which a discourse can be understood as a cooperative activity aiming at answering a question that is accepted by participants of the discourse as the *question under discussion*. In this paper I will focus on cases where the QUD is explicitly raised, but I intend the asymmetry and my explanation of it to apply to cases where the QUD is implicit as well.

[2] In this paper, I will stick to linguistic data involving "sure" and "not sure." But the same asymmetry can be observed for "certain" and "not certain." I choose to focus on "not sure" because it's a more frequently used phrase.

[3] The standard approach in contemporary epistemology takes certainty to be the limiting case of graded belief or confidence, namely, maximal credence (Halpern, 1991; Clarke, 2013; Moss, 2019). Even sophisticated recent attempts to rehabilitate certainty continue to analyze it in terms of maximal credence (Beddor, 2020). A typical introduction to formal epistemology will take for granted that certainty is maximal credence. For example, the *Routledge Encyclopedia of Philosophy* explains maximal credence as "corresponding to the strongest possible belief or certainty" (Pettigrew, 2023). Some traditional epistemologists who reject nearby theses also presuppose that certainty is maximal credence. For example, Jackson (2020) rejects identifying belief with credence one because it is "in tension with the plausible idea that we hold many of our beliefs with less than maximal certainty."

[4] Two notable exceptions are Horgan (2017) and Goodman and Holguín (2022).

The credence account of certainty struggles to explain why the opinion inference arises across practical and theoretical domains, why certainty talk patterns with settling talk, and why the content of the opinion inference concerns one's progress in inquiry.

One might attempt to explain away the opinion inference as an instance of scalar implicature, which would have little to do with the nature of certainty. But in §4, I show that the opinion inference is best modeled as a presupposition, i.e., information that is taken for granted in an utterance, as opposed to an implicature. This means that we have to take the opinion inference seriously. In fact, as I will show in §5, the opinion inference is not an isolated linguistic puzzle, but the tip of the iceberg of what is wrong with the credence account. We have no choice but to rethink the nature of certainty.

In §6, I propose my alternative account that can explain the opinion inference, the inquiry account of certainty: the state of being certain *just is* the state of taking inquiry to be closed (or equivalently, taking a question to be settled). This is the state of being settled on an answer, where one relies on this answer in further deliberation, takes oneself to know the answer, and loses any disposition to inquire further. This state is constrained by, but ultimately orthogonal to, one's credences. Instead, certainty is a distinctive way of relating to inquiry, and our certainty talk tracks the status or progress of one's inquiry. This helps explain the opinion inference. In asking where one is, we presuppose that one is *somewhere*. So in asking whether one is at the end of inquiry or denying that one is there, we presuppose that one is *somewhere* locatable in a line of inquiry, i.e., one has made *some* progress in inquiry.

My account of certainty also clarifies the relation between certainty and knowledge. It explains why, although knowledge does not entail certainty, it is normatively in tension with uncertainty. In particular, this helps to explain some puzzling linguistic data reminiscent of Moore's puzzle. I discuss these upshots in §7.

2 Additional Data for the Opinion Inference

I take it to be clear from the initial cases that denials of certainty do generate an opinion inference. But in case it is not, here are three ways to sharpen the intuition. First, compare the following responses to the question, What kinds of fish live in Lake Superior?

(5) I don't know. I know absolutely nothing about fish or lakes.

(6) #I am not sure. I know absolutely nothing about fish or lakes.

(5) sounds fine, but (6) is infelicitous. The contrast persists in third-personal cases:

(7) Jennie doesn't know (what kinds of fish live in Lake Superior). She doesn't know anything about fish or lakes.

(8) #Jennie is not sure (what kinds of fish live in Lake Superior). She doesn't know anything about fish or lakes.

This contrast suggests that denial of certainty is in tension with admitting that the subject is completely clueless about what the answer might be.

Second, imagine asking a highly technical or specific question of someone who clearly lacks the relevant background. For example, suppose I ask my mother (who knows nothing about philosophy): *Does Gibbard's plan expressivism successfully avoid the Frege-Geach problem?* It would be strange if she answered "I am not sure;" but it would be fine if she answered "I don't know." This shows that denying that one is sure is in tension with it being common knowledge that one is completely clueless about what the answer might be.

Lastly, compare the following:

(9) #I don't know. *Actually*, I have no idea.

(10) I am not sure. *Actually*, I have no idea.

(9) sounds oddly redundant, but (10) sounds fine. Here is a plausible explanation. The expression "actually" is often used to *take back* some pragmatically communicated content. For example, uttering (11) alone communicates that it is not the case that Pedro is *definitely* coming to the party:

(11) Pedro might be coming to the party.

But uttering (12) allows the speaker to take back that content:

(12) Pedro might be coming to the party. *Actually*, he definitely is.

The same explanation plausibly applies to (10): in saying that they are not *sure*, the speaker communicates that they have *some* idea regarding the QUD, which they then take back by clarifying that, actually, they have no idea. This also explains why (9) sounds redundant: denying that one *knows* does not communicate that one has *some* idea regarding the QUD, so there is no content to take back.

2.1 Opinion inference in the practical realm

A striking feature of the opinion inference is that it arises whether the QUD is theoretical or practical. Suppose someone asks you: What are you going to make for the department potluck? And suppose that you have absolutely no idea—you haven't thought about this at all, you don't know anything about cooking, and you have no idea how to even begin deliberating. In this case, too, it would be weird for you to respond with (14) rather than (13):

(13) I don't know.
(14) #I am not sure.

Again, the same applies to third-person cases:

(15) Pat doesn't know what to make for the department potluck.
(16) Pat isn't sure what to make for the department potluck.

Unlike (15), (16) strongly suggests that Pat is not completely clueless. This means that the opinion inference arises both when the QUD is theoretical, i.e., about what is the case, and when it is practical, i.e., about what to do.

This observation is puzzling if we accept the credence account of certainty. The problem is not just that it fails to explain why the opinion inference arises when the QUD is practical. Instead, it is that it cannot even make sense of what sentences like (16) are saying: credences apply to propositions, but the answers to a practical question are plans, or decisions. After all, credences are degrees of belief, and belief is simply the wrong attitude for plans. However, it is very natural to ascribe uncertainty toward practical questions, as in (16), as well as theoretical questions, and when we do, the opinion inference arises uniformly in both cases.

2.2 Patterning with settling talk

Another striking observation is that, while denials of knowledge or denials of belief do not generate an opinion inference, both denials of certainty and denials of a settled attitude do. For example, both (17) and (18)-(20), which are various kinds of settling talk, generate the opinion inference that Adam is not completely clueless about where Jeita Grotto is:

(17) Adam is not sure where Jeita Grotto is.
(18) Adam has not settled where Jeita Grotto is.
(19) Adam has not resolved the question of where Jeita Grotto is.
(20) Adam has not made up his mind about where Jeita Grotto is.

By contrast, (21)-(22) do not generate the opinion inference:

(21) Adam does not know where Jeita Grotto is.
(22) Adam does not have a belief about where Jeita Grotto is.

This puts certainty talk and settling talk on one side of the opinion inference, and knowledge talk and opinion talk on the other. Again, the credence account struggles to explain why certainty talk would differ from belief talk while patterning with settling talk. As I will show, these observations ultimately lead us toward an alternative account, according to which certainty is essentially a question settling attitude.

3 The Content of the Opinion Inference

Denying that a subject is sure generates the opinion inference that the subject has something in mind that bears on the QUD. In this section, I will make the content of this inference more precise, by spelling out the conditions for felicitously reporting a subject as not sure.

At first glance, one might think that the opinion inference is simply that the subject *believes* some answer to the QUD. This is too demanding, however. After all, one can felicitously utter:

(23) I am not sure. I am still suspending my judgment./I haven't figured out what to think./I don't have an answer in mind yet.

In fact, the subject does not even have to know what the possible answers are to the QUD. Consider:

Dutch politics expert
Blake is very knowledgeable about the political history of the Netherlands. She has a deep understanding of Dutch politics in general. But she did not learn who the candidates of the last Dutch election were. When someone asks Blake, *Who won the last Dutch election?*, she says:

(24) I am not sure. I don't know who the candidates were, but whoever won must belong to one of the more conservative parties.

Since Blake doesn't know what the possible answers are, she does not believe (think likely, or have high credence in) any of the answers. Nevertheless, she can be felicitously reported as not sure (about who won the last Dutch election). This counts against any view on which the content of the opinion inference entails that the subject has an attitude toward some *particular* answer to the QUD.

A natural thought is that maybe the subject just needs to be able to *comprehend* the QUD. This suggestion would explain why exchanges like (25) are odd:

(25) Sofia: Do quarks move fast enough for relativistic effects to be significant?
Joe: #I am not sure. What is a quark? What is a relativistic effect?

Since Joe does not know anything about quarks and relativity, he lacks the conceptual resources to comprehend the question. As a result, he cannot felicitously report himself as not sure. Others have proposed a similar condition as a necessary condition for being unsure: Goodman and Holguín (2023) propose that a subject is unsure whether p just in case (i) the subject understands the question whether p and (ii) the subject is neither sure that p nor sure that not-p.

That the subject comprehends the QUD is only a *necessary* condition for felicitously reporting them as not sure, however. There are many cases in which it would be infelicitous to report the subject as not sure, even though they comprehend the QUD. For example, someone who knows nothing about Dutch politics still comprehends the question, *Who won the last Dutch election?*, yet they cannot felicitously respond with "I am not sure." Or consider the following sentence:

(26) #Sally is not sure whether the total number of atoms in the universe is odd or even.

If we take for granted that, like most of us, Sally has no special insight into the total number of atoms, (26) strikes us as quite odd. But surely Sally comprehends the question. Therefore, denying that a subject is sure communicates something *stronger* than just that the subject comprehends the question (and that it is not the case that the subject is sure).[5]

Perhaps the subject needs to be *inquiring into* the question, rather than merely comprehending it. This might account for the oddness of (26): since the question never arises for Sally, she never inquires into the question. Plausibly, to inquire into a question involves aiming to answer that question.[6] So perhaps one cannot be felicitously reported as not sure about a question unless one aims at answering the question.

However, this condition predicts both too much and too little. A middle schooler can inquire into the question, *What cures cancer?*. They could be fairly invested in trying to figure out what the answer to the question is. Still, it would be odd to say:

(27) #The middle schooler is not sure what cures cancer.

On the other hand, someone who can be felicitously reported as not sure about a question need not be inquiring (or have inquired) into that question. Suppose that someone asks me: *How many letters are in the word, "supercalifragilisticexpialidocious"?*. I might not bother actually inquiring into the question, since I find it boring or uninteresting. Still, I can felicitously respond with "I am not sure." Or suppose that someone asks me, *If you could bring three items with you to live on a deserted island for a year, what would you bring?*. I can respond with "I am not sure" and *then* proceed to think about my answer to the question. In this case, I can felicitously report myself as not sure *before* I begin to engage in any inquiry.

[5] This is not a criticism of Goodman and Holguín's (2022) account of unsurety. Goodman and Holguín are not concerned with the linguistic behavior of 'sure' and 'not sure'. And while they acknowledge, in a footnote, that both 'Hania is unsure whether Biden and Trump's phone numbers end in the same digit' and 'Hania isn't sure whether Biden and Trump's phone numbers end in the same digit' sound a bit odd, they do not take that to be evidence against their view. They do not discuss this observation further, after conjecturing that there must be a pragmatic explanation.

[6] Providing a full account of inquiry is beyond the scope of this paper. For two plausible accounts, see Daniel Friedman (2025) and Jane Friedman (2013, 2017). For the purpose of this paper, it is important that both these accounts agree that an inquiring agent is in a goal-directed state of mind, whose goal is to answer the question.

So believing some answer to the QUD is too strong, but merely comprehending the QUD is too weak, and actively inquiring into the QUD is sometimes too strong and sometimes too weak. Nevertheless, connecting uncertainty to inquiry seems to me to be the right direction. Denying that one is sure seems to suggest that one is *somewhere locatable in inquiry*. My proposal, then, is the following:

> (Content of the opinion inference) Denying that a subject is sure about the QUD communicates that the subject has made *progress in inquiry* into the QUD.

I would like to make three clarifications. First, one does not count as having made progress if all that one has learned is basic common sense, i.e., something that can be reasonably taken for granted. For example, I have not made progress in inquiry into who won the last Dutch election if all that I know is that it was a human. Secondly, sometimes one can make progress in inquiry without having actually inquired into the question. For example, suppose that a geologist who has spent many years studying emeralds is asked to decide whether a particular stone is an emerald. Although the geologist has not inquired into the question, *Is this stone an emerald?*, as soon as they consider that question, it seems that they have already made lots of progress just by being an expert on emeralds. Finally, I am using "progress" in a non-factive sense: describing someone as having made progress in inquiry does not entail that they have gotten anything right. I can count as making progress in inquiry by ruling out a possible answer, for example, even if that is in fact the correct answer.

Now that we have clarified the content of the opinion inference, it poses an immediate problem for the standard credence account of certainty. On that account, being certain that p consists in having maximal credence in p, while being uncertain that p consists in having less than maximal credence in p. But the inference associated with "S is not sure" concerns S's *progress in inquiry*: it presupposes that S has made some progress toward answering the question. Credence, however, has no intrinsic connection to progress in inquiry. According to standard decision theory, one can have a credence in any proposition one can understand, regardless of whether one has made any progress toward answering it. The full range of credence, including credence one, is open to a subject who has made no progress at all in inquiry. The credence account

cannot explain why "I am not sure" is infelicitous when the speaker has made no progress in inquiry.

4 Not Scalar Implicature

The opinion inference poses a serious problem for the credence account of certainty. A defender of the credence account might attempt to explain away the opinion inference in a way that does not require revising our understanding of certainty. One way to do that is by explaining the opinion inference as a *scalar implicature*. Despite its initial promise, however, the scalar implicature explanation does not work. We have no choice but to rethink the nature of certainty.

Scalar implicature is a kind of conversational implicature generated on the basis of one of Grice's (1989) maxims of quantity:

(Quantity) Make your contribution *as informative as is required*.

Based on the assumption that the speaker is cooperative and obeying the maxims, the hearer can make inferences by comparing what the speaker *in fact says* with what they *could have said*. Consider a paradigm case of scalar implicature:

(28) Some of the students passed the exam.

implicates (28a) Not all of the students passed the exam.

The reasoning is as follows. The speaker says that *some* of the students passed the exam, but they could have said something stronger (more informative), namely that *all* of them passed the exam. Assuming that the speaker is obeying (Quantity), it must be that they have a reason not to say the stronger thing. It must be, then, that it is not true that all of the students passed the exam.[7]

Building on Grice's framework, Horn (1972) and Gazdar (1979) have developed more systematic accounts of scalar implicature by introducing the notion of *quantitative scales*.[8] A quantitative scale is a set of linguistic expressions that are (i) ordered by their strength, and (ii) equally lexicalized, of the same word class, from the same register, and

[7] This is a rough story of how the reasoning goes. The details are controversial. For example, one might think that the conclusion should only be that the speaker does not know that all of the students passed the exam, and that other contextual factors are needed to license the further inference that it is not true that all of the students passed. For a representative discussion of the possibilities here, see Levinson (2000).

[8] Different names have been given to quantitative scales: e.g., 'Horn scales', 'linguistic scales'.

from the same semantic field.[9] Given a quantitative scale, when the speaker makes a statement with a weaker expression from that scale, it implicates that the corresponding statement involving the stronger expression is not true. For example, <all, some> is a quantitative scale, which explains why uttering (28) implicates (28a). Importantly, negation reverses quantitative scales (Horn, 1972; Atlas & Levinson, 1981): if <φ, ψ> forms a quantitative scale, then the same goes for <¬ψ, ¬φ>. For example, since <all, some> is a quantitative scale, so is <not some (none), not all>, which licenses the inference from (29) to (29a):

(29) Not all of the students passed the exam.
implicates (29a) Some of the students passed the exam.

Intuitively, "sure" is stronger than "(have) some idea." So one might think that <sure, some idea> forms a quantitative scale. Since negation reverses scales, this means that <no idea, not sure> also forms a quantitative scale. As a result, the following inference from (30) to (30a) would be licensed:

(30) S is not sure about Q.
implicates? (30a) It is not the case that S has no idea about Q.
or (30b) S has some idea about Q.

This would provide an explanation for the opinion inference. However, this account faces three problems.

4.1 Too much and unequal morphological complexity

The first problem is that <sure, some idea> fails to meet *all* of the constraints on quantitative scales. Two of the constraints on quantitative scales concern the morphological complexity of the expressions: (i) they should be relatively morphologically simple, and (ii) they should have the same degree of morphological complexity. These constraints explain why <if and only if, if> fails to be a quantitative scale: "if and only if" is morphologically complex, and it is significantly more so than "if." For the same reason, <some but not all, some> is not a quantitative scale, even

[9] The strength of linguistic alternates are typically cashed out in terms of entailment. An expression, x_1, is stronger than another, x_2, iff $S(x_1)$ unilaterally entails $S(x_2)$, where S is an arbitrary simple sentence frame. For a more detailed summary of (ii), see Levinson (2000).

though <all, some> is. Without these constraints, we run into the notorious symmetry problem.[10]

The proposed <sure, some idea> fails both of these constraints. "Some idea" is itself morphologically complex. And if we look for alternative expressions that are roughly synonymous, they turn out to be just as complex: e.g., "some opinion," "some guess," "not clueless." In addition, "sure" and "some idea" are not equally complex. "Sure" is a morphologically simple expression, not just in English but across many languages—e.g., "seguro" in Spanish, "确定 (quèdìng)" in Chinese, whereas "some idea" is more complex. So the proposed quantitative scale <sure, some idea> is unlike canonical quantitative scales like <all, some>, <and, or>, and <necessary, probable, possible> because it consists of linguistic items that fail to have the right degree of complexity.

Admittedly, what constraints there should be for quantitative scales is not entirely uncontroversial. Hirschberg (1991) argues, for example, that there should be very little constraint on them; in fact, she argues that, given the right context, any poset (partially ordered set) of expressions can license a scalar implicature.[11]

> People can and do form scales from items that bear little intrinsic relationship to one another; these scales are hard to predict outside the context in which they are evoked. (Hirschberg 1991, p. 114)

So given the right context, <sure, some idea> can form a quantitative scale, too.

There are prima facie reasons to doubt whether the phenomenon Hirschberg observes should be treated as a kind of scalar implicature. But more importantly, this response would only create more problems for the scalar implicature account, for if any poset can form a quantitative scale, then both <sure, some idea> and <know, some idea> can be quantitative scales. This predicts that (31) would implicate (31a):

(31) S doesn't <u>know</u> Q.

implicates? (31a) It is not the case that S has <u>no idea</u> about Q.

or (31b) S has <u>some idea</u> about Q.

[10] The problem is explaining why, for example, "some" implicates "not all," but not "some but not all."
[11] A poset is a set ordered by a relation which is reflexive, transitive, and antisymmetric.

As a result, the scalar implicature account would only explain one side of the asymmetry: it would explain why denials of certainty generate the opinion inference, but it would fail to explain why denials of knowledge do not generate the same inference. In fact, it would falsely predict that they do.

4.2 Mischaracterized Content

The scalar implicature account also mischaracterizes the content of the opinion inference. Scalar implicatures are licensed by quantitative scales, and the expressions that make up a quantitative scale determine the content of the scalar implicature. Since the proposed quantitative scale is <sure, some idea>, the scalar implicature account predicts that the content of the opinion inference is that the subject has *some idea* (about the QUD). While this was an acceptable statement of the opinion inference initially, it is no longer accurate now that we have made its content precise. Consider, again, the case of the Dutch politics expert, Blake. Given that Blake does not know who the candidates of the Dutch election were, it would be odd for her to utter (32).

(32) #I have some idea who won the Dutch election.

But it is okay for her to report herself as not sure about who won the Dutch election. So the scalar implicature account falsely predicts that Blake cannot be felicitously described as not sure.

4.3 Projection

Lastly, the opinion inference projects in ways that are unusual for implicatures and typical of presuppositions. Both presuppositions and implicatures are pragmatically communicated content of an utterance. But unlike implicatures, presuppositional content is typically not at-issue; instead, it is information that the speaker *takes for granted* in making the utterance. For example, "John knows that the car is repaired" presupposes that the car is repaired, and "Sally stopped smoking" presupposes that Sally smoked before. One of the most distinctive diagnostics for presuppositions is that they survive under negation and questioning, whereas implicatures do not (Simons, Beaver, Tonhauser & Roberts, 2010). For example, (33)-(35) all seem to commit the speaker to the content that Sally smoked before:

(33) Sally stopped smoking.

(34) Sally didn't stop smoking.

(35) Has Sally stopped smoking?

On the other hand, scalar implicatures fail to project in this way. Consider:

(36) Josephine ate some of the cookies.

implicates (36a) Josephine didn't eat all of the cookies.

(37) Did Josephine eat some of the cookies?

*does **not** implicate* (37a) Josephine didn't eat all of the cookies.

Now we can apply this diagnostic to the opinion inference and "sure."

(38) Patrick is not sure that he can park here. (negation)

(39) Is Patrick sure that he can park here? (question)

(40) If Patrick is sure that he can park here, he should tell Amy to park here too. (antecedent of a conditional)

(41) Patrick might be sure that he can park here. (epistemic modal)

All of (38)-(41) seem to communicate that Patrick has made some progress as to whether he can park here. To see that more clearly, we can imagine that it is common ground that Patrick is completely clueless about where he can park. We can imagine, for example, that Patrick knows nothing about parking rules or that he does not even drive. Then (38)-(41) would sound odd. The opinion inference projects from questions, antecedents of conditionals, and epistemic modals, which suggests that it is not an implicature, but rather a presupposition.[12]

5 Certainty is Not Maximal Credence

The credence account of certainty cannot explain the opinion inference; nor can the inference be explained away as an instance of scalar implicature. This is a problem for the credence account. One might worry that a linguistic puzzle like this is too minor to shake a foundational and otherwise compelling assumption in epistemology like the credence account. After all, the natural response to a linguistic puzzle is to look for a

[12] I focus on questioning, rather than negation, since negation has a special effect on scalar implicatures.

linguistic solution, not an epistemological one. This is a natural thought: we certainly don't want to throw out the baby with the bath water. However, while in many cases a linguistic puzzle remains just a linguistic puzzle, this is not one of those cases. The failure to explain the opinion inference is not an isolated problem, but a revealing symptom of what is wrong with the credence account of certainty. In this section, I show that the linguistic data we have seen actually points to a much deeper problem with reducing certainty to credence: certainty applies to both decisions and judgments, whereas credence only applies to judgments. In addition, there are at least two independent reasons why we should reject the credence account: (i) certainty is governed by norms that go beyond norms of credence; and (ii) our intuitive judgments about credence and certainty come apart.

5.1 Certainty about decisions

We have seen that the opinion inference arises when the QUD is a practical question, i.e., a question about what to do. This is not just an arbitrary fact about our language. Rather, it suggests that certainty applies to both practical decisions and theoretical judgments. Consider the following:

> **Ride**
> Alice and Bill are at a party together. Alice warns Bill that if he becomes drunk, he won't be able to drive home. Bill gets drunk anyway. Alice is trying to decide whether to give Bill a ride. On the one hand, Alice is convinced that giving him a ride is the right thing to do, and will lead to the best overall result. On the other hand, it is very late and she really doesn't want to drive him home. Plus, she is upset at Bill for being so imprudent. For a while, she is unable to decide.

Intuitively, Alice is uncertain. It is natural to imagine her reporting her predicament to a friend: "I am really not sure what to do. I know I should give him a ride home, but it's so late and he did this to himself!" If Alice is uncertain, what is her uncertainty about?

One might be tempted to look for a judgment, yet no judgment seems to be the right answer. Alice is not uncertain about a *normative or moral* judgment: she is sure that she should give Bill a ride home, that giving him a ride is the right thing to do, and so on. Alice is not uncertain about whether she *wants* to give Bill a ride, in the sense of desiring or preferring to drive him home, for she is sure that she does not. Perhaps Alice

is uncertain whether she will in fact give Bill a ride, but unless Alice suspects herself of weakness of the will, what Alice will do rests entirely on what Alice decides to do. If Alice is *just* uncertain what she will in fact do, she can easily come to be sure by making a decision. Yet Alice's uncertainty seems more stubborn than that: she wants to make up her mind, but she just can't.

It seems that Alice is not uncertain about a judgment at all. Instead, she is not sure *what to do*, despite being sure about all the relevant facts of the matter, i.e., about what is the case.[13] In other words, Alice is uncertain about her plan.

In general, we can be sure or unsure about what to do as well as what is the case. This is reflected in the fact that we often combine "sure" with infinitival interrogatives, i.e., phrases of the form "wh- to + verb." Typically, these express practical questions, questions about what to do, rather than theoretical questions, questions about what is the case. Consider:

(42) Carl is not sure what to do.

(43) Daisy is reasonably sure which move to make next.

(44) Emma is unsure whether to go to the party.

(45) Fabio is not sure whether to take the job offer or wait another year.[14]

Whatever certainty is, it should be an attitude we can hold toward both judgments and plans.

A defender of the credence account might suggest that there really are two distinct attitudes picked out by our ordinary use of "sure" and "certain." One applies only to judgments, which is maximal credence, and one applies only to decisions, which is something else—perhaps maximal expected utility. This suggestion risks

[13] I am using "facts" in a minimal sense. In this sense, all judgments concern facts. I am putting aside the distinction between judgments that are in some sense fully factual and ones that are not.

[14] One complication is that, while "sure" naturally embeds interrogatives under negation, it doesn't embed interrogatives as naturally without negation. For example, (42a) sounds considerably less natural than (42):

(42a) #Carl is sure what to do.

Adding a preposition makes it sound better, but still not completely natural:

(42b) ?Carl is sure about what to do.

This is a puzzling feature of "sure", but it is important to note that it applies to both finite interrogatives (denoting theoretical questions) and infinitival interrogatives (denoting practical questions). See:

(46) Carl is not sure what happened.
(46a) #Carl is sure what happened.
(46b) ?Carl is sure about what happened.

unnecessarily proliferating attitudes, but more importantly, it is hard to square with the fact that we frequently combine certainty about judgments and certainty about decisions in the same report. For example, consider (47):

(47) Gao is not sure whether it will rain tomorrow and whether to bring an umbrella.

Here the theoretical question of whether it will rain and the practical question of whether to bring an umbrella are under the scope of the same "sure." If there were no unified state of being sure that applies to both what is the case and what to do, it would be unclear what (47) could even mean. However, (47) sounds perfectly coherent and natural.

This is a problem for the credence account of certainty. We have credences toward judgments, not plans or decisions. After all, credences are degrees of beliefs, and belief is simply the wrong attitude for plans. Beliefs and credences are truth- or accuracy-apt, whereas decisions and plans are not.[15] The problem is that no amount of credence (even credence one) that one assigns to a possibility can amount to one's certainty about a plan to make that possibility obtain. For instance, no amount of credence that I assign to the possibility that the world will come to an end will amount to my being sure about *making* the world come to an end; conversely, from the fact that I do not have maximal credence that I *will* in fact manage to get to the train station on time, it does not follow that I am not sure whether *to* get to the train station on time. Credence is simply too narrow to capture certainty.

Now one might be tempted to think that, although the credence account fails to capture the practical dimensions of certainty, it does a great job capturing the theoretical one. However, I will now argue that credence fails to capture certainty even in the theoretical realm.

[15] An expressivist like Gibbard (2003) might agree with me that we can be sure about decisions, but argue that deciding to do X just is judging that X is the thing to do, or that X ought to be done, all things considered. If so, there may be no real distinction between decisions and judgments, and that opens up some room for the possibility of having credences toward decisions. Addressing the broader debate over the relationship between judgments and decisions is beyond the scope of this paper. For now I will assume that they are distinct attitudes, as this distinction seems intuitive. In particular, the equivalence between deciding to do X and judging that X is the thing to do seems difficult to maintain, since we can coherently judge that X is the thing to do without having made up our mind to do X, as in the following:

(48) I know quitting is the thing to do, but I just can't make up my mind to do it.

5.2 Norms of certainty

When are we justified in being sure? If certainty is credence one, then the Bayesians offer a simple answer: when conditionalizing on one's evidence yields probability one. That in turn depends on (a) the prior probability that one assigns to each possibility, (b) laws of probability, and (c) conditionalization. A particular kind of Bayesian might add the *Principal Principle*, according to which one's credence in a proposition should match the objective chance of that proposition, absent other, non-chancy evidence. On its own, this is a simple but powerful picture of norms of credence: having the right level of credence is just a matter of getting the probability right.

However, this picture appears too barren, if we are hoping to capture norms of certainty, for certainty is plausibly governed by at least two epistemic norms. First, certainty is governed by a *knowledge norm*:

> (KNC) One should be sure about a question only if one knows the answer to that question.

The normative connection between certainty and knowledge is hardly news. Ayer (1983) argues that knowledge gives one the right to be sure, for example. More recently, Goodman and Holguín (2022) argue that being sure is the least demanding doxastic attitude subject to a knowledge norm: i.e., one should only be sure that p if one knows that p.[16] They argue that we should accept KNC because it would explain both why (49) is infelicitous and, with the help of a certainty norm of assertion (CNA)—one should only assert p if one is sure of p—why (50) is infelicitous.

(49) #I am sure that it's raining, but I don't know that.

(50) #It's raining, but I don't know that.

They are infelicitous because uttering (49) amounts to flouting KNC, and uttering (50) amounts to flouting either KNC or the certainty norm of assertion.

Secondly, whether one is justified in being sure also depends on whether one has the right *kind* of evidence, not just the strength of one's evidence. This means that certainty is governed by an *evidential norm*:

[16] For this reason, Goodman and Holguín (2022) also reject the credence account of certainty, but their view differs from mine in a few respects. For example, they accept that certainty is a doxastic attitude, while I argue that certainty crosscuts the doxastic and non-doxastic distinction.

(ENC) One should not be sure about Q unless one has the right kind of evidence for Q.

Consider the following pair of cases (cf. Lawlor 2014, 5):

> **Lottery$_1$**
> Bill bought a lottery ticket. Bill knows that there are N tickets in total and therefore that his ticket has an extremely low probability of winning. On this basis, Bill comes to be sure that he lost, without checking what the winning ticket is.

> **Lottery$_2$**
> Bill bought a lottery ticket. Bill has no idea how many tickets there are in total. But he looks up the number of the winning ticket and sees that it is not his. On this basis, he comes to be sure that he lost.

Let's suppose that Bill's ticket did in fact lose. Let's also stipulate there is a small, non-zero chance (equal to 1/N) that Bill's checking in **Lottery$_2$** went wrong, for example because Bill misread the number. This means that given Bill's evidence, the probability that Bill's ticket lost is the same in both cases. In other words, Bill ought to have the exact same credence that his ticket lost in **Lottery$_1$** and **Lottery$_2$**.

However, intuitively there is a difference in whether Bill is justified in being sure. It seems fine for Bill to be sure in **Lottery$_2$**: that just involves trusting his ability to read the newspaper. But it seems less rational for Bill to be sure in **Lottery$_1$**: he has done nothing to check whether he is the winner, and is simply basing his certainty that he lost on the sheer number of tickets. He seems to be doing something epistemically irresponsible in coming to be sure that he lost without checking the winning ticket, even though he ought to have the same credence in both cases. Imagine if Bill decides to throw away the lottery ticket in **Lottery$_1$**. It is very natural to object by saying: "But you haven't even checked the winning number yet. How can you be sure that you lost?" Contrast that with **Lottery$_2$**: if Bill decides to throw away the ticket after checking the winning number, it would sound a bit delusional to say "But there might have been a mistake. How can you be sure that you lost?" This shows that certainty is sensitive to the *kinds* of evidence one has, not just the amount of credence it warrants.

Can the credence account as it stands make sense of KNC and ENC? It seems not. No combination of conditionalization, laws of probability, and the Principal Principle would generate a corresponding knowledge norm or evidential norm of credence. In fact, it would be a little puzzling if credence one or some other level of credence were subject to such norms. When an agent generates a credence in p in accordance with conditionalization and laws of probability, how is the fact that one knows or does not know p supposed to bear on that credence? Suppose I come to have credence 1. If my credence is fixed by my total evidence and the probabilistic norms governing rational belief revision, it is unclear what additional role knowledge could or should play. The same applies to the *kind* of evidence we have. Once my total evidence determines a probability, it is unclear what additional role is left for the *kind* of evidence to play in determining my credence.

Of course, there is some room for knowledge to play a role in constraining credence. Bayesianism is silent on what constrains the prior probabilities, so a defender of the credence account could always supplement norms of credence with some additional constraint on the priors, such that an agent should assign maximal credence to and only to propositions that they know.[17] In effect, this amounts to a knowledge requirement on maximal credence. However, this has a side effect: it risks making knowledge too demanding. Intuitively, one can know many things that are not entailed by their evidence. Consider, for example, perceptual knowledge or knowledge from memory. I know that I went to Switzerland in 2019 from memory, even though my evidence does not rule out the possibility that I am wrong. If I am a good probabilistic reasoner, I should not assign maximal credence to the proposition that I went to Switzerland in 2019. This places the credence account in a dilemma: either it makes knowledge too demanding, or it makes maximal credence too easy.

It would be even harder to accommodate ENC by putting a constraint on one's evidence. What counts as the right kind of evidence could differ depending on the subject matter of the proposition: the number of tickets in a lottery is not the right kind of evidence for whether Bill has lost, but it is the right kind of evidence for whether the

[17] Though it is controversial, we can understand Williamson (2000) as endorsing something like this.

lottery is fair. Putting a blanket ban on certain kinds of evidence, such as statistical evidence, would not constrain certainty with the right level of fine-grainedness.

More importantly, the issue is not just whether KNC and ENC are compatible with, or can be somehow accommodated by, the credence account of certainty; rather, it is whether the credence account can *explain* them: whether identifying certainty with credence one can shed light on why certainty should be constrained by knowledge and the kind of evidence one has. So if there is an alternative account of certainty that does explain these constraints, as I will argue there is, that is a strong reason to reject the credence account.

5.3 Mismatches between certainty and credence

Certainty comes in degrees. One can be *more* sure, *less* sure, *absolutely* sure, *pretty* sure, *extremely* unsure, and so on. What credence would one need to have in a proposition p to count as *maximally unsure* or *not sure at all* that p, according to the credence account? Let's consider two natural candidates: credence 0 and credence 0.5.

Clearly, it cannot be credence 0: being not sure at all that p should not entail being absolutely sure that not-p. I might have no clue at all whether it is going to rain. In that case, I am not sure at all that it will rain, and I am not sure at all that it won't.

However, it cannot be credence 0.5 either. Consider two cases:[18]

Horse Race
Four horses are racing: A, B, C, and D. Liz has credence 0.5 that horse A will win, credence 0.2 that B will win, 0.15 that C will win, and 0.15 that D will win.

Tiny Horse Race
Two horses are racing: A, and B. Maddy has credence 0.5 that horse A will win, and credence 0.5 that horse B will win.

Liz and Maddy have equal credence, credence 0.5. However, intuitively Liz might be fairly sure that horse A will win; by contrast, Maddy seems not sure at all that A will win, despite having the exact same credence that Horse A will win.

[18] These are reminiscent of Yalcin's (2010) cases, which Yalcin uses to show that ordinary speakers' intuitions about "probably" statements are affected by the salient alternatives. The data for this observation come from Windschitl & Wells (1998).

A defender of the credence account might propose a more complicated version of the view: instead of a particular level of credence, perhaps certainty amounts to a certain distribution of credences. That is, one is *maximally* uncertain iff one has equal credence in each of a question's answers, and one becomes more certain when one assigns higher credence to one of the answers. Here is a toy version of this approach:

> (Credence distribution account of certainty) Given a set of answers to a question, one's degree of certainty is determined by the gap between the answer that one has the highest credence in and the answer that one has the second-highest credence in. The larger the gap, the more certain one is.

The credence distribution account generalizes the credence account: the gap is maximally large when one has credence one in one of its answers and credence 0 in the rest. On top of that, it predicts that Maddy is more unsure than Liz, which seems plausible. However, now we simply encounter some new problem cases. For example, consider:

> **Rabbit Race**
> Six rabbits are racing. Sophie has credence 0.7 that Rabbit A will win, credence 0.3 that Rabbit B will win, and credence 0 that any of the other 4 rabbits will win. Meanwhile, Jenny has credence 0.5 that Rabbit A will win, and credence 0.1 that each of the other 5 rabbits will win.

According to our toy credence distribution account, Sophie and Jenny should be equally certain about which rabbit will win, given that the difference between their most confident answer and second-most confident answer is the same, namely, 0.4. However, intuitively Sophie might be more sure than Jenny, given that she has ruled out 4 out of the 6 candidates.

In addition, it is not clear that an even distribution of credences guarantees maximal uncertainty either. Consider:

> **Bridge**
> Jim is investigating whether a bridge is safe. He has done extensive research, and has determined that whether it is safe just depends on whether it was built before 2003. Jim does not know whether it was—for all he knows, it is equally likely to have been built before 2003 as after, but he knows where to find the record for the bridge. In fact, Jim is now on the way to the city archives to look at the records.

Jim has credence 0.5 that the bridge was built before 2003 and credence 0.5 that it was built after. As a result, he has credence 0.5 that the bridge is safe and credence 0.5 that it isn't. (We can imagine that all his other evidence cancels out in terms of credence.) We certainly don't want to say that Jim is sure about whether the bridge is safe, but would we say that Jim is *not sure at all* whether the bridge is safe? After all, Jim has made a lot of progress in figuring out whether it is safe, and he is steps away from being completely sure. If we feel any hesitation ascribing Jim maximal uncertainty, it should make us question the credence distribution account. After all, if we are just looking at Jim's credences, he is as uninformed as it goes.

Admittedly, what we are working with is just a toy version of the credence distribution account. There are many other potential maneuvers we can try, and perhaps one of them might just give us all the results we want. One reason to be pessimistic, though, is that similar problems arise for predicting judgments about "probably." As Yalcin (2010) points out, it is notoriously difficult to give a semantics of "probably p" in terms of the probability distribution over p and its salient alternatives which gets all the intuitive judgments right. So far, it seems that reducing certainty to credence would put us in a bad position to articulate what it takes to be maximally uncertain.

6 Certainty is Taking Inquiry to Be Closed

I have shown that failure to explain the opinion inference is but one manifestation of the ways in which certainty and credence come apart. Credence is ill-suited to play the role of certainty, and we are in need of an alternative. But if certainty is not maximal credence, what is it? In this section, I propose my alternative inquiry account of certainty: to be sure about a question is to *take that question to be settled* (or equivalently, to *take inquiry into that question to be closed*). This state is constrained by, but ultimately orthogonal to, one's credences.

6.1 Taking Inquiry to Be Closed

We inquire into all sorts of questions: where the nearest grocery store is, who will present next in the reading group, what makes life meaningful, and so on. Some inquiries are involved and sophisticated: for example, a philosopher's inquiry into

whether compatibilism is true will involve inquiring into many related questions—what a plausible account of free will is, what compatibilism entails, and which objections have been raised and addressed already. Some inquiries are simple and low effort: inquiring into when the movie starts might just involve checking your ticket. In fact, some inquiries might be happening mostly outside of our intentional control: we might sometimes engage in what Siegel (2025) calls "wandering inquiry," where we make progress in inquiry by mind wandering.

Many take there to be a *causal* or *normative* connection between psychological certainty and inquiry: when one is certain, one *will* or *ought* to close inquiry. I propose that there is instead a *constitutive* connection between certainty and inquiry:

> (Inquiry account of certainty) The state of being certain with regard to a question *just is* the state of taking inquiry into that question to be closed, and the state of being uncertain *just is* the state of taking inquiry to be open.

What is the state of taking inquiry to be closed? Most importantly, it is a *settled* attitude toward a question. When our mind is *made up*, when we decide that something is *conclusive*, when we think a matter is *settled*, when we are *set* on a decision, and so on, we take inquiry to be closed. While taking inquiry to be closed is a way of stopping inquiry, it is different from other ways of stopping inquiry, such as suspending or abandoning inquiry: one can stop inquiring for a variety of reasons without taking inquiry to be closed; but to take inquiry to be closed, one has to take oneself to have achieved the aim of inquiry. So long as one takes inquiry to be closed, one has no disposition to inquire. This is another respect in which someone who takes inquiry to be closed might differ from someone who merely suspends or abandons inquiry: the latter might still have a disposition to inquire, which is overpowered by other dispositions, while the former has none to begin with. Taking inquiry to be closed entails that one is at peace, as it were, with one's answer to the question and feels no pull at all toward further inquiry. This means that taking a question, q, to be settled entails (i) endorsing an answer to q (in virtue of forming a belief or an intention), (ii) being disposed to rely on this answer as a premise in further deliberation, and (iii) thinking that one knows it.[19]

[19] One might worry about what it is to know the answer to a practical question. I think spelling out what it is to know the answer to a practical question is an interesting task, but beyond the scope of this paper. I take for granted that we do know such answers, since we frequently describe each other as knowing (or not knowing) what to do: e.g.,

In particular, (iii) is important for distinguishing taking inquiry to be closed from a more temporary and conditional attitude like *acting as if* inquiry is closed, or treating inquiry as closed *for the purpose of the conversation*. A member of the jury might choose to yield a verdict of not guilty, for example, simply because they do not have enough permissible evidence in court to prove that the suspect is guilty. Once outside the legal context, however, the member of the jury will not take inquiry to be closed. Even when delivering the verdict, they might be wondering privately whether the suspect really did it, precisely because they do not take themselves to know the answer. Because one has to take oneself to know the answer in order to take inquiry to be closed, one does not have full voluntary control over whether one is sure: one cannot simply decide to be sure in the same way that one can decide to raise one's arm.

I hope I have said enough to pick out a familiar state, the state of taking inquiry to be closed. My proposal is that this state *just is* certainty.

6.2 Explaining the opinion inference

The inquiry account of certainty explains why, with respect to the opinion inference, certainty talk patterns with settling talk, not with knowledge or belief talk. After all, if certainty just is taking a question to be settled, then it is natural that our certainty talk and settling talk would exhibit the same patterns.

My account also explains why denials of certainty would generate an opinion inference at all. We have established earlier that the opinion inference is a presupposition. It is generally accepted that presuppositions are triggered by particular linguistic expressions. In other words, "sure" is a presupposition trigger, and the presupposition it triggers is that the subject has made some progress in inquiry. Now we are in a better position to understand why that would be the case. If certainty is taking inquiry to be closed, then our notion of certainty does not really track credence, but rather the status of one's inquiry. It is natural, then, that applying the notion of certainty to a subject would presuppose something about the subject's relation to inquiry.

"Jonathan didn't know what to do, but now he does." For an in-depth discussion of knowing what to do, see Jerzak & Kocurek (2025).

Here is a more detailed story. "Sure" is a gradable adjective: one can be *more or less* sure, *very* sure or only *somewhat* sure.[20] According to the standard semantics, gradable adjectives are associated with scales.[21] For example, "tall" is associated with a scale that consists of degrees of tallness. In general, using a gradable adjective (whether in its positive form or embedded) presupposes that the object in question can be located on the corresponding scale. Describing something as tall or not tall, for example, presupposes that it is somewhere locatable on the shortness/tallness scale, i.e., that it is something with height. This explains why it would be odd to apply "tall" to something without height at all, or "full" to something that is not a container:

(51) #This piece of paper is tall.
(52) #This piece of paper is not tall.
(53) #The front door is full.
(54) #The front door is not full.

According to my proposal, our notion of certainty tracks progress in inquiry. It is natural, then, to take the scale for "sure" to consist of degrees of progress in inquiry. Just as using "tall" presupposes that one can be located on the shortness/tallness scale, using "sure" or "certain" presupposes that one can be located on the scale of progress in inquiry. In other words, it presupposes that one has made *some* progress in inquiry.

In addition, "sure" is somewhat special because it is an *upper-closed absolute* gradable adjective: its corresponding scale has an upper bound, but no lower bound. This is shown by the fact that we can modify "sure" but not "unsure" with degree modifiers that indicate a closed end point:

(55) Jimmy is absolutely/completely sure.
(56) #Jimmy is absolutely/completely unsure.

Since the scale for "sure" tracks progress in inquiry, this means that there isn't a point on the scale that corresponds to the absolutely minimal degree of progress in inquiry. In other words, every point on the scale involves *some* progress in inquiry. To be evaluable

[20]Everything I say here applies to "certain" as well.
[21] Note that this is a different kind of scale from quantitative scales. The latter are ordered sets of expressions, but the scales for gradable adjectives are degrees or orderings of the relevant property.

along the dimension of certainty, therefore, requires that one has made some progress in inquiry.

Finally, the inquiry account makes sense of why the opinion inference arises when the QUD is practical and, more generally, why we can apply certainty to both practical and theoretical questions. We inquire into practical as well as theoretical questions: just as we deliberate about what is the case, we also deliberate about what to do. Insofar as inquiry into theoretical questions is continuous with inquiry into practical questions, taking inquiry to be closed is already a unified state between the theoretical and the practical domains. Since certainty is taking inquiry to be closed, it naturally applies to both theoretical judgments and practical decisions. And invoking the notion of certainty presupposes that one has made progress in inquiry, whether that inquiry concerns what is the case or what to do.

6.3 Norms of certainty

The credence account struggles to explain why certainty is constrained by both knowledge and the kind of evidence one has. The inquiry account, by contrast, offers a unified explanation. According to the inquiry account, certainty is taking inquiry to be closed, so norms of certainty are really norms of closing inquiry. There are *independently motivated* norms for closing inquiry that explain the knowledge and evidential norms of certainty. Kelp (2021), for example, takes knowledge to be necessary and sufficient for appropriately closing inquiry. Even more (Jane Friedman, 2024; Kelp, 2021; Carter and Hawthorne, 2024) take knowledge to be the aim of inquiry. Intuitively, one should not take a goal-directed activity to be complete until one has reached the goal. So given that inquiry aims at knowledge, we have the following knowledge norm of inquiry:

> (Knowledge norm of inquiry) One should take inquiry into q to be closed only if one knows the answer to q.

In addition, Daniel Friedman (forthcoming) argues that different inquiries come with their associated kinds of evidence. For example, scientific research is answerable to experiments, crime investigations are answerable to fingerprints and autopsies, philosophy is answerable to arguments. Without the right *kind* of evidence, it is

irresponsible to close an inquiry, even if one's evidence is otherwise quite strong. This gives us the evidential norm of inquiry:

> (Evidential norm of inquiry) One should take inquiry into q to be closed only if one has the right kind of evidence.

If certainty is taking inquiry to be closed, these two norms of inquiry straightforwardly generate, and thus explain, the corresponding norms of certainty.

7 Knowledge and Certainty

I have argued that our notion of certainty is constitutively tied to inquiry. Knowledge, on the other hand, is not constitutively tied to inquiry in this way. This difference between knowledge and certainty explains why their negations behave differently. In attributing knowledge or ignorance to someone, we are not necessarily interested in specifying where they are in their inquiry, or whether they take their inquiry to be closed. Instead, we could be interested in whether they have the right kind of access to truth, or whether they can confer knowledge to us. This explains why denying knowledge does not generate the opinion inference: doing so does not presuppose anything about the subject's position in inquiry.

Apart from explaining the asymmetry between denials of knowledge and denials of certainty, my view that certainty is taking inquiry to be closed also has broader implications for understanding the relation between knowledge and certainty. On my view, knowledge does not entail certainty: certainty, as I have argued, is a matter of taking an inquiry to be closed, and one can know the answer to a question without treating inquiry into that question as closed.[22] However, since there are normative connections between knowledge and inquiry, there are normative connections between knowledge and certainty. We have seen that the inquiry account of certainty explains why there is a knowledge norm of certainty. We might think that, in addition, certainty has a normative connection with knowledge in the other direction: it is plausible that knowledge is normatively in tension with uncertainty.

[22] It is beyond the scope of this paper to defend the idea that knowledge does not constitutively entail certainty. For a detailed defense, see Stanley (2008).

Following Jane Friedman (2017), I take it that knowing the answer to a question is *normatively in tension* with inquiring into that question. In other words, inquiry is subject to the ignorance norm:

> (Ignorance norm of inquiry) If one knows the answer to a question, then one ought not to inquire into that question.

The ignorance norm of inquiry is prima facie plausible. After all, inquiry is a goal-directed activity, where the goal is to know the answer to the question. In general, once the goal of an activity has been achieved, it would be irrational to continue that activity. So once the subject knows the answer to a question, it would be irrational for them to continue inquiring into that question.[23] In Whitcomb's words, inquiring after attaining knowledge is like "continuing to eat after being nourished" (Whitcomb 2010, p. 674).

So according to the inquiry account of certainty, knowledge does not entail certainty, but it is normatively in tension with uncertainty. This is an advantage of my view, because it allows us to explain some puzzling linguistic data. Consider the following sentences:[24]

(57) #I know when the movie starts, but I am not sure (when it starts).
(58) #I know that the movie starts at 7, but I am not sure (that it does).
(59) Hannah knows when the movie starts, but she is not sure (when it starts).
(60) Hannah knows that the movie starts at 7, but she is not sure (that it does).
(61) I knew when the movie starts, but I wasn't sure (when it starts).
(62) I knew that the movie starts at 7, but I wasn't sure (that it does).

(57) and (58) sound infelicitous, whereas (59)-(62) sound much better. To explain the contrast between them, we cannot simply take knowledge to entail certainty, since that would predict incorrectly that (59)-(62) are also infelicitous.[25]

Stanley (2008) explains the first-/third-person contrast between (58) and (60) by appealing to CNA. Given CNA and the factivity of knowledge, the first part of (58) is only

[23] Jane Friedman (2017) argues for a stronger version of the ignorance norm along these lines.
[24] To my knowledge, Unger (1975) was the first to make the observation that sentences like (58) are infelicitous. He cites it as evidence for the claim that knowledge requires certainty.
[25] Stanley (2008) makes this point as an objection to Unger's (1975) argument that the infelicity of sentences like (58) show that knowledge requires certainty.

appropriate if the speaker is certain that the movie starts at 7; however, the second part of (58) denies precisely that. So (58) is infelicitous because it flouts a norm. On the other hand, asserting (60) would not violate any norms: it is the speaker who has to be certain that the movie starts at 7, not the subject.

But Stanley's explanation only goes so far: while it explains the asymmetry between (58) and (60), it cannot explain the asymmetry between (57) and (59). In asserting (58), the speaker avows, if not asserts, the proposition that the movie starts at 7; but in asserting (57), the speaker does not avow any particular proposition about when the movie starts. So uttering (57) would not be a flouting of CNA, even if CNA is true. But (57) seems to be infelicitous in exactly the same way as (58): the only difference lies in whether "know" and "sure" embed declarative complements or interrogative complements. It would be surprising if the infelicity of (57) and that of (58) require distinct explanations.

In addition, CNA faces some serious objections. Mandelkern and Dorst (2022) point out, for example, that if CNA is true, "How are you sure?" should be a legitimate follow-up to any assertion; but compared to "How do you know?" or "Why do you think that?", "How are you sure?" seems presumptive. It is much better to ask "Are you sure?", which does not presuppose that the speaker *is* sure.

(63) Ann: The movie starts at 7. We should get ready soon.
(63a) How do you know?
(63b) Are you sure?
(63c) #How are you sure?/What makes you sure?/How can you be sure?[26]

This observation suggests that we do not generally expect a cooperative speaker to be sure of what they assert, though we do expect them to believe and sometimes know what they assert. CNA seems too strong.[27]

My view that certainty is taking inquiry to be closed offers a more unified explanation for (57)-(62) that does not appeal to CNA. Instead, it appeals to the ignorance norm of inquiry. In uttering the first part of (57) or (58), the speaker reports

[26] The intuition is not shared by everyone. Goodman and Holguín (2022, p. 635) actually think that 'What makes you sure?' is a felicitous follow up to an assertion.
[27] For a more comprehensive discussion of CNA, see Mandelkern and Dorst (2022).

that they have knowledge—i.e., that the aim of their inquiry has been achieved; in uttering the second part of (57) or (58), the speaker reports that they are still uncertain—i.e., that they are still taking inquiry to be open. (57) and (58) are infelicitous, therefore, because they amount to flouting the ignorance norm of inquiry. To borrow Whitcomb's (2010) analogy, it is like saying "I am full, but I am still eating." On the other hand, (59)-(62) are fine because uttering them amounts to pointing out that *someone else* (including one's past self) violates the ignorance norm of inquiry.

8 Conclusion

In this paper, I have introduced a puzzle about the asymmetry between denials of knowledge and denials of certainty: the latter generates an opinion inference—that the subject has made some progress in inquiry—while the former does not. This simple observation has been overlooked in the existing literature on certainty, but as I have shown, it is an important clue to the epistemological question of what it takes to be sure. Not only does the standard credence account of certainty fail to predict the opinion inference, but it also struggles to explain how we can be sure about both judgments and decisions, why we shouldn't be sure when we lack knowledge or the right kind of evidence, and what maximal certainty or uncertainty amounts to.

In place of the credence account, I have proposed my alternative inquiry account of certainty: certainty is taking inquiry to be closed, and uncertainty is taking inquiry to be open. This account also helps clarify the relation between knowledge and certainty: knowledge does not entail certainty, but it is normatively in tension with uncertainty. This provides a unified explanation of some otherwise puzzling linguistic data—"S knows X, but S is not sure" sounds Moore-paradoxical as a first-personal present-tense report but not as a third-personal or past-tense one.

The inquiry account allows us to understand certainty as a distinctive way of relating to inquiry, and thereby as an important attitude in its own right. As a result, it also transforms some longstanding questions about certainty and other epistemic attitudes. To understand the connection between certainty and belief, for example, we cannot treat them as simply attitudes at different points of the same scale. Instead, we must ask how belief relates to inquiry.

More generally, I see my account as contributing to a larger shift toward placing inquiry at the center of epistemology. As Jane Friedman (forthcoming) argues, epistemologists should investigate questions about how to inquire, in addition to questions about what to believe. Among other things, my view of certainty demonstrates what that could look like: by understanding certainty in terms of inquiry, questions about certainty become questions about one's relation to inquiry.

References

Atlas, Jay. D., & Levinson, Stephen. C. (1981). "It-Clefts, Informativeness and Logical Form: Radical Pragmatics (revised standard version)." In P. Cole (Ed.), *Radical Pragmatics* (1–62). New York: Academic Press.

Ayer, Alfred J. (1983). Knowing as having the right to be sure. In Steven Davis, *Causal Theories of Mind: Action, Knowledge, Memory, Perception and Reference*. Berlin, Boston: De Gruyter. pp. 131–134.

Beddor, Bob. (2020). "New Work for Certainty." *Philosophers' Imprint*, 20(8).

Carter, Sam & Hawthorne, John. (2024). Dogmatism and inquiry. *Mind* 133 (531):651–676.

Clarke, Roger. (2013). Belief Is credence one (in context). *Philosophers' Imprint* 13:1–18.

Friedman, Daniel. (2025). "Intentions and Inquiry." *Mind*, 134(533):85–106.

Friedman, Daniel. (forthcoming). Cooperation and Shared Inquiry. *Australasian Journal of Philosophy*.

Friedman, Jane. (2013). "Question‐Directed Attitudes." *Philosophical Perspectives*, 27(1), 145–174.

Friedman, Jane. (2017). "Why Suspend Judging?" *Noûs*, 51(2):302–326.

Friedman, Jane. (2024). The aim of inquiry? *Philosophy and Phenomenological Research* 108 (2):506–523.

Friedman, Jane. (forthcoming). Zetetic epistemology. In Baron Reed & A. K. Flowerree, *Towards an Expansive Epistemology: Norms, Action, and the Social Sphere*. Routledge.

Gazdar, Gerald. (1979). "Pragmatics, Implicature, Presupposition and Logical Form." *Critica*, 12(35).

Gibbard, Alan. (2003). *Thinking How To Live*. Cambridge, MA: Harvard University Press.

Goodman, Jeremy & Holguín, Ben. (2022). Thinking and being sure. *Philosophy and Phenomenological Research*, 106: 634–654.

Grice, Paul. (1989). *Studies in the Way of Words*. Cambridge, MA: Harvard University Press.

Halpern, Joseph Y. (1991). The relationship between knowledge, belief, and certainty. *Ann Math Artif Intell*, 4, 301–322.

Hirschberg, Julia L. B. (1991). *A Theory of Scalar Implicature*. New York: Garland.

Horgan, Terry. (2017). Troubles for Bayesian formal epistemology. *Res Philosophica* 94 (2):233–255.

Horn, Lawrence. R. (1972). *On the Semantic Properties of Logical Operators in English*, Ph.D. Thesis, UCLA.

Jackson, Elizabeth. (2020). Belief, credence, and evidence. *Synthese* 197, 5073–5092.

Jerzak, Ethan & Kocurek, Alexander W. (2025). Knowing what to do. *Noûs*, 59, 160–190.

Kelp, Christoph. (2021). Theory of inquiry. *Philosophy and Phenomenological Research* 103 (2):359–384.

Lawlor, Krista. (2014). Exploring the stability of belief: resiliency and temptation. *Inquiry*, *57*(1), 1–27.

Levinson, Stephen. (2000). *Presumptive meanings: the theory of generalized conversational implicature*. Cambridge: MIT Press.

Mandelkern, Matt & Dorst, Kevin. (2022). "Assertion is Weak." *Philosophers' Imprint*, 22 (19): 1–20.

Moss, Sarah. (2019). Full belief and loose speech. *Philosophical Review* 128 (3):255–291.

Pettigrew, Richard. (2023). Credences. In *The Routledge Encyclopedia of Philosophy*.

Roberts, Craige. (2012). "Information Structure in Discourse: Towards an Integrated Formal Theory of Pragmatics." *Semantics and Pragmatics*, 5(6): 1–69.

Siegel, Susanna. (2025). "Wandering Inquiry." *Journal of Philosophy*.

Simons, Mandy; Beaver, David; Tonhauser, Judith & Roberts, Craige. (2010). What projects and why. *Semantics and Linguistic Theory* 20:309-327.

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

Williamson, Timothy. (2000). *Knowledge and its Limits*. Oxford University Press, USA.

Windschitl, Paul D. & Wells, Gary L. (1998). The alternative-outcomes effect. *Journal of Personality and Social Psychology*, 75, 1411–1423.

Unger, Peter. (1975). *Ignorance: A Case for Skepticism*. Oxford: Oxford University Press.

Whitcomb, Daniel. (2010). "Curiosity was Framed." *Philosophy and Phenomenological Research*, 81(3): 664–687.

Yalcin, Seth. (2010). Probability operators. Philosophy Compass 5 (11):916–37.