KNOWLEDGE AND CERTAINTY

Jason Stanley
Rutgers University

What is the connection between knowledge and certainty? The question is vexed, in part because there are at least two distinct senses of "certainty". According to the first sense, *subjective certainty*, one is certain of a proposition if and only if one has the highest degree of confidence in its truth. According to the second sense of "certainty", which we may call *epistemic certainty*, one is certain of a proposition p if and only if one knows that p (or is in a position to know that p) on the basis of evidence that gives one the highest degree of justification for one's belief that p. The thesis that knowledge requires certainty in either of these two senses has been the basis for skeptical arguments. For example, according to one kind of skeptical argument, knowledge requires epistemic certainty, and being epistemically certain of a proposition requires having independent evidence that logically entails that proposition. Since we do not have such evidence for external world propositions, we do not know external world propositions. According to another kind of skeptical argument, due to Peter Unger (1975), knowledge requires subjective certainty, and we are never subjectively certain of any proposition. So, we never know any proposition.

Some authors have responded to these skeptical arguments by adopting fallibilism about certainty, the doctrine that having the highest degree of justification for a belief does not involve the possession of logically entailing, non-question begging grounds for that belief (Miller (1978), Klein (1981, Chapter 3)). But my interest in this paper does not lie in rebutting skeptical arguments based upon the assumption that knowledge entails certainty. Rather, my purpose is to establish that knowledge does not require certainty in either of these two senses. Even if we are certain of many things, knowing that p does not entail subjective or epistemic certainty.

Since the claim that knowledge requires certainty (in either sense) is closely associated with detrimental conclusions, the central case for the

34   Jason Stanley

connection has typically been made on intuitive rather than normative grounds. For example, it seems that someone is speaking falsely when they assert "I know that Bush is president, but I'm uncertain of it." The kinds of arguments I will focus on in this paper involve such appeals to our intuitions about various claims about the relation between knowledge and certainty. So I will not be giving a normative argument about the nature of knowledge or the nature of the various kinds of certainty. My goal is rather to contrast the *intuitive* grounds for the claim that knowledge requires certainty (in either sense) with the *intuitive* grounds for the claim that knowledge requires truth, belief, or justification. I will argue that the evidence that has been marshaled in support of the thesis that knowledge requires certainty is instead far better explained by the hypothesis that we adhere to norms that connect subjective and epistemic certainty with the speech act of assertion.

In the hands of the contextualist in epistemology, the arguments for the thesis that knowledge requires certainty become arguments for the thesis that "knowledge" requires "certainty". My own view is that contextualism about knowledge attributions is false, but that "certain" is a context-sensitive adjective. So I am a contextualist about "certain", but not about instances of "know that p". Be that as it may, my arguments that knowledge requires certainty are also arguments that "knowledge" does not require "certainty".

1.

Peter Unger's argument for his novel form of skepticism involves two premises, only the first of which will be of interest to me in this paper. The first premise is that "If one knows, then it is all right for one to be certain" (1975, p. 98). We do not need to be concerned here with what Unger means by the somewhat normative notion of being "all right for one to be certain". My concern is rather with Unger's claim that knowledge in fact requires certainty. As Unger writes (Ibid.):

> The very particular idea that knowing entails its being all right to be certain is suggested, further, by the fact that knowing entails, at least, that one *is* certain. As we saw in section 9 of the preceding chapter, that this is a fact is made quite plain by the inconsistency expressed by sentences like 'He really *knew* that it was raining, but he *wasn't* absolutely *certain* it was.' Such a sentence can express no truth: if he wasn't certain, then he didn't know.

Unger's focus is on the notion of *subjective* certainty, rather than *epistemic* certainty. Though the term "certain" is ambiguous between subjective and epistemic certainty, there are constructions in which it only can be read as the former than the latter, and vice-versa. When one speaks of a *person being*

certain of a proposition, it is subjective certainty that is at issue. In contrast, when one speaks of a *proposition being certain*, it is epistemic certainty that is at issue. So, in a sentence such as "John is certain that Bush is president", it is subjective certainty that is at issue, whereas when one says "It is certain that Bush is president", what is at issue is the epistemic certainty of the proposition that Bush is president, relative to one's own body of evidence. Unger's arguments concern subjective certainty, as do his examples.

How do we convince our undergraduates that knowledge requires truth, belief, and justification? One way is via consideration of the oddity of sentences such as "He knew that snow is purple, even though snow isn't purple", or "He knew that snow is white, but he didn't believe that snow is white", or "He knew that snow is white, but he had no reason to believe it". That is, the intuitive basis for the thesis that knowledge requires truth, belief, and justification is the oddity of instances of the schemas in (1)-(3):

(1) X knows that p, but p is false.
(2) X knows that p, but X doesn't believe that p.
(3) X knows that p, but X has no reason to believe that p.

Unger's argument that knowledge requires subjective certainty is of the same sort. Just as instances of (1)-(3) sound odd, so instances of (4) sound odd:

(4) I know that p, but I'm not certain that p.

Unger is right that instances of (4) sound odd. Given that the oddity of instances of (1)-(3) is good evidence that knowledge requires truth, belief, and justification, the oddity of (4) provides good evidence that knowledge requires subjective certainty.

Unger's focus is on the notion of subjective certainty. But there are analogous arguments to Unger's for the conclusion that knowledge requires epistemic certainty. Instances of (5) are just as odd as instances of (1)-(4):

(5) I know that p, though it isn't certain that p.

In fact, an utterance of a sentence such as (6) is also quite odd:

(6) He knows that Bush is president, though it isn't certain that Bush is president.

Together, these facts provide an argument of the very same form for the conclusion that knowledge requires *epistemic* certainty, in addition to *subjective* certainty.

There are reasons philosophers have given to reject connections between knowledge and justification, or even knowledge and belief. So, the fact that

instances of (2) and (3) sound odd might be part of a folk theory of knowledge that we may reject on philosophical grounds. Nevertheless, these facts provide evidence that, according to our folk theory of knowledge, knowledge does require belief and justification. Many philosophers have a similar attitude towards the connection between knowledge and certainty. They grant that the oddity of instances of (4) and (5) provide evidence that our folk theory of knowledge connects knowledge and subjective and epistemic certainty. They argue that we should instead adopt a fallibilist conception of knowledge that severs these connections.

But the case that knowledge requires certainty in either sense was never meant to establish that knowledge *ought* to require certainty. Rather, the argument is that the oddity of instances of (4) and (5) shows that we intuitively think that knowledge does *in fact* require certainty. Most epistemologists who have emphasized such connections have used it to argue for distressing conclusions. Their attitude towards attempts to explicate a fallibilist notion of knowledge that lacks this connection to certainty has been to emphasize that the connection between knowledge and certainty is so deeply ingrained, that any attempt to eliminate it will not preserve central features of the knowledge relation (e.g. the discussion of fallibilism in Lewis (1996)).

2.

There are a number of normative considerations in favor of the thesis that knowledge requires truth, belief, and justification. But it also has an intuitive basis. It clearly is odd to speak of someone knowing something clearly false, or knowing something without believing it. The case for the alleged relation between knowledge and certainty does not rest upon some desirable feature of the knowledge relation. Typically, it is a premise in arguments for skeptical conclusions. Because of this, we should expect that the descriptive case for the link between our actual concept of knowledge and the notion of certainty is at least as strong as the descriptive case for the constitutive link between knowledge and truth.

If knowing a proposition requires that proposition to be true, we would expect (7) to sound like an assertion of a trivial conceptual truth and (8) to sound like an assertion of an obvious falsity:

(7) Everything anyone knows is true.
(8) There is something someone knows that isn't true.

(7) is obviously true, and (8) obviously false. Similarly, if knowing a proposition requires believing that proposition, then we should expect (9) to be a trivial truth and (10) to be obviously false:

(9) Everything someone knows she believes.

(10) There is something someone knows that she doesn't believe.

Finally, if knowing a proposition requires having evidence for that proposition, we would expect (11) to sound like a trivial truth and (12) to sound obviously false:

(11) If someone knows something, she has a reason to believe it.

(12) There is something someone knows that she doesn't have any reason to believe.

An assertion of (11) certainly seems true, and (12) seems false.

If it is intuitively obvious that knowledge requires subjective certainty, we should expect (13) and (14) to seem like banal truths and (15) to seem obviously false:

(13) I'm certain of everything I know.

(14) Everyone is certain of everything she knows.

(15) There are some things I know, of which I'm only fairly certain.

However, (13) and (14), unlike (7), (9), and (11), do not sound like banal truths. An utterance of (15) also does not share the obvious sense of falsity of (8), (10), and (12). Similarly, if knowledge requires epistemic certainty, we should expect (16) to be a banal truth, on a par with (7) and (9), and we should expect (17) to seem clearly false, on a par with (8), (10), and (12):

(16) Everything I know is certain to be true.

(17) There are some things I know, which are only fairly certain to be true.

But (16) does not seem like a banal truth, and (17) seems perfectly in order.

A further dissimilarity between the case for the factivity of knowledge (the entailment from knowledge to truth) and the case for the entailment from knowledge to *subjective* certainty involves the existence of an asymmetry in our intuitions between first person and third person reports. Unger is right that an instance of (4) sounds odd:

(4) I know that p, though I'm not completely certain that p.

But third-person reports of the form of (4) are not as odd as utterances of (4). For example, compare (18) with (19):

(18) John knows that Bush is a Democrat, though Bush isn't a Democrat.

38   Jason Stanley

(19) John knows that Bush is a Republican, though, being a cautious
    fellow, he is only somewhat certain of it.

(18) is obviously false. By contrast, (19) is not at all obviously false. So there
is an asymmetry between first-person and third-person ascriptions in the
case of the alleged relation between knowledge and certainty, an asymmetry
that is lacking in the case of the relation between knowledge and truth (or
knowledge and belief).

There is an asymmetry between first and third person reports of
knowledge without subjective certainty. But this asymmetry is not present
where "certain" has an epistemic use. For example, (21) is just as odd as (20):

(20) I know that Bush is president, though it's not certain that he is.
(21) John knows that Bush is president, though it's not certain that he
    is.

So, while there is an asymmetry between first and third person reports
of *subjective* certainty. But this asymmetry disappears in the case of *epistemic*
certainty.

If one proposition obviously entails another, it will feel redundant to
follow an assertion of one with an assertion of the other. So, redundant
conjunctions provide evidence of entailments. In the case of the relation
between knowledge and truth, we clearly see such evidence of entailment, as
is witnessed by the oddity of (22):

(22) I know that Bill came to the party. In fact, he did.

We also see similar evidence in the case of the relation between knowledge
and belief, as (23) is just as odd as (22):

(23) I know that Bill came to the party. In fact, I believe he did.

The reason (22) and (23) are odd is that knowing entails truth and belief. No
new information is conveyed by assertions of the second sentences in (22)
and (23). So, the utterances seem pointless, and the discourses odd.

In contrast, we see no such evidence of entailments in the case of
subjective and epistemic certainty. In both cases, there is no similar sense
of redundancy:

(24) I know that Bill came to the party. In fact, I'm certain that he did.
(25) I know that Bill came to the party. In fact, it's certain that he did.

The discourses in (24) and (25), in contrast to those in (22) and (23), do
not seem odd at all. The assertions of the second sentences seem to add

new information to the information expressed by "I know that Bill came to the party". If knowledge entailed subjective and epistemic certainty, this fact would be a mystery. The discourses in (24) and (25) would be just as odd as the discourses in (22) and (23). If knowledge does not entail subjective and epistemic certainty, this fact is explicable.[1]

Another way we convince our undergraduates that knowledge requires truth is we ask them to consider claims such as:

(26) John knows that snow is purple.

(26) is clearly false. Similarly, if it is quite clear that John doesn't at all believe that snow is white, then one cannot ascribe to him the knowledge that snow is white. But correspondingly lucid intuitions are lacking in the case of the alleged connection between knowledge and certainty. It does no violence at all to our use "know" to ascribe knowledge of a proposition to someone who is somewhat uncertain of it. As Unger points out (Unger, 1975, p. 83-4):

> Thus typical in the contemporary literature is this sort of exchange: An examiner asks a student when a certain battle was fought. The student fumbles about and, eventually, unconfidently says what is true: The Battle of Hastings was fought in 1066. It is supposed, quite properly, that this correct answer is a result of the student's reading. The examiner, being an ordinary mortal, and so unconfident of many things himself, allows that the student knows the answer; he judges that the student knows that the Battle of Hastings was fought in 1066. Surely, it is suggested, the examiner is correct in his judgment, even though this student clearly isn't certain of the thing...

Unger's example shows that we easily accept attributions of knowledge to people who are not subjectively certain.

The intuitions linking knowledge and subjective and epistemic certainty seem to provide a case for an entailment between knowledge and certainty parallel to the entailment between knowledge and truth, or knowledge and belief. But once we canvass a broader range of cases, the parallel between knowledge and certainty, on the one hand, and knowledge and truth, or knowledge and belief, on the other, becomes considerably less compelling. There is nothing odd about ascribing knowledge of a proposition to someone who has less than the highest degree of confidence in her belief in it. In contrast, it is clearly incorrect to ascribe knowledge of a false proposition to someone. In so doing, one has asserted something false. There are many asymmetries between the relation that holds between knowledge and truth, on the one hand, and the alleged relation between knowledge and certainty, on the other. Though there are a series of complicated facts that Unger has drawn to our attention that require an explanation, it is now evident that the hypothesis of a constitutive connection between knowledge and subjective or epistemic certainty is ill-suited to be part of it.

40   Jason Stanley

Unger is aware of some of these points. As we have seen, he readily concedes that many of his "paradoxical" sounding sentences are in fact quite clearly intuitively acceptable. But Unger maintains that this is due to the fact that we normally allow language to be used loosely. According to Unger, in order to focus upon the actual meaning of a term, we need to place *focal stress* on that term. Whereas (19) is acceptable, a sentence such as (27) is not:

(27) He actually *knows* it's a Cadillac, but he's not absolutely *certain*.

Unger argues that focal stress is a way of drawing attention to literal meaning, and when we place focal stress on "know" and "certain", the oddity of knowledge without subjective certainty emerges. Unger thereby inadvertently draws attention to a *further* disanalogy between the relation between knowledge and truth (or belief), on the one hand, and knowledge and subjective certainty, on the other. We do not need the device of focal stress at all to detect the falsity of ascriptions of knowledge of clearly false propositions.

3.

Unger's reaction to the fact that we do not flinch when someone is described as knowing something of which she is not certain is to argue that in such cases, we are not "focused on actual meaning". As Unger writes (Ibid., p. 74), in response to a hypothetical objector dubious of Unger's "flatness" skepticism (according to which hardly any surface is actually flat):

> According to your account, he notes, 'flat' means the same as 'absolutely flat'. Thus, the simplest hypothesis for you is that sentences of the form 'x is flat' always are equivalent in meaning to ones of the corresponding form 'x is absolutely flat', at least when 'x' is not meant to pick out some old beer, etc. This creates something of a problem for you: When people say things like 'That is flat; the other is flatter', your account would have them saying something which is *inconsistent*, something which must always be false. But, the sentence doesn't *sound* inconsistent. And, indeed, it really does seem that sentences of this sort are often used to say things which are true. How are to you account for such blatant discrepancy?

Unger's response to the hypothetical objector is to argue that in cases in which we accept assertions such as 'That is flat, the other is flatter', it is because we allow terms to be used in loose and distorted ways, ways that are not reflective of "actual meaning". Similarly, when we allow someone who is uncertain of p to be described as knowing that p, we are allowing the term "know" to be used loosely.

Unger's own meaning hypotheses are principally motivated by appeal to our intuitions about various cases. It is therefore incumbent upon Unger to provide tests to distinguish intuitions that reflect "actual" meaning from intuitions that are corrupted by our tendency to allow speakers to use language loosely. Unger does so by providing methods of improving what he calls our "Focus on Actual Meaning". These tests centrally involve *focal stress*. As Unger writes (1975, p. 76):

> Emphasis does not change the meaning of words or sentences to which it is applied. Thus the sentence 'He killed her' means the same as the sentence 'He *killed* her'. The first will express what is true if and only if the second does. . .In any case, the importance of emphasis does not derive from any effect on meaning or truth. Quite the contrary, its importance in our language derives from its having no such effect. Shouldn't we have some device(s) to attract attention to a term which work in a way that does not effect (sic) the term's meaning, or the meanings of sentences of which that term is a part?

For Unger, the function of focal stress in our language is to "focus on actual meaning".

Unger acknowledges that it is a fact that we do not flinch when attributing knowledge that p to those who are clearly not certain that p. His response is to argue that we are speaking loosely (not in accord with the literal meanings of the terms), and the fact that we are speaking loosely emerges when we place focal stress on the relevant expressions, as in examples like:

(28) He actually *knows* it's a Cadillac, but he's not absolutely *certain*.

The fact that (28) sounds contradictory is explained by the fact that the function of focal stress in natural language is to focus on actual meaning, and the actual meaning of "know" involves absolute certainty.

As Unger recognized, his claim that stress "does not effect (sic) the term's meaning, or the meanings of sentences of which that term is a part" had been seriously threatened in work done by Fred Dretske (1972), who argued that stressing a term does affect the meaning of sentences of which that term is a part. Dretske emphasized the intuitive difference in meaning between (29) and (30):

(29) The reason Clyde *married* Bertha was to qualify for the inheritance.
(30) The reason Clyde married *Bertha* was to qualify for the inheritance.

Suppose that Clyde must marry within a year to qualify for his inheritance. He has no wish to give up his bachelor lifestyle, so he decides to marry Bertha, his lesbian friend, in order to meet this requirement. Intuitively, (29) is true

relative to this situation, and (30) is false. Conversely, suppose that Clyde is dating Bertha and Hannah. Clyde decides he wants to marry one of them, and has a preference for Hannah. But he will not qualify for his inheritance if he marries a Jew, and so he decides to marry the Christian Bertha instead. Relative to this situation, (30) is true and (29) is false. Dretske's work suggests that the purpose of focal stress is not to "focus on actual meaning", but rather to interact with salient sets of contrasts to the focal element in the sentence. Operators such as "the reason" are sensitive to this interaction.

Unger (Ibid., pp. 78ff.) provides a *pragmatic* account of the intuitions concerning (29) and (30). By stressing "married" in (29), the speaker is drawing our attention to one aspect of the event being explained— namely, the marrying part of the event (similarly, by stressing "Bertha" in (30), the speaker is drawing our attention to another aspect of the event being explained—namely that its patient was Bertha). Qualifying for the inheritance is the sole explanation for why the event of Clyde marrying Bertha was a *marrying* event. But it is not the sole reason for the event of Clyde marrying Bertha. Therefore, according to Unger's pragmatic account of (29), though the speaker's utterance is felicitous, what she asserts is nevertheless false, since qualifying for the inheritance was not *the* reason Clyde married Bertha.

Unger's pragmatic account of Dretske's data does not do justice to the generality of the phenomenon. Consider the well-known contrast between (31) and (32):

(31) Bill only introduced *John* to Frank.
(32) Bill only introduced John to *Frank*.

An utterance of (31) is true if and only if Bill introduced John to Frank, and Bill did not introduce anyone other than John to Frank. In contrast, an utterance of (32) is true if and only if Bill introduced John to Frank, and Bill did not introduce John to anyone other than Frank. It is unclear how to generalize Unger's pragmatic explanation of the felicity of (29) and (30) to (31) and (32). Just as "the reason that" interacts with the focused expression in (29) and (30), "only" interacts with the focused expression in (31) and (32), and the interaction results in a difference in truth conditions.[2]

Unger's claim that placing focal stress on a term does not affect the truth-conditions of larger sentences containing it is false.[3] But more germane to our concerns is Unger's contention that placing focal stress on a term does not alter its actual meaning.[4] According to Unger, pronouncing sentences with flat intonation results in a content communicated that is thoroughly affected by pragmatic processes, and the way to eliminate the effects of these pragmatic processes is to appeal to focal stress. This is a strange view, from a contemporary perspective. We certainly do not yet understand all the ways in which focus interacts with context. But one thing all linguists agree

upon is that focusing an expression introduces a myriad variety of pragmatic effects, depending upon the type of intonation, effects which are eliminated if one pronounces the relevant sentence with flat intonation. Though matters are never straightforward where focus is concerned, if anything it is flat intonation which allows us to "focus on actual meaning".

Unger's general view of focus is incorrect. But he is right that it is odd to say, of someone who has a good deal of evidence that Bush is president in 2006, but is somewhat unconfident:

(33) He just *knows* that Bush is president.

Similarly, sentences such as (34) also seem quite odd:

(34) He really *knows* that Bush is president, though he's not completely *certain*.

It is one thing to dismiss Unger's theoretical justification for taking such utterances at face-value; it is quite another to explain their oddity without impugning the doctrine that knowledge does not require certainty.

I do not have an explanation for what is going on with stressed uses of "know". But consideration of a wider range of cases might make us suspicious of allowing such uses to guide our hypotheses about literal meaning. For example, suppose, reading a newspaper on the day after the election, I discover that Bush in fact won the election. It would be quite ordinary for me to respond:

(35) How could this have happened? I *knew* that Kerry was going to win.

Or suppose, having found out that my friend did not after all cheat me:

(36) Boy, I just *knew* you had cheated me. I'm so relieved to find out I was wrong.

One would surely not want to take these uses of "know" as showing that the actual meaning of "know" did not require truth. But even if we did take stressed uses of "know" to be indicative of actual meaning, Unger's thesis that knowledge is incompatible with uncertainty would be no more plausible. There are more direct arguments against his dual contentions that focal stress brings out actual meaning, and that this undermines the thesis that knowledge is consistent with subjective uncertainty. Here is an example from a different context.[5] Suppose that I am watching my favorite college basketball team on what I know to be tape-delay on television, and they are

44   Jason Stanley

down by 15 points with 10 minutes to go. No one has told me the final result. But, knowing my team, I felicitously say:

> (37) I *know* they are going to lose. But I guess there is a *chance* they'll pull it out in the end, so I'll continue watching.

Unger thinks that focal stress is a way to focus on actual meaning. So he is committed to taking the stressed use of "know" at face-value. Surely, the epistemic possibility of ∼p is incompatible with knowing that p. So Unger needs to take the stressed use of "chance" to be a way of talking of subjective credence. So, by Unger's lights the apparent truth of (37) shows that knowledge does not require certainty after all.

4.

Though the case for the thesis that knowledge requires certainty is extraordinarily weak, Unger is correct that instances of (38) are invariably odd, as are instances of (39):

> (38) I know that p, but I'm not certain that p.
> (39) I know that p, but it's not certain that p.

This is the last vestige of the argument that knowledge requires subjective and epistemic certainty. We have seen that it cannot support the conclusion that knowledge requires certainty. So there must be some other explanation of the oddity of utterances of instances of (38) and (39).

In the case of the relation between knowledge and truth, there is no asymmetry between first person claims and third person claims. In other words, instances of (40) and (41) are clearly odd:

> (40) I know that p, though p is false.
> (41) She knows that p, though p is false.

The oddity of instances of both (40) and (41) seems to emerge from the same source, namely obvious falsity. In contrast, as we have seen, instances of (42) are not anywhere nearly as odd as instances of (38) (imagine an instance of (42) used to describe a cautious yet eminently renowned expert on the subject matter in question):

> (42) She knows that p, though she's only somewhat certain of it.

If knowledge entailed subjective certainty, as knowledge entails truth, then the existence of the asymmetry between first person and third person cases would be utterly mysterious. So there is some other explanation of the

oddity of instances of (38). We have also seen that there is a clear contrast between quantified versions of (40) and (41), and quantified cases involving knowledge and subjective and epistemic certainty, such as (13) and (16). This pattern of intuitions completely vitiates an explanation that appeals to entailments between knowledge and subjective or epistemic certainty.

A clue to the correct explanation of Unger's evidence comes from the fact that sentences such as (43) and (44) are just as odd as instances of (38) and (39):

(43) Dogs bark, but I'm not certain that they do.
(44) Dogs bark, but it's not certain that they do.

(43) and (44) are instances of "Moore's Paradoxical" utterances, such as:

(45) Dogs bark, but I don't know that they do.
(46) Dogs bark, but I don't believe that they do.

Assertions of (43)-(46) are odd, and their oddity presumably is due to some fact about conversational norms. The fact that (38) and (39) are just as odd as cases of Moore's paradox with "certain" suggests the desirability of a uniform explanation.

Unger (1975, p. 259) quite clearly recognizes the fact that the oddity of an assertion of (38) is due to the same source as the oddity of an assertion of (43):

> We may explain the apparent inconsistency of 'It's raining, but I'm not absolutely sure that it is' by making an assumption about knowing and being absolutely sure. We need only suppose that one's knowing something to be so entails one's being absolutely sure of it.

Unger (Ibid., p. 260) proposes the following account of the fact that an utterance of (43) is odd:

> We may think, for example, that the person just represented himself as knowing that p, and not as being absolutely certain that p. Understanding this representation, we in turn quickly realize that what he represented entails that he is absolutely certain that p, which he then goes on to deny.

This account proposed by Unger appeals to the knowledge account of assertion, which is:

The Knowledge Norm for Assertion: Assert that p only if you know that p.

According to the account, the knowledge norm for assertion, together with the principle that knowledge entails subjective certainty, explains the fact that

utterances of (43) are odd; an analogous entailment in the case of epistemic
certainty would explain the oddity of utterances of (44). But the principle
that knowledge entails certainty (in either sense) is false. So this account is
unhelpful in explaining the oddity of (43) and (44).

The fact that knowledge does not entail certainty shows that we cannot
derive an explanation of the oddity of (43) by appeal to the Knowledge
Norm for Assertion. But perhaps we can reverse the strategy, and account
for the lure of the Knowledge Norm by appeal to The Certainty Norm for
Assertion:[6]

The Certainty Norm for Assertion: Assert that p only if you are certain
that p.

The Certainty Norm for Assertion explains the oddity of (43). If there were
a certainty norm for assertion, one could both explain away the intuitions
behind the claim that knowledge entails certainty, as well as the existence
of Moore's paradoxical utterances involving "certain". However, since there
is good reason to reject an entailment from knowledge to certainty, one
cannot derive this norm, as Unger suggests, from the knowledge account of
assertion.

Utterances of both (43) and (44) are odd. This suggests that accounting
for all of the evidence requires two certainty norms of assertion:

The Subjective Certainty Norm for Assertion: Assert that p only if you
are subjectively certain that p.

The Epistemic Certainty Norm for Assertion: Assert that p only if you
are epistemically certain that p.

The existence of these norms for assertion straightforwardly explains the
oddity of (43) and (44), in the same way in which the knowledge account of
assertion explains the oddity of (45). It remains to be seen how these norms
of assertion can be brought to bear to explain the oddity of all versions of
Moore's paradox.

Subjective certainty is a context-sensitive matter, and its relation to full
belief is vexed. Nevertheless, I take it that whatever the level of subjective
certainty is in a context, it is at least as strong as the level of confidence
required for full belief. Since subjective certainty entails full belief, the
subjective certainty norm for assertion can explain the oddity of instances of
Moore's Paradox involving "believes", such as (46), in exactly the same way
as the proponent of the knowledge account of assertion explains them (e.g.
Williamson (2000, p. 254)).

The issue is slightly more complex in the case of instances of Moore's
Paradox involving "know", such as (45). Consider the proposition that there

are no large Jewish elephants in my bedroom. This may have been an epistemic certainty for me five minutes ago, even though I did not know that there were no large Jewish elephants in my bedroom. I did not know that there were no large Jewish elephants in my bedroom, because I did not believe it, and I did not believe it simply because it didn't occur to me ever to entertain that possibility. Nevertheless, in this case, if I had entertained the proposition that there are no large Jewish elephants in my bedroom, I would have known it. The reason this counterfactual is true is because it is an epistemic certainty for me that there are no large Jewish elephants in my bedroom. So the fact that a proposition is an epistemic certainty for a person does not entail that the person knows that proposition. If a proposition is an epistemic certainty for a person at a time, then it does follow that the person is in a *position to know* that proposition. Being in a position to know a proposition is to be disposed to acquire the knowledge that that proposition is true, when one entertains it on the right evidential basis. Since epistemic certainty entails possession of this dispositional property, utterances of (45) are odd.[7]

So, all versions of Moore's paradox (with belief, knowledge, and certainty) would be explained by the invocation of the dual certainty norms for assertion. In contrast, the knowledge account of assertion, according to which assertion is governed by a norm for knowledge, can only explain the belief versions of Moore's paradox, such as (45). In order to account for the certainty versions of Moore's paradox, such as (43) and (44), the advocate of the knowledge account of assertion must embrace a *contextualist* account of knowledge ascriptions, according to which we have "a reluctance to allow the contextually set standards for knowledge and certainty to diverge" (Williamson, 2000, p. 254). But this is problematic. First, contextualism about knowledge ascriptions is a problematic view about the semantics of knowledge ascriptions. Secondly, it forces the advocate of the knowledge account of assertion to adopt the position that a knowledge ascription can only express a truth relative to a context of use if the corresponding certainty ascription expresses a truth in that context. But, as we have seen, neither the view that knowledge entails certainty, nor the view that "knowledge" entails "certainty", is plausible.

The certainty norms for assertion also explain the patterns of symmetry and asymmetry between first and third person reports. As Unger realized, a third-person report of knowledge despite lack of subjective certainty is perfectly in order:

(19) John knows that Bush is a Republican, though, being a cautious fellow, he's somewhat uncertain of it.

The fact that (19) is fine, but instances of (38) are not, is straightforwardly explained by the fact that we adhere to a norm of subjective certainty for

48   Jason Stanley

assertion. These norms also explain why there is no such asymmetry in the case of epistemic certainty:

(20) I know that Bush is president, though it's not certain that he is.
(21) John knows that Bush is president, though it's not certain that he is.

An utterance of (21) is just as odd as an utterance of (20), because the use of "certain" is linked to the knowledge base of the person making the assertion; an utterance of (21) expresses the proposition that John knows that Bush is president, though, given the knowledge base of the person making the assertion, it's not certain that Bush is president. The oddity of this assertion is explained by the epistemic certainty norm for assertion, together with the fact that knowledge is factive. In order for Hannah to assert that John knows that Bush is president, she must be epistemically certain of it; but then, by the factivity of knowledge, she must also be epistemically certain that Bush is president. So the oddity of Hannah's utterance of (21) shows that Hannah must be epistemically certain that John is president in order to assert it. But it does not show that *John* must be epistemically certain that Bush is president in order to know that Bush is president. So there is no route from the oddity of utterances of (21) to a constitutive connection between knowledge and epistemic certainty.

Finally, there are non-Moorean considerations in favor of the certainty norms for assertion. One natural way to challenge an assertion is, as advocates of the knowledge norm of assertion have emphasized (Williamson (2000, pp. 252-3)), to utter "How do you know?" But it is equally natural, when confronted by someone making an assertion on inadequate grounds, to respond with "Are you sure?" Without appeal to contextualism about knowledge ascriptions, the knowledge norm of assertion can only explain the former sort of challenge, but not the latter. If however the norms for assertion involve in the first instance certainty, rather than knowledge, the naturalness of both of these kinds of challenges is explicable.

The certainty account of assertion can explain all the data that motivates the knowledge account of assertion, but not vice-versa. Furthermore, attempts to explain the data that motivate the knowledge account of assertion with weaker norms of assertion, such as the truth norm (Weiner, 2005) or the reasonable-to-believe norm of assertion (Lackey, forthcoming) do not generalize to account for the kind of data that motivates the certainty norms for assertion. For example, according to Weiner (2005), the truth norm of assertion can account for Moore's paradoxical utterances involving "know". According to him, this is because someone who asserts that p while disclaiming knowledge that p is construed as not possessing warrant to follow the truth rule. However this explanation may fare in explaining Moore's paradoxical utterances involving "know", it does not generalize to

explain Moore's paradoxical utterances involving "certain", such as (43) and (44), nor does it explain why we often demand certainty, and not merely knowledge, of our interlocutors.[8]

5.

Perhaps, like the knowledge account of assertion, there is only one certainty norm for assertion, and the other is derivative. For example, perhaps assertion primarily requires *subjective* certainty. One could explain the work done by the norm of epistemic certainty by appealing to norms governing subjective certainty. One such norm is:

The Epistemic Certainty Norm for Subjective Certainty: Be subjectively certain that p only if you are epistemically certain that p.

On the assumption that epistemic certainty is a norm for subjective certainty, it may be possible to reduce the twin norms of assertion to a single norm of subjective certainty. Here is how the defender of the single subjective certainty norm for assertion would explain the oddity of Moorean utterances such as (45). First, an utterance of "Dogs bark" would implicate that the speaker is subjectively certain that dogs bark. If the speaker is adhering to the norms of subjective certainty, then she is also epistemically certain that dogs bark, and hence knows that dogs bark. The oddity of (45) is due to the fact that asserting "dogs bark" implicates that the speaker knows that dogs bark, via the subjective certainty norm for assertion, together with the epistemic certainty norm for subjective certainty.

One worry for this approach is that the fact that epistemic certainty is a norm for subjective certainty might not 'transfer' to explaining the Moorean oddity of utterances of (45). For example, it is plausible that knowledge is the norm of belief. Nevertheless, utterances of (47) are perfectly in order:

(47) I believe that dogs bark, but I don't know it.

The fact that utterances of (47) are in order suggests that the fact that knowledge is a norm of belief does not mean that an utterance of an instance of "I believe that p" implies that the subject knows that p. Similarly, even if epistemic certainty is a norm for subjective certainty, it would not follow that presumed adherence to a norm of subjective certainty would implicate that the agent is epistemically certain of the relevant proposition. However, the function of using "I believe" in (47) is to qualify support for the truth of a proposition, rather than endorse it. In short, such uses of "believe" are not cases in which one reports a belief that p at all; they are rather cases in

50   Jason Stanley

which one reports that one has weak reasons in support of the truth of a proposition.[9]

An alternative way to pursue the thought that there is just one certainty norm for assertion is to defend the view that assertion primarily requires *epistemic* certainty. The requirement of subjective certainty would be a derivative norm. Subjective certainty, unlike epistemic certainty, is under the rational control of an agent. Rational agents who seek to adhere to the norm of epistemic certainty would manifest their adherence by only asserting propositions of which they were subjectively certain. Instead of governing the *act* of assertion, the norm of subjective certainty would emerge from rational requirements on an agent's *adherence* to the norm of epistemic certainty.[10]

In what follows, I will assume for the sake of discussion that the fundamental constitutive norm governing the act of assertion is the epistemic certainty norm for assertion, and that the subjective certainty norm has a derivative status.

The familiar rivals to the knowledge account of assertion are *less* demanding norms for assertion, such as the truth norm, the belief norm, and the warranted belief norm. In contrast, the epistemic certainty norm for assertion is an even *more* demanding norm than the knowledge norm. But the argument for the more demanding norm for assertion is of the same form as the argument for the knowledge norm for assertion. The certainty norm for assertion explains all versions of Moore's Paradox, and does so without commitment to contextualism. Furthermore, there are additional reasons for the more demanding norm, which stem from the nature of testimony.

It is an intuitive thought about testimony that someone who acquires a belief via testimony has less justification for it than the person who transmitted the belief. But this thought is in tension with the thesis that communication is a reliable means for transmitting knowledge. If A knows that p, and informs B, and neither A nor B has any defeating reasons for their belief that p, then B thereby comes to know that p, even if A is her only source. But then if the person who asserts that p only just satisfies the conditions for knowing that p, then the belief formed on the basis of her assertion will not amount to knowledge. For communication to be a guaranteed method of transmitting knowledge between agents with no defeating reasons, the assertor must be in a *better* epistemic position with regard to p than just knowing that p. In short, the epistemic certainty norm for assertion is one way to explain the consistency between the thesis that assertion is a method of transmitting knowledge, and the thought that testimony does not preserve strength of evidential position towards a belief.

If knowledge is the norm for full belief, then the epistemic certainty norm for assertion places a greater demand upon assertion than the norms governing full belief. This is problematic, if occurrent belief is thought of as a kind of inner assertion. But this model of occurrent belief ignores the public character of the practice of assertion. When someone asserts something, she

takes on the commitment of transmitting her knowledge of it to her interlocutor. On the assumption that acquiring a belief via testimony results in having a weaker epistemic position towards the content of that belief than the one from whom one has acquired the warrant to believe it, it makes perfect sense that assertion would have a more demanding norm than full belief.

One might worry that the twin norms of certainty are *impossibly* demanding. For example, if subjective certainty is construed as credence 1, then being subjectively certain of a proposition requires that one would bet on it, no matter what the odds. Similarly, if epistemic certainty is construed as knowing a proposition on non-circular logically entailing grounds from a priori premises, then epistemic certainty requires Cartesian certainty. Thus construed, the claim that there are certainty norms governing assertion is impossibly demanding.

The main point of my paper has been to argue that knowledge does not entail certainty of any kind. If one holds that certainty is not achievable, then if the certainty norm is what explains instances of Moore's Paradox, it follows that no one can satisfactorily meet the norms governing assertion. Skepticism about certainty undermines warranted assertion, but since knowledge does not entail certainty, skepticism about certainty does not entail skepticism about knowledge. On the picture I have defended, skepticism about certainty would suggest that we mistakenly demand overly restrictive conditions on asserting, ones that we in fact never meet.

But skepticism about certainty should be rejected. Unger's skeptical argument that we are never certain of anything depended crucially upon his "method of focusing on actual meaning". Unger's skeptical argument has two premises: first, that knowledge requires certainty, and secondly, that we are certain of almost nothing. We have already seen that the failure of this method undermines Unger's case for the first premise. But it also undermines Unger's case for the second premise.

According to Unger, the adjective "certain" is what he calls an *absolute term*. An absolute term, by Unger's lights, is one that only holds of things at the end of the relevant scale. According to Unger, "flat" is an absolute term; it holds only of the flattest thing in the universe. Unger's argument that "flat" is an absolute term is that by "focusing on actual meaning" we recognize that if x is flatter than y, then y isn't flat at all. Unger also argues that "certain" is an absolute term. There is good reason to believe that the class of adjectives Unger picks out as absolute form a linguistically distinctive class.

The fact that absolute adjectives form a distinctive class does not save Unger's skeptical argument that absolute adjectives are never in fact true of anything. As we have seen, Unger is utterly clear that an unstressed use of (48) is perfectly in order, as is an unstressed use of (49):

(48) That is flat; that other is flatter.
(49) Though John is certain that Bush is president, he's even more certain that he's not going to vote Republican anymore.

52 Jason Stanley

The way that Unger argues that (48) and (49) actually are false is by appeal to his method of focusing on actual meaning. The fact that this is not a way to identify actual meaning undermines Unger's skeptical argument for regarding absolute adjectives. The fact that "flat", "straight", and "certain" are associated with endpoints in the way that they are in no way entails that it is not the case that plenty of things satisfy "flat", "straight", and "certain".

Here is a brief sketch of a context-sensitive semantics for "certain". On either of its interpretations, "certain" accepts modifiers such as "absolutely" or "completely". So "certain", in both its subjective and epistemic uses, is an absolute gradable adjective. "Certain" is context-sensitive in two ways. First, it is context-sensitive as regards the *kind* of scale relevant to its interpretation (a scale of confidence or a scale consisting of degrees of justification). Secondly, it is context-sensitive as regards the degree on the scale that is required to satisfy the property it expresses, relative to that context. A person's belief satisfies the property expressed by a subjective use of "certain" relative to a context if and only if that person holds that belief at or above the contextually salient degree of confidence; *mutatis mutandis* for epistemic certainty and degrees of justification.

Just as many beliefs may satisfy "certain", many beliefs may, in context, satisfy "absolutely certain". The semantic function of "absolutely" is to raise the degree on the scale above that for "certain". So in any context, it will be harder for a belief or a proposition to satisfy "absolutely certain" than it will be for it to satisfy "certain". But it still is possible for a belief to satisfy "absolutely certain" in one context, while not satisfying "absolutely certain" (or even "certain") in another.

Assuming this contextualist semantics for "certain", the norms governing assertion are naturally to be understood meta-linguistically; assert that p in a context only if your confidence in p meets the contextually salient standard for "certain" at that context, as well as the contextually salient standard of justification that is required for epistemic certainty. When I assert that the store is open, I am adhering to the norm of subjective certainty for assertion provided only that I would also be licensed to declare in that context that I am sure that the store is open, i.e. that I meet the contextually salient standard for "certain". Construed this way, the notions of certainty relevant for the norms of assertion are nowhere near as demanding as willingness to bet on a proposition no matter what the odds, or having Cartesian grounds for its truth.

The fact that I am sympathetic to the thesis that "certain" is context-sensitive does not mean I endorse preserving the view that knowledge requires certainty by appeal to semantic ascent. As Unger himself noted, it is perfectly in order to describe someone whose belief that p does not meet the contextually salient standard of certainty as knowing that p. Knowing requires neither certainty nor "certainty".

Conclusion

I have defended fallibilism via the familiar strategy of arguing that there is a pragmatic explanation for the data that suggests that knowing that p entails being subjectively or epistemically certain that p. The problem with previous attempts is that their proponents have never provided a clear description of the nature of the pragmatic explanation, or a decisive refutation of the explanation in terms of an entailment between knowledge and certainty. The appeal to pragmatics has therefore seemed like an ad hoc attempt to explain away obvious entailments.

I have argued that the attraction of the entailment thesis is due to an insufficient diet of examples; when we cast our nets more broadly, we see it cannot be correct. Furthermore, the intuitions that have been marshaled in favor of the entailment thesis uniformly have the character of instances of Moore's Paradox, which are typically used as evidence in favor of various norms of assertion. Following this strategy leads directly to the adoption of a certainty norm for assertion. On the supposition that there is a certainty norm (or norms) for assertion, we can explain all instances of Moore's Paradox, including the evidence in favor of an entailment between knowledge and certainty. So not only is the entailment thesis clearly incorrect, but we have in place a compelling pragmatic explanation of the data seemingly supporting it, an explanation of a character familiar from both classical and recent discussions of Moore's Paradox.

The view of the relation between knowledge, certainty, and assertion that I have defended here has a number of consequences. For example, it undermines a recent argument by Keith DeRose (2002) for contextualism about knowledge attributions. According to DeRose, the standards for assertion vary from context to context. If the knowledge account of assertion is correct, then contextualism about knowledge attributions seems to follow. As DeRose writes:

> What of the advocate of the knowledge account of assertion who does *not* accept contextualism? Such a character is in serious trouble. Given invariantism about knowledge, the knowledge account of assertion is an untenable attempt to rest a madly swaying distinction upon a stubbornly fixed foundation. Less metaphorically, it is an attempt to identify what is obviously a context-variable standard (the standard for the warranted assertion of "P") with what one claims is a context-invariable standard (the relevant truth-condition of "S knows that P," according to the invariantist). The knowledge account of assertion demands a contextualist account of knowledge and is simply incredible without it.

Given what I have argued, one can agree with DeRose that there are varying contextual standards for assertion, while rejecting contextualism

54   Jason Stanley

about knowledge attributions. Since the norms for assertion involve *certainty*, if certainty is also a context-dependent matter, the fact that there are varying contextual standards for assertion is consistent with invariantism about knowledge.

Fallibilism in epistemology is often formulated as the doctrine that knowledge is compatible with lack of certainty. It is a widespread belief among philosophers that the folk concept of knowledge has an infallibilist character, and fallibilism is the doctrine to which theorists of knowledge are forced to retreat when confronted with its implications. One charge against fallibilism has been that it seems to entail that I can know that p, despite it being possible that ~p. The other charge is that it seems to entail that I can know that p, despite being less than certain that p. In previous work (XXXX), I have argued that fallibilism does not entail that I can know that p, despite it being possible that ~p. In this paper, I have argued that while fallibilism does entail that I can know that p despite being less than certain that p, it follows from independent facts about norms for assertion that I cannot *say* that I know that p and am less than certain that p. Therefore, no revision in our ordinary conceptual scheme is required to embrace fallibilism. Our folk concept of knowledge gives no succor to the skeptic.

Notes

1. I argue below that asserting "I know that Bill came to the party" pragmatically imparts that the utterer is certain that Bill came to the party, and that it is epistemically certain for the utterer that Bill came to the party. Thus, (24) and (25) are similar to cases in which an implicature is *reinforced*, as in the example "John has two children. In fact, he has exactly two children.", or "I ate some of the cake, but I didn't eat all of it." Pragmatically imparted information can be reinforced; entailments cannot (Sadock, 1978). Thanks to John Bengson here.
2. There are of course many other cases in which focus affects truth-conditions. Consider, for example, the distinction in truth-conditions between "*Hurricanes* arise in this part of the Atlantic" and "Hurricanes arise in *this* part of the Atlantic". The former is true if and only if at least some hurricanes arise in the demonstrated portion of the Atlantic. The latter is true if and only if it is a generic property of hurricanes that they arise in the demonstrated part of the Atlantic.
3. Unger has a second argument against Dretske, involving indirect speech reports (p. 79), which is undermined once one recognizes that focus is a context-sensitive semantic phenomenon.
4. There is a good deal of intuitive evidence against this. Suppose I need to cut through something thick, and I ask my friend for a knife, and he brings me a butter-knife. I might respond by saying "I need a *knife*, not this thing."
5. I recall learning about an example similar to this one from Brian Weatherson's weblog (http://tar.weatherson.org).
6. This strategy is also one Unger (Ibid.) contemplates.

7. Alternatively, the right reaction to these cases may be to subsume the apparent counterexamples to the thesis that epistemic certainty entails knowledge to cases of dispositional belief.

8. It bears mention that the certainty norm of assertion is in tension with other commitments I have incurred in my work. For example, in XXXX, one of the claims we defend is that under certain conditions, knowing that p is sufficient for acting on p.

9. Thanks to John Hawthorne for discussion. In unpublished work, Michael Huemer has persuasively argued for a similar conclusion with respect to examples such as (50).

10. Thanks to Michael Martin and Ian Rumfitt for discussion here.

References

DeRose, Keith (2002): "Assertion, Knowledge, and Context," *Philosophical Review* **111**: 167–203.

DeRose, Keith (2004): "Single Scoreboard Semantics," *Philosophical Studies* **119**: 1–21.

Dretske, Fred (1972): "Contrastive Statements", *Philosophical Review*, **81.4**: 411–437.

Klein, Peter (1981): *Certainty: A Refutation of Skepticism* (Minnesota: University of Minnesota Press).

Lackey, Jennifer (forthcoming): "Norms of Assertion", *Nous*.

Lewis, David. 1996. "Elusive knowledge" *Australasian Journal of Philosophy* **74**: 549–567.

Miller, Richard (1978): "Absolute Certainty", Mind, **87**: 46–65.

Pryor, James (2000): "The Skeptic and the Dogmatist", *Noûs* **34** (2000), 517–49.

Sadock, Jerry (1978): "On Testing for Conversational Implicature". In *Syntax and Semantics* **9**: *Pragmatics*, edited by Peter Cole: 281–297.

Unger, Peter (1975): *Ignorance* (Oxford: Oxford University Press).

Weiner, Matthew (2005): "Must We Know What We Say?" *Philosophical Review*, April 2005(vol. 114 no. 2).

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