What is Evidence in Natural Language?*

Eric McCready

September 8, 2011

+++DRAFT +++ COMMENTS WELCOME+++

Abstract

This paper tries to understand the proper notion of evidence to use in the semantic analysis of natural language evidentials. I review various notions of justification from the epistemological literature, and consider how they relate to the use of evidentials and related constructions. I then consider how (some) evidentials behave under Gettier scenarios. The conclusion is that the required notion of evidence is one which is weaker than (many accounts of) knowledge, involves increase of speaker credence, but which is necessarily first-person. I thus settle on a view based on a self-ascription of probability increase due to knowledge of propositions that increase credence after conditionalization.

Keywords: evidentials, evidence, probabilities, justification, epistemology, semantics, indexicality

1 The problem

Evidentials have been widely studied in linguistics for many years. They can be defined roughly as expressions which indicate the basis of the claim made by a speaker: in short, the source of the evidence on which the speaker's claim is based. In early descriptive work the existence of evidentials was clearly acknowledged, though not given much concentrated attention (e.g. Sapir 1922). Subsequently, researchers have addressed themselves to complex questions about the nature of evidentials and how to define the category: for example, much attention has been given to the distinction between evidentials and miratives (e.g. DeLancey 1997) and between evidentials and epistemic modals (de Haan, 1999; Matthewson et al., 2007). Most recently of all, scholars working in formal semantics and pragmatics have worked to place evidentials into the mathematical analysis of linguistic meaning; after the seminal work of Garrett (2001) and Faller (2002), a rather large literature has developed (e.g. Ogata 2005; Chung 2005; McCready and Asher 2006; McCready

*I owe thanks to very many people over the course of this project. First is Norry Ogata, with whom I started working on this topic. It is too bad we could not have written this paper together. I would also like to thank . . . and audiences at LENLS, GLOW, . . .

and Ogata 2007; Matthewson et al. 2007; Faller 2007; Davis et al. 2007; Murray 2010; Bronnikov 2011, among many others). From this work, a great deal has been learned about evidentials: for instance, that evidentials in different languages can show different properties with respect to scope and 'level of application', some evidentials act much like (epistemic) modals, while others appear more like elements which alter the speech acts performed with the sentences that include them, and that different evidentials show distinct preferences for what kinds of content they can apply to.

However, some issues remain open. One involves what hearers ought to do with evidentially marked content after processing it. Should it be updated with in the usual way? Or should it be treated more like content modified by e.g. modal operators? The vast literature on dynamic semantics tells us that the two options yield very different results (cf. Veltman 1996). Some authors have taken explicit positions on this issue; for instance, McCready and Ogata (2007) claim that Japanese evidentials behave like modals, and so update with evidentially marked content should act as global tests on hearer information states, but simultaneously introduce new discourse referents for information sources. On the other hand, Davis et al. (2007) take evidentials to modify a contextual standard for assertability, according to which update on the evidentially marked content proceeds as usual. Part of the reason for this difference in focus presumably involves the different sorts of evidential systems examined by the two papers. But the general problem remains. In some cases, it can be shown to lead to other complex issues: McCready (2011b) considers the general situation for update with hearsay evidentials in light of widely accepted pragmatic principles, showing that it is not at all simple, regardless of the evidential system in question. A more general treatment of this issue is needed.

This paper, however, focuses on another, still more pressing problem. What is evidence, actually? Put in another way, what is the content introduced, or manipulated, by evidentials? A great deal of the formal literature on evidentiality takes the notion of evidence as a pure primitive. We often find definitions like the following one (slightly modified; emphasis mine):[1]

(1) $[\![k'a(f)(B)(w)(\varphi)]\!]$ is only defined if for all worlds $w, w' \in B(w)$ **the inferential evidence in** w **holds in** w', and f is a choice function of type $\langle st, st\rangle$ such that $f(B(w)) \subseteq B(w)$.
If defined, $[\![k'a(f)(B)(w)(\varphi)]\!] = 1$ iff $\forall w' \in f(B(w)) : [\![\varphi]\!](w') = 1$.
(Matthewson et al., 2007)

Here, it is assumed that we can identify whatever evidence is being presupposed to exist. But this is not a trivial task. Compare a standard case of presupposition like (2): here there are no problems finding out whether the presupposition is or is not satisfied. We merely check for the existence of a King of France.

(2) The King of France enjoys oatmeal with blueberries.

Life is not so easy for (1). Without a definition of evidence, how can we be sure that we have found the piece the speaker had in mind, or a substantial enough piece of evidence, or indeed any evidence at all? It seems to me that we can't. This, in my view, constitutes a serious problem for (most) current theories of evidentiality.

[1] The use of the choice function here is meant to account for quantificational variability in the St'at'imcets evidentials. See Matthewson et al. 2007 for details.

The issue becomes even more worrying when one considers accounts of evidentiality like that of Chung (2005), for whom evidence is (as I read things) something to be found in the external physical environment. She defines a notion *v-trace* along the lines of the 'temporal trace function' of Krifka (1992; further modified by Faller 2004 to a function on spatiotemporal traces for use in the analysis of a particular kind of past tense marker). Such a function is sensible: we know what it is for an event or an individual to occupy (space)time, and we know how to identify the spatiotemporal region that event or individual has occupied. But consider the *v-trace* function, defined as (3) in the version of Chung (2007).

(3) $v\text{-}trace(e) = \{\langle t, l\rangle | \exists v[EVIDENCE - FOR(v, e) \wedge AT(v, t, l)]\}$, where $AT(v, t, l)$ is true iff the evidence v for the occurrence of the eventuality e appears at a location l at time t.

So the 'evidential trace' of an event corresponds to the spatiotemporal coordinates of all evidence for that event. Clearly, it is difficult or impossible to identify such evidence without knowing what evidence is meant to be. The essential difference between the cases is that, unlike the King of France, the notion of evidence is not epistemologically innocent. There are many assumptions lurking in the background, and, without clarifying those assumptions, it does not seem to me that we can claim to have an actual theory.

Further, I think that leaving crucial terms undefined leads to strange predictions in our theories. Considering (3) again, what happens in the case of inferential evidence? Suppose that I see you at home, and I know that you ordinarily stay home if it is raining, and otherwise don't. Together this is enough to infer that it must be raining now. But what is the spatiotemporal trace of this evidence? It seems odd to say that it is the same as the trace of your being at home, for my background knowledge also comes into play. But where is the spatiotemporal trace of this knowledge? The obvious answers are highly counterintuitive, and I doubt that anyone would really want to accept them; further, we are already into dubious philosophical territory with the assumption that evidence corresponds to *events*. While this might make sense for many kinds of evidence, such as the observable fact that you are at home, an existentially quantified statement about an event, it makes much less sense for propositions such as the second premise above, that you ordinarily are at home only if it is raining: how can one find the spatiotemporal coordinates of a generic statement? Clearly, the lack of a proper definition of evidence is leading to some pernicious consequences at this point, which I believe can be avoided by being more explicit about what evidence is.

So far my argument has been limited mostly to what might be described as ontological qualms. One might respond: why don't we just leave it to the epistemologists? Why do linguists have to worry about this issue? The answer is: we don't, if we are willing to let the epistemologists define for us how our theories are supposed to work. More precisely, there is no obvious reason to assume that the notions of evidence relevant to work in epistemology or philosophy of science are identical to those we require for doing semantics. Why should natural language care about scientific evidence? There seems no a priori reason why it should, and perhaps many reasons why we should expect that a different kind of evidence is at issue. The goal of this paper is therefore to find the right notion of evidence for work on evidentiality. This task is of independent interest beyond the task of, as it were, picking up after the formal semanticists: I hope to show that we learn something substantial about the nature of evidentiality in natural language from the attempt.

The structure of the paper is as follows. In section 2 I will lay out some background relating to the main topic of the paper: first, I will summarize some existing positions in the linguistics literature about the nature of evidence (some implicit in the formal analyses that have been proposed), and then turn to some facts that have been used to argue for particular views of evidence in natural language (McCready, 2008, 2010). These facts primarily relate to two factors: the thesis that evidence must be knowledge (cf. Williamson 2000), and Gettier scenarios (Gettier, 1963), broadly construed as involving the distinction between subjective and objective understandings of what counts as evidence and knowledge. Section 3 discusses some distinctions made in epistemology relevant to the issues, along with some extra necessary conditions on an account of linguistic evidence. Section 4 takes the facts adduced in section 2.2 and considers various possible stories about linguistic evidence in their light. The conclusion is that natural language evidence is a concept which involves speaker beliefs about whether or not the truth of one proposition increases the probability of another in a necessarily *de se* manner. In section 5.1, several phenomena are exhibited which are of interest with respect to the proposal, involving externalist versus internalist concepts of evidence and true causality. I show that in each case some light is shed by the proposed framework. Section 6 concludes the paper.

2 Subjectivity, skepticism, and Gettier scenarios

In this section I examine accounts of evidence from the linguistic literature and show that they leave certain questions unanswered. I focus on the account of McCready and Ogata (2007), who give the most explicit story of which I am aware; looking at other available analyses, I show that even the most (formally) explicit have as their main goal explaining the assertability conditions of evidential sentences, taking the notion of evidence essentially for granted. This is done in 2.1. Section 2.2 discusses some data that bear on the kind of evidence used in evidentials having to do with the use of evidential sentences in skeptical and Gettier scenarios.

2.1 Implicit accounts of evidence

What does evidence do? The obvious answer is that it provides justification for certain beliefs. One way to think about this justification is by means of changes in the probabilities assigned to the content of those beliefs. Those authors working on evidentiality who are explicit about the concept of evidence all take this perspective, as far as I know. Here I will briefly review the most explicit theory I am aware of in linguistics, that of McCready and Ogata (2007). I will then indicate what points are left untouched, underspecified, or insufficiently argued for by these authors.

McCready and Ogata (2007) defined the function of evidence in (roughly) a Bayesian manner. This paper provides an analysis of certain Japanese evidentials, of two types: inferential evidentials, and hearsay evidentials.[2] The inferential evidentials were modeled using an operator $\triangle_a^i$, where i indexes an evidence source and a is an agent, whose effect can be stated informally as

[2]I will not have much to say about hearsay evidentials here; McCready and Ogata (2007) analyzed them as simple tests for the existence of an 'observation sentence' in the Quinean sense (Quine, 1960) of hearsay type, with content identical to that in the scope of the evidential. See McCready (2011b) for some additional discussion.

follows: $\triangle_a^i \phi$ is true given a world w, time s, agent a and probability function μ iff: (i) ϕ was less likely according to a as determined by μ at some time preceding s (before introduction of some piece of evidence i), (ii) ϕ is still not completely certain for a at s (given i), and (iii) the probability of ϕ for a never decreased between the time a became aware of the evidence i and s as a result of the same piece of evidence i (i.e., the probability of ϕ given i is upward monotonic). Even more informally, then, observation of i made ϕ likely but not certain, and i never subsequently had the effect of lowering the likelihood of ϕ. The observation of the evidence itself was modeled with a predicate E. This predicate also serves a complex function. Informally, $\mathsf{E}_a^i \varphi$ (i) changes the probabilities assigned to every proposition ψ (excluding φ itself) in the current information state σ by replacing them with the conditional probability of ψ given φ, if defined, and (ii) replaces the modal accessibility relation with one restricted to worlds in which φ holds. This account is meant as a treatment of what evidence does in a context; it changes the probability of other propositions that are related to it, and revises the set of accessible possibilities to one containing only those possibilities that make the content of the evidence true.[3]

In the definitions of these two operators, we have an implicit account of what a piece of evidence does: it alters the probability of certain propositions holding true, and also induces changes in the general epistemic space of the agent acquiring the evidence. These aspects seem uncontroversial, at least in an intuitive sense. A number of questions remain unanswered, however. We might ask: is mere increase in subjective probability sufficient for use of evidentials? Whose subjective probability function is at issue? Need all evidence be propositional, as the account implies? Are there no other requirements on evidence than that it monotonically increases the probability of the evidenced? The further question of how to give a general identification of the propositions which, by this definition, count as evidence (those which raise the probability of whatever we are interested in) I will have to put aside in this paper, though it might be statable in terms of defeasible reasoning: evidence talk is statable in terms of defeasible reasoning : ψ is evidence for ϕ iff $\psi > \phi$ as normality conditional (Pollock and Cruz, 1999). Obviously, the connections between reasoning in terms of probabilities and using defeasible inference go deep (Halpern, 2003). I will not explore this connection here, but it isn't obvious to me how distinct the two kinds of theory really are in terms of empirical predictions (as opposed to philosophical commitments).

Let us now consider the theory of Davis et al. (2007). This theory, like those discussed in the first section, takes evidence to be a primitive; however, it is quite explicit about what the evidential itself is meant to do with respect to justificational requirements for assertion of sentences including evidentials. On the natural assumption that evidence ties directly to justification, the theory therefore makes indirect claims about what evidence *does* (as opposed to what it *is*).

What are the justificational norms of assertion? This topic remains controversial in philosophy. Some possibilities: assertions might require simple belief of the speaker about the asserted proposition (Bach and Harnish, 1979), they might require knowledge (Williamson, 2000), or they might require something weaker.[4] Davis et al. (2007) assume a version of the last view on which the context makes available an assertability threshold; this is a real number in [0,1] corresponding to

[3]McCready and Ogata (2007) did this by revising the accessibility relation; it could also be done by revising the set of available possibilities *qua* set, as was done in the analysis of modality of Asher and McCready (2007).

[4]See the papers in Brown and Cappelen (2011) for extensive and interesting recent discussion of this issue.

a probability. Individuals are associated with (subjective) probability functions. If the probability the speaker assigns to φ is higher than the threshold, then φ is assertable for that speaker. The basic strategy is borrowed from Lewis (1980); the realization here is similar to work on gradable adjectives by Kennedy (2007) and others.

What is the function of evidentials in this theory? According to Davis et al. (2007), evidentials serve to shift the assertability threshold. Direct evidentials leave the threshold unchanged; hearsay evidentials and other kinds of evidentials marking indirect information sources lower it. The upshot is that evidentials can strengthen or weaken justificational requirements on assertion; this part of their meaning can be characterized as pragmatic. Semantically, they simply 'commit the speaker to' the existence of evidence of the relevant type. Exactly what this commitment amounts to is not clear (presupposition? assertion? conventional implicature?), but this is perhaps as desired by the authors, since these aspects of evidential meaning seem to vary from language to language, and indeed from evidential to evidential within a single language. With this move, the authors avoid the central question of this paper, as they take evidence to be primitive. However, their proposal is of interest as it concerns directly the question (raised in the introduction) of how evidential information is to be processed. Fuller discussion can be found in McCready (2011b).

One can find other notions of evidence in the literature on evidentials that take more controversial philosophical positions. For instance, Bronnikov (2011) provides a general model of cognition and inference; this model assumes a version of the representational theory of mind (cf. Fodor 1987, 2000) and characterizes semantic content based on its behavior with respect to the proof theory of such representations. Evidentials then are dealt with in one of two ways: either (a) as existence checks on the statement of the evidence-relevant content in the (representational) mental model, or (b) as attempts to prove that the sentence representations are present in distinct 'boxes' corresponding to different means of acquiring information. In effect, there is no longer a need for a separate theory of evidence, only for a theory of how bits of content behave in the proof theory of mental representations: a fairly radical move. Whether this looks desirable will depend on one's attitudes toward the representational theory of mind in general, and how exactly one spells out the proof theory of evidence, something about which Bronnikov (2011) is not very specific. For present purposes, the proposal is not spelled out enough to be viewed as a proposal about the nature of the evidence manipulated by evidentials.

2.2 Some relevant facts

In this section I would like to exhibit some data that speaks to the kind of evidence required for use of natural language evidentials. This data is of two types. The first involves situations where an agent alters subjective probabilities based on purely mental phenomena. It shows that simple subjective probability increase is not enough to ground use of evidentials. The second type involves calling into question the justificatory base itself in various ways. It shows that the putative evidence must be taken by the evaluator to properly match the external world in order to be a basis for use of evidentials.

2.2.1 Subjective probability is not enough

So far we have seen views on which evidence corresponds to changes in subjective probability. An obvious question that arises with respect to such views is this: do all changes in probabilities follow from observation of evidence? Clearly not: consider, for instance, a situation in which Jerome hopes to get a particular job, and, although he is unqualified for it, allows his hope to influence his perceived likelihood of employment to the degree that it affects his actions. Another instance, even more irrational, involves a person who lets her desire to win big slowly raise the likelihood she assigns to winning the lottery to a point where she actually believes that she will win, after her purchase of one ticket.[5] Such cases easily lead to irrational behavior.[6] One response to this sort of case is to identify evidence with knowledge: one certainly would not like to say that Jerome knows he will get the job, though he may believe it with a high degree of credence. This is the $E = K$ thesis of Williamson (2000). This means that, if evidentials do indeed look for evidence in the usual epistemological sense, the relevant evidence must be known by the speaker. This position was taken by McCready (2008, 2010). It is supported by the data adduced in the next section, at least on one interpretation. However, $E = K$ by itself means little without an account of what knowledge *is*, itself a contentious issue; as we will see, given the facts about use of evidentials, $E = K$, while it may be correct, does not do much to advance our understanding of the issue.

2.2.2 Evidentials in skeptical scenarios

How to tell knowledge from belief? Clearly, for the truth of both $B_s\phi$ and $K_s\phi$ the speaker must assign an extremely high degree of credence to ϕ.[7] Here is a traditional answer from epistemology: knowledge is justified true belief. I can be said to know p if I believe p, p is true, and I have good reason to believe p This answer looks reasonable, and many people have espoused some version of it. But it is wrong, as the epistemologists well know. Gettier (1963) discovered examples in which all the conditions above are met, but still there is no knowledge. Here is a scenario in the Gettier style. Johnny is traveling in the country when he sees what looks to him like a horse on top of a hill and hear a horse neigh. However, what he sees is a horse-shaped rock, and the neigh is just the wind whistling through that pipe over there. But there is—coincidentally—a horse standing behind the rock. Now consider this sentence:

[5] These cases are modeled after an example in Williamson (2000).

[6] See Fantl and McGrath (2009) for extensive discussion of the relation between knowledge and practical reason, as well as Hawthorne (2004).

[7] I assume a probabilistic account for consistency with the previous (Bayesian) discussion; one could also transpose this view to a more linguistically-traditional possible worlds-based picture of attitudes, so that

$$[\![B_a\phi]\!] = 1 \text{ iff } \frac{card(\{w' : R_aww' \ \&\ \phi(w')\})}{card(\{w' : R_aww'\})} > s,$$

i.e. the proportion of the agent's epistemically accessible worlds verifying ϕ exceeds some contextually specified degree. The requirement for knowledge would then be to further increase the required proportion (to something approximating 1) or to add extra conditions, as in the main text below. I do not see much to choose between the two pictures, at least for the present application.

(4) Johnny knows there is a horse on top of the hill.

This statement seems false–though the conditions listed are satisfied: Johnny believes that there is a horse on top of the hill, there is in fact a horse there, and Johnny has good reason—in fact two good reasons—to believe there is one there, at least from his own perspective.

The above considerations suggest a way to distinguish knowledge from belief: if one can destroy the justification for the putative piece of knowledge, yet there is no change in the (subjective) cognitive status of the object of the attitude, then it is belief. If the cognitive status of the putative knowledge changes together with the change in justification—if it becomes uncertain or eliminated—then the putative knowledge is knowledge indeed. Obviously, this won't work for every sort of knowledge (for instance, it has nothing to say about a priori knowledge, by definition). Still, this way of distinguishing things seems a good first approximation, and we will see the result of working with this assumption in this section. It is a bit too simple, though: I have obviously left out the problem of determining what the relevant justification is supposed to be, and for what agent or agents its status changes. As we'll see, both these points turn out to be crucial for the full story. But this characterization allows us to proceed to get some initial data to go ahead with. [8]

The strategy, then, is to call into question the justification for the evidence. We will first use the most extreme form of this general strategy: the *skeptical argument*. Skeptical arguments call into question the foundations of all our knowledge (for some given area). They have the following general form: one introduces possibilities which falsify all—or some relevant portion of—our putative knowledge and cannot be conclusively eliminated. Because we cannot eliminate them, possible flaws in the foundations of our knowledge enter our awareness. In view of these potential errors, we become uncertain about the solidity of our knowledge. As a result, our knowledge disappears. One can think of this effect in various ways, for instance as a change in the contextual standards for knowledge attribution (e.g. DeRose 1992; Lewis 1996; DeRose and Warfield 1999 provides an overview of other possibilities; general background on epistemological stances can be found in Pollock and Cruz 1999). For our purposes we need not take a stand on which, if any, of the currently available positions is correct.

Skeptical scenarios usually look highly implausible. Some traditional examples include the possibility that you might be deceived by an evil demon into believing that you are receiving certain kinds of perceptual input, such as that you are drinking a cup of coffee; that you might be a brain in a vat, with your perceptual centers stimulated by electric impulses, or in a Matrix-like situation; that you might be dreaming everything you are perceiving, or be in a catatonic state. The common characteristic to these scenarios is that, in each, the sensory data you receive is not trustworthy as a guide to what actually is. Note the similarity to the Gettier scenarios. The difference between skeptical scenarios and Gettier cases is that, in the skeptical scenarios, there is no possibility for the individual in the scenario to either learn that he is in fact in such a scenario or to conclusively prove

[8]The linguistically minded reader may now be wondering why we need to go to all this trouble. After all, isn't knowledge factive, and belief not? That means that the object of knowledge is presupposed, but not so in the case of belief. If this is so, then why must we worry about justification and the foundations of knowledge? There is some initial plausibility to this objection, but it rests on a confusion. The verbs *know* and *believe* are factive and not factive respectively, but here we are not interested in knowledge or belief as it is linguistically expressed. Rather, we are interested in evidence, as the object required for the felicitous use of evidentials. This content is not explicitly expressed in language. To find out its properties, we must take a more indirect route.

that he is not in one, because all of his sensory input is open to question, while in the Gettier cases, a world-internal observer could make the Gettiered individual aware of his error. To anticipate the later discussion, this distinction turns out to play a role in the use and evaluation of evidentials.

We can also find scenarios that look more common-sensical, especially when we confine ourselves to scenarios that only cast doubt on certain types of knowledge or knowledge in certain domains. For instance, consider a scenario on which you fell down a moment ago and hit your head, and the resulting damage caused you to hallucinate your current state—you appear to be reading this paper, but in fact you are lying on the floor outside your office viewing an internal projection of what you had planned to do before your injury. This situation seems quite normal compared to those above, but only calls into question your knowledge of your present activities, rather than of your entire set of memories.

What all skeptical scenarios have in common is the property that—if taken seriously—they destroy knowledge. For any p that one putatively knows (or for any salient p, for limited skeptical scenarios like the above), one may retain the belief that p but this belief can no longer be conclusive. There is always a possibility of error. Such beliefs are thus no longer knowledge in the strict sense. As a result, skeptical arguments can be viewed as tests for knowledge, when used on susceptible speakers. By running a skeptical argument on someone who is willing to consider them seriously, one can test whether a particular bit of their cognitive state is knowledge or belief, in the following sense: if the skeptical argument has no effect on the cognitive status of the content of interest, that content is merely believed.

To believe something, one must assign it a degree of subjective probability higher than whatever the threshold for belief is taken to be. In general, this threshold is contextually determined in the usual way familiar from degree predicates (cf. Kennedy 2007 on degree predicates and Stanley 2005 on belief in particular); skeptical arguments are implausible enough that they will not (barring an extremely high contextual standard) rule out belief, for they lessen degrees of subjective probability in a very minor way. Thus beliefs can survive skeptical scenarios, but knowledge cannot. We thus have a way to distinguish knowledge from belief. One application of this tool, the one I am concerned with in this paper, is in determining whether mere belief is acceptable for the use of evidentials. I now turn to this application.

How can one use skeptical arguments in the desired way? The idea is straightforward. First, give a speaker a piece of evidence supporting some conclusion φ in the intuitive sense. As before, I temporarily sidestep issues concerning exactly what should count as evidence for some conclusion; I will discuss some of these issues later. After providing the evidence, ask whether $Evid(\varphi)$ is true (or assertable, depending on the language). This step ensures that the piece of evidence is the right kind to license the evidential in general. Here is an application of this test to the case of Japanese evidentials. Under ordinary circumstances, the observation that the street is wet outside in the morning leads to a rise in the probability that it rained the night before. So, by the definitions above, it should count as evidence, and be sufficient to license the inferential evidential *mitai*.

(5) michi-ga nureteiru. kinoo-no ban ame-ga futta mitai
street-Nom wet yesterday-Gen night rain-Nom fell INF.EVID

'The street is wet. It must$_{inf}$ have rained last night.' (spoken in the morning)

(6) kinoo ame-ga futta soo-da
yesterday rain-Nom fell EVID.HRSY-Cop

'[I heard that] it rained yesterday.' (spoken after John said that it rained yesterday.)

This is correct. In this case, the evidential sentence is assertable. So the sequence $E; Evid(S)$, where E is the evidence and $Evid(S)$ the sentence containing the evidential, is a felicitous one.

The test for knowledge comes when we introduce a skeptical scenario after the evidence. Here is an English version.

(7) The street is wet. But perhaps there is no street—perhaps you are just dreaming. (Anyway,) It rained last night–Evid$_{inf}$.

(8) John said that it rained yesterday. But maybe you were just hallucinating when you saw John. It rained yesterday–Evid$_{horsy}$.

Now ask the speaker: Is the new sequence $E; S; Evid(S)$ acceptable, where S is the skeptical scenario? Or, for languages where we can consider the evidentials primarily truth-conditional, is the sentence containing the evidential judged true in this new context? If the new sequence is acceptable, and the sentence containing the evidential is true, then the evidence required does not need to be actual knowledge: belief is sufficient. We know this because the skeptical scenario, if taken seriously, destroys knowledge; so if the evidence must be knowledge, then the sentence with the evidential would be bad. Conversely, if the new sequence is not acceptable, or the sentence with the evidential is false, then knowledge is required.

So that is the test. What are its results? I have tried this test on a number of Japanese speakers. A few 'skeptical' individuals were unwilling to take the skeptical arguments seriously. Disregarding these subjects, no speaker allows sentences with evidentials after the skeptical scenario is introduced. This suggests very strongly that the evidence needed for Japanese evidentials is not as simple as plain vanilla subjective belief.

Here is a possible objection to the test. It might be suggested that my informants are just balking at asserting anything about the world, given that I have called into question all their knowledge of it, and its very existence. This objection has some initial plausibility, but when examined closely, lacks force. It contains two subarguments. The first involves assertion: the unstated assumption is that, without full confidence, one cannot assert anything. This unstated assumption is false. To assert, knowledge is not necessary—we do not even need total belief. Belief beyond reasonable doubt is sufficient, where the level depends on context (again, see Stanley 2005 or Davis et al. 2007 for more discussion).[9] In any case, the objection depends on the particular skeptical scenario chosen above, which did in fact call into question everything about the world. But it is easy enough to change the scenario in such a way that we limit its application to the case at hand. Here is an instance. I give only the English version for readability.

[9]This position is in opposition to the knowledge-norm view of Williamson (2000) and others. On this view, speakers should, normatively, only assert the things they know. I believe that this view is far too strong—it is tantamount to forcing people to have full credence in any content they assert, which is virtually impossible in the real world.

(9) The street is wet. [But you may have a brain tumor that causes all streets to look wet, even though they are not. You cannot be sure if the street is truly wet or not.] It rained-Evid$_{inf}$.

This new scenario only calls into question the speaker's knowledge of street wetness. The rest of the world remains untouched. Nonetheless, speakers are reluctant to use evidentials in scenarios like these as well. I conclude that the apparent flaw in the test is only apparent, and that evidence for evidentials—in Japanese at least—must be knowledge.

2.2.3 Gettiered evidentials

We are now evaluating the claim that the evidence needed for evidentials must be stronger than increase in subjective probability. When the evidence was called into question via the application of skeptical arguments, evidentials could no longer be used (in Japanese). A natural question to ask now is whether the same thing happens in Gettier cases. This section explores the facts in this domain.

Recall one primary difference between Gettier cases and skeptical scenarios: the Gettiered individual is Gettiered because of facts about the world, whereas the victim of a skeptical argument has his knowledge destroyed precisely because the facts about the world became uncertain. The crucial point is that, while the justification the Gettiered individual has for his beliefs is not well-founded, this lack of justification can be apparent to other individuals in the Gettier case. Thus we see that being Gettiered is a perspective-dependent problem: only the Gettiered individual is necessarily Gettiered. In skeptical scenarios, however, there is no way to determine whether the skeptical argument is true; all individuals have an equal lack of access to the 'actual' situation. (In the case of skeptical scenarios limited to cases like the brain tumor scenario above, this assessment remains valid: the individual who sees the street as not wet in (9) might equally be the victim of a brain tumor of a different kind.) This discussion suggests that we may expect to find differences between felicitous uses of evidentials in the two kinds of cases.

This expectation is fulfilled. Unsurprisingly, the Gettiered individual can assert an evidential with respect to his putative knowledge:

(10) ano oka-no ue-ni uma-ga iru mitai da
that hill-Gen top-Dat horse-Nom exists EVID Cop

'There appears to be a horse on top of that hill.' (said by the Johnny of (4))

For the outside observer the situation is a bit more complex. We can distinguish two cases involving only failure of warrant. I leave aside cases where the object of belief is in fact false.

1. The observer knows that Johnny's warrant for belief is no good, but does not know whether there is actually a horse.

2. The observer knows both that Johnny's warrant is no good and and that there is a horse.

In both of these cases, (10) is bad. But it is bad for different reasons. In Case 1, it is bad because of the first clause (i) of the definition of the inferential evidential; the outside observer has no

piece of evidence—that is, no piece of *knowledge*, since Johnny's putative evidence is useless—that increased the probability that there is a horse on the hill to the necessary level. In Case 2, the observer runs afoul of the second clause (ii); since the observer *knows* that there is a horse, the probability she assigns to there being a horse approximates 1; she is completely certain that there is a horse, and the evidential sentence cannot be used. The situation of Case 2 is not particularly relevant for the present discussion as it involves something closer to a Gricean violation, modeled in the theory of McCready and Ogata (2007) as something akin to Veltman's (1996) examples with epistemic modalities:

(11) # It is sunny . . . It might be sunny.

If we know that it is sunny, it is not helpful to assert the possibility. The evidential case is analogous.

A question immediately arises when one considers the behavior of evidentials in Gettier scenarios, concerning the distinction between assertability and truth evaluation. The issue turns on the question of whose perspective is taken, both to the (putative) evidence itself and to the evaluation of the evidential-containing sentence. Consider again (10). Is this sentence assertable? By Johnny, yes; in his Gettiered state, he believes that he has evidence enough to make it true, so he can utter it sincerely. By a non-Gettiered observer, however, it is not assertable: in Case 1, the observer knows the putative evidence to be incorrect, so the evidential is false; in Case 2, the evidential statement is just inappropriate given the observer's knowledge. So we see that the perspective taken matters for assertability. Now consider what happens when we evaluate the truth of (10). Here again, Johnny himself will take (10) to be true—as will anyone Gettiered along with Johnny—but the outside, omniscient observer will take it to be false. So perspective matters for truth evaluation as well. I will return to this issue in much more detail in section 3.

2.3 Summary

In this section we have established the following. Japanese (inferential) evidentials cannot be used when skeptical scenarios call the evidence they rely on into question. They can be used by Gettiered individuals, but not by individuals who observe the Gettier scenario from an external perspective. These facts taken together suggest that the evidence required for evidentials is stronger than increase in subjective probability, and is indeed a kind of knowledge. However, we can't make it too strong: if we do, we eliminate the possibility for evidentials to be used by confused or Gettiered individuals, who clearly can use them; though such uses might be judged inappropriate or false by observers with the relevant discriminatory powers. But what other possibilities exist? To answer this question, we need to look at accounts of knowledge and justification that have been proposed in the epistemological literature. This is the task of the next section.

3 Evidence in epistemology

The way I will approach the problem of characterizing 'evidence for evidentials' is as follows: the question of what evidence is for natural language amounts to the question of how natural language

makes use of the notion of justification, i.e. the justification for assertion, and the evaluation of that justification. Thus we are interested in what notion of justification is at play in natural language. In this section, I will have a look at some existing views (in broad outline) with an eye toward determining what might be appropriate for natural language analysis.

3.1 Justification and knowledge

One major line of demarcation between theories of evidence and justification is that between internalist and externalist theories. Internalist theories, roughly speaking, are those which take justification to depend exclusively on the state of the agent, i.e. on her mental states; externalist theories conversely put some part of the responsibility for justifying in the external environment. Which kind of theory looks preferable for natural language evidentials?

We can start by considering Gettier cases. Here, Gettiered individuals are willing to assert evidentials. This means that we cannot have a purely externalist conception of evidence, for if we do, we predict that such assertions would be unwarranted. However, external observers judge such uses of evidentials false. One possible conclusion to draw is that two distinct conceptions are at work: assertions are performed according to an internalist conception of evidence, and evaluation is done according to an externalist conception. This sounds a bit odd. Can it be right?

Let us first consider what we might want for evaluation on this kind of view. We can imagine several different kinds of externalism that might work for evidentials. Two are identified by Fantl and McGrath (2009): 'radical' and 'moderate' externalism.[10] The two types differ in the degree to which they require externally verifiability. Radical externalism requires the existence of a genuine reliable method for knowledge acquisition. This reliability also must be externally determined, which of course means that Gettiered contexts are irrelevant; the method of justification used by a Gettiered individual is by assumption unreliable. The result is that evidential statements based on evidence gathered according to the methods used by Gettiered individuals will be evaluated as false in general. This prediction is obviously wrong. A second possible position is moderate externalism. On this view, knowledge requires justification, but justification is understood externally. This position might look correct for evaluation; but consider a small complication of the case. Suppose we have two Gettiered individuals. One makes an evidential assertion, and the other evaluates it: here, the evidential sentence is judged true or appropriate, but from an external perspective there is obviously no proper evidence. The conclusion is that, in fact, both use and evaluation of evidential sentences use the same metric, which then leads to a requirement for an internalist view of justification (and knowledge, if we adopt the $E = K$ thesis).

There is a third option: the knowledge-level justification of Fantl and McGrath (2009:97). Having knowledge-level justification means having a (justified) belief that you are justified enough to know, though you may in fact be mistaken. This amounts to having higher-order beliefs about the reliability of one's own justification. This view also comes in internalist and externalist versions; here, I will adopt an externalist version. However, knowledge-level justification is a weak enough notion that it will still apply in to Gettiered agents: such agents believe that they are justified

[10]These varieties of externalism are set up by Fantl and McGrath specifically for the analysis of knowledge, and so are adapted slightly to the equivalents for justification, our concern here.

enough to have knowledge.

On this view, though, how can we account for the difference in judgments about felicity of agents inside and outside Gettier scenarios? Assertion and evaluation are both perspective-dependent. We assert what we believe, or think we know; but doing so depends on what we know at the time of assertion. However, evaluation can take place with respect to a different perspective, in much the way explored for epistemic modals by Egan (2006) and MacFarlane (2011), or indeed the entire body of work on epistemic logic (Fagin et al., 1995; van Ditmarsch et al., 2007). Does this mean that we must follow the relativists and take modal, evidential, or indeed plain vanilla statements to be assessor-sensitive? I think not: evaluating something as true does not mean that it is true. The difference between objective truth and subjective truth evaluation seems to be a point of confusion in this entire literature. But this is not the place to pursue this point in detail. In the rest of the paper, I'll focus on the proper conditions for assertion and evaluation of evidentials. What we've seen so far is that, for assertion, an internalist model is nearly sufficient, i.e. Fantl and McGrath's 'knowledge-level justification', which, given that it is determined by agent beliefs about the reliability of their warrants, cannot but be essentially internalist in conception). However, the relevant notion must also involve some higher-order beliefs about the match between justification and external world. The same holds, mutatis mutandis, for the evaluation of evidential sentences for truth; but, since such evaluation is necessarily agent-dependent (since no evaluation happens except by the agency of some individual), this evaluation need not track actual truth at all. I will discuss the relation between these two aspects of linguistic practice in section 4.

3.2 Awareness

Before moving to a discussion of the options we have for modeling natural language evidence, there is one more point to consider, related intimately to perspective-dependence. A piece of evidence cannot count as evidence for an agent unless that agent is aware that the evidence is indeed evidence. In other words, if I cannot recognize something as evidence for something else, it fails to be so. This situation differs from evidence in e.g. philosophy of science, where we are more interested in evidence as an absolute; here, awareness of the evidence *qua* evidence is strictly required.

Failure to recognize evidence as evidence can be rooted in at least two sources. One might fail to recognize the relationship between the evidential proposition and the proposition which it is evidence for. A simple case of such a situation might involve a person in a new country who is not aware of all the relevant social conventions. For example, suppose a European moves to Japan to work in a Japanese company. She sees one employee bow deeply to another. However, she is not aware that the bow is *too* deep and has a sarcastic character indicating lack of respect; instead, she concludes that the first employee respects the other.[11] The problem here is that she was not sufficiently familiar with the relevant conventions and so was unable to recognize the particularized evidential relationship.

The other obvious source of failure is the lack of a relevant concept. This situation is brought out clearly by an example from Audi (2002), who writes "if a child has no concept of an insurance

[11] As in other societies, excessive use of honorific speech or 'honorific behavior' is naturally interpreted ironically.

adjuster, then seeing one examining a damaged car and talk to its owner about deductibles will not function as a source of justification for the proposition that this is an insurance adjuster" (Audi 2002:89). We could view this as a special case of failure of recognition. Still, this issue is not one that will arise often in the linguistic case, simply because lacking the relevant concept is enough to (virtually) guarantee that the speaker will not utter any evidentially marked sentence which contains terms denoting the content in question.

Plainly, if an agent fails to recognize something as evidence she should not be licensed to use an evidential. We would therefore like to build this restriction into the concept of evidence we use for evidential terms. This need ties closely together with the view of evidence as perspective-based that Gettier scenarios already make clear that we require. I will explore ways to spell this out in the next section.

4 Options for evidential evidence

4.1 Desiderata

We are now in the following situation. We require a theory of evidence on which evidence is sensitive to the awareness and perspectives of agents, in order to account for cases of failure to recognize evidence and the behavior of evidentials in internal Gettier cases. Still, as we have seen, in skeptical scenarios use of evidentials is not so good, and they are also judged false or inappropriate by external observers of Gettier scenarios; this means that we need some means of tracking the external environment as well—or, on the internalist view, tracking individuals' beliefs about their relation to the external environment. In the rest of this section I will explore a couple of ways of resolving this tension, and finally settle on a specific proposal. However, this proposal will itself not be the end of the story. I will show in section 5.1 that different evidentials can be subject to differing requirements on justification and reliability.

4.2 Knowledge

The first option is to accept Williamson's $E = K$ thesis and take the evidence required for evidential use to necessarily be knowledge. This move is straightforward and, in conjunction with a theory of the effect of evidence (e.g. raising subjective probabilities) yields a reasonable looking picture. In earlier work, I adopted this view *simpliciter* (McCready, 2008, 2010). However, I now believe that more is required.

The main reason involves the relation between this view and evidence. If one follows virtually all work in epistemology and assumes that knowing something requires more than simply believing it, difficulties arise for the fit between felicitous use of evidentials and true knowledge of the required evidence. Consider internal Gettier scenarios again. What does the agent know in cases like these? Recall the horse scenario from (4). Here, Johnny's evidence is his perceptual: hearing what he takes to be the sounds of a horse, and seeing what he takes to be the form of a horse. Does he know these propositions? I suppose it depends on what exactly they are. Is his evidence that he sees a horse, or that he sees a horse-like form? If the former, he does not know it, if the latter,

clearly, he does.[12] The conclusion of McCready (2008, 2010) that $E = K$ rested on missing this particular point. I now think the right answer is that his evidence is the latter proposition. His belief that there is a horse follows from a Millarian 'quasi-inference', which amounts to making the assumption that perception is reliable (Millar, 1991). If this is correct, then assuming that evidence must be knowledge does not help. Johnny does know his evidence. The problem lies in the reliability of the inference from perception of the horse-like form to the conclusion that there is a horse. This problem is completely separate from the question of whether or not the $E = K$ thesis is true. If the above reasoning is correct, we must look elsewhere for the characterization we need.

For completeness, we should consider the other possibility: that what Johnny knows is that there is a horse. If this is the evidence at issue, then, again, knowledge is not going to help us much: in this case, it will amount to mere belief, since there is no horse at all. In other words, the required characterization of belief is nothing more than Fantl and McGrath's 'knowledge-level justification', sufficient justification for the agent to believe that he knows and nothing more. But if this is all we have, then knowledge won't do much to help decide the case. The $E = K$ thesis thus seems to be completely inert in the present case.

4.3 Higher-order beliefs about probabilities

Here is a possibly better way to go. Suppose that we take the acquisition of evidence to be the self-ascription of an increase in subjective probability based on the putative evidence. (Let me put aside for a moment possible incoherencies in the above notion.) We then get the assertability of evidentials in internal Gettier cases: based on the putative evidence, the agent self-ascribes a rise in probability (which she would not do if she had all the relevant information). Conversely, from an external perspective, it is clear that the putative evidence is merely a confusion, so the external observer does not self-ascribe any changes in probability distributions. As a result, the evidential sentence will be assertible for the Gettiered individual, yet the uttered sentence will be judged false by the external observer. This is precisely as desired.

However, is the proposed definition coherent? The possible issue is this: subjective probability distributions are based on the degrees of belief agents assign to various propositions (whatever those may be; see Hajek and Eriksson 2007). Is it conceptually possible to have a subjective probability distribution without being aware of the probabilities that are assigned? In other words, are we always aware by definition of our subjective probability distribution? If we are so aware, then self-ascription follows and the theory fails to discriminate between cases.

There are a number of possible fixes for this problem. The most obvious is to simply assume that agents are not necessarily aware of what probabilities they assign to various propositions. This amounts to abandoning positive introspection for beliefs. The details of how this would go depend on the background theory of belief, but in standard Kripke semantics we have to give up the 4

[12] This issue is a difficult one for theories of direct evidentiality as well; how reliable is perception? I think the right move here is to assume something like Faller's 'best possible grounds', by which we can evade the problem. Surely perceiving a horse(-like form) is the best possible grounds one could have for believing that there is a horse, regardless of whether or not there actually is one. One also wonders in this context about the evidential basis of Wittgensteinian hinge propositions (Wittgenstein, 1991), and what evidentials are used with them. I do not have data that speaks to this issue.

axiom (transitivity on states). While this is nonstandard, there is not anything necessarily wrong with it. Still, one would prefer a theory of evidentiality which did not require making assumptions about the nature of belief; if anything, the implications ought to go the other way.

A second option would be to introduce more complex notions of introspection. For example, one might suppose that agents are introspective, but that there is some higher-order uncertainty about their belief states. We might model this as uncertainty about possible subjective probability distributions; agents could be associated with a (convex) range of such distributions, in terms of which introspective knowledge could be defined, for example in terms of the mean of a sampling value (cf. Kruschke 2010). The downside to this move is the now highly complex nature of belief states and subjective probabilities. Perhaps such states are empirically necessary, but, again, one would prefer to do without them if possible.

A third option would be to move away from subjective probabilities and have agents self-ascribe properties involving probabilities of other kinds. For instance, one might self-ascribe being in a world where the objective probability of ϕ increased on the basis of the truth of ψ, ie. a world in which ψ is evidence for ϕ. We now have something resembling the picture developed by David Lewis in Lewis (1979), though he did not talk about probabilities there. This seems like a reasonable possibility. I will return to it shortly, but first let us consider what might appear to be a simpler and more attractive way of ensuring self-ascription of probability increases.

4.4 Subjective probabilities with judges

Suppose that we assume subjective probability distributions as usual, but take the relevant agents to be, not simply individuals, but *judges*. The idea of a judge has been used pretty extensively in recent work on such phenomena as predicates of personal taste (Lasersohn, 2005). The basic idea is that terms like these have truth conditions that are, essentially, relativized to individual interpreters; thus, we can disagree about whether (for example) horse sashimi is tasty without either of us being strictly speaking wrong. This view has been extended to epistemic modals by Stephenson (2007), motivated by two factors: the apparent agent-dependence of the truth of statements including epistemic modals, and the availability of shifts in agent in certain contexts. As I said above, the former consideration does not strike me as compelling, due to the difference between truth proper and truth evaluation. The latter parallel is telling, though. Both predicates of personal taste and epistemic modals shift in terms of 'judge' under attitude predicates and in questions; we can find similar facts with Japanese experiencer predicates and certain indexicals (McCready, 2007).

(12) a. Horse sashimi is tasty.

b. It might be raining. (judge = speaker)

(13) a. Robert thinks horse sashimi is tasty.

b. Robert thinks it might be raining. (judge = Robert)

(14) a. Is horse sashimi tasty?

b. Might it be raining? (judge = hearer)

With embeddable evidentials in e.g. Japanese, we find precisely similar shifts: under attitudes, the source is the attitude holder, and in questions the hearer (in general), but in simple sentences, it is the speaker. (I omit specific examples for space reasons.) The parallel indicates that—if it is reasonable to assume judge sensitivity in personal taste predicates and modals—it is reasonable to assume judge sensitivity for evidentials as well.[13]

If so, we can derive certain aspects of the preferred interpretation of evidentials. But recall that the original motivation was to simplify the self-ascriptive aspect of linguistic evidence. The question to as therefore is: do we get *de-se*-ness for free on this theory? I think that the answer must be negative. For an analogy, we may ask whether people are necessarily aware of their tastes. Since taste predicates (on these theories) involve judge-sensitivity, we would anticipate that the same sort of self-ascriptive quality found with evidence should be found with taste predicates as well. If it is not, much of the reason to adopt a judge-based theory of evidence disappears. Note that this is not to say that evidentials themselves do not make use of whatever mechanism determines the perspective holder in taste predicates or epistemic modals, only that judge sensitivity is not in play at the deeper level of (subjective) evidence itself.

Can one be unaware of one's tastes? Clearly, the answer is yes. Imagine a woman or a man who likes only men who are bad for her/him but doesn't realize it,[14] or someone who finds themselves eating junk food all the time despite believing that she prefers healthy food. The behavior of such individuals shows that they have certain tastes, though consciously these tastes are not accessible to them, and, if asked, they might well deny those tastes. The view that people's behavior shows their tastes better than their overt opinions is the theory of revealed preference (Samuelson, 1938), which is widely accepted in economics, though not completely uncontroversial. Tastes, then, do not require self-ascription, though (I have argued, with others) evidence does. It thus seems to be an error to identify self-ascription with judge-sensitivity.

The reader might be worrying at this point about the connection between assertion and self-ascription. Someone who is not aware of a given preference would not assert that preference, just as someone who is not aware of some evidence would not assert the existence of that evidence. This is certainly true. However, as I said above, I am not disputing that both evidentials and taste predicates exhibit something like judge-sensitivity in interpretation.[15] The question here is the proper way to characterize the required notion of evidence. I have argued that, for this, we need something like the self-ascription of probability increases. The fact that self-ascription is not required for preference shows that the parallel with taste predicates cannot be the proper way to approach this particular question.

[13]Is it actually reasonable at all? That is a different question entirely. The fact that shift is not obligatory in questions, only preferred, begins to make it appear that it might *not* be desirable, given that one must assume something like optional application of some monstrous operator to derive the facts. Better might be a fully pragmatic story: from aspects of the meaning of the 'judge-sensitive' expressions, derive a preferred interpretation on the basis of independently motivated rules or knowledge of the language and how it is typically used. See McCready (2011a) for a framework that might be applicable in this context, given the right setup of lexical entries.

[14]Thanks to Jason Quinley (p.c) for this example.

[15]Or whatever the proper way is to characterize the relevant dependencies; see footnote 13.

4.5 Proposal

Accordingly, let me make a proposal which directly references self-ascription. The idea will be to take the usual account of evidence as increase in probability via conditionalization, and to self-ascribe the property of being in a world in which the required increase occurs. The proposal thus comes in two parts: the change in probabilities, and the self-ascription of that change.

Describing evidence in terms of probability is completely straightforward. What we want is just the following:

(15) φ is evidence for χ iff $\mathsf{P}(\chi|\varphi) > \mathsf{P}(\chi|\neg\varphi)$, where $(\chi|\varphi)$ is the conditionalization of χ on φ.

Should we think of the above conditionalization in terms of subjective probabilities or something else, say objective or logical probabilities? As discussed in the last sections, using subjective probabilities here probably requires a highly nonstandard logic of belief, and should be avoided if possible. However, using other theories of probability may have strange consequences, as it is perfectly possible for ψ to be evidence for ξ without any human being aware of this fact. If we tied the felicity of using evidentials to this probability change, wrong predictions about truth evaluation would result.

However, this problem can be avoided by introducing self-ascription. We need some notion of self-ascription (or reflexive awareness) anyway in order to account for other intuitive requirements on evidence, as with the Audi (2002) example of the insurance adjustor. A straightforward technical implementation of this idea is due to Stalnaker (2008). On this theory, propositional content consists of sets of world-individual pairs,[16] rather than the familiar sets of worlds: so, roughly, $[\![\phi]\!] = \{\langle w, i\rangle : \phi(w) \text{ for } i\}$. This corresponds to one way of understanding Lewis's (1979) view of 'propositions as properties,' which allows self-location in worlds in which certain propositions hold, as required for modeling the cognitive status of indexical content (Perry, 1979). Stalnaker ensures this by modeling belief by a strengthened version of the usual accessibility relations on worlds, so that $\langle w, i\rangle R\langle w', i'\rangle$ holds only if i believes in w that she is in world w' and that she is i'.

These two elements can be combined as follows. Suppose that subjective evidence, as required for use of evidentials, amounts to self-ascription in a world where such evidence exists. Let us define a predicate $EVID$ for convenience as follows:

(16) $EVID(\chi, \varphi) \longleftrightarrow \varphi$ is evidence for χ (cf. 15)

Then, for φ to count as subjective evidence for χ with respect to agent a, we require:

(17) $SE(a, \chi, \varphi) \longleftrightarrow$
$\forall\langle w, a\rangle[\langle w_@, a\rangle R_{dox}\langle w, a\rangle \rightarrow \langle w, a\rangle \in \{\langle w, i\rangle : EVID(\chi, \varphi)(w) \text{ for } i\} \cap \{\langle w, i\rangle : K_i\varphi(w) \text{ for } i\}]$

The above simply says that φ is subjective evidence for χ for agent a just in case, in all of the agent's belief-accessible worlds, φ is indeed evidence for χ (defined in terms of probability increase), and a knows, in that world, that φ. The latter amounts to implementing knowledge-level justification.

[16] Officially, Stalnaker's formulation uses pairs of *centers* and worlds, where centers are pairs of individuals and times. In this paper, I am not concerned with temporally dependent propositions, so I will redact this aspect of the theory.

Does this analysis meet the desiderata presented at the beginning of this section? Obviously it accounts for awareness: it is specifically designed to do so. Does it give the right results for the cases we considered? For basic cases of use of evidentials, it will give the right results: if the piece of evidence is actually true and actually a piece of evidence by the criterion in (15), use of the evidential will both be licensed and deemed true. For internal Gettier scenarios, the evidential user will be justified in use of an evidential sentence, given belief in the apparent evidence φ; this is as desired. Conversely, for an external observer, φ will be false in that observer's doxastically accessible worlds, if resolved to the proposition believed by the Gettiered individual; if resolved to the actual, misperceived fact, it will be true, but no longer evidence. This is also the right prediction. Finally, for the case of skeptical scenarios, the evidential user will no longer be able to self-ascribe the proper evidential relation. All in all, this notion of evidence looks the proper one for the analysis of natural language evidentials.

5 Two complications

That said, it is not the case that no problems remain, or that the above analysis exhausts all complications related to the notion of evidence. This section discusses two such complications. The present framework has at least a partial answer for both of them.

5.1 Japanese inferential evidentials and externalism

This paper has focused for its data on inferential evidentials in Japanese. However, not all such inferentials are created equal. Consider the following contrast. These sentences should be imagined as uttered in a scenario where the speaker is trying to extricate himself from a skeptical scenario: i.e., he is asking himself whether or not he is actually dreaming. In this situation, it is reasonable to say (18a), but not (18b).

(18) a. yume janai yoo-da
dream Cop.Neg YOO
'This seems not to be a dream.'

b. ?? yume janai rashii
dream Cop.Neg RASHII
As above; impossible when read as inferential evidential.

Why might this be? I see two possibilities. The first is that different evidence sources are required, despite both *yoo-da* and *rashii* being inferential evidentials (on this reading). I do not think that this can be correct, mainly because it seems impossible to characterize any differences that might exist between the evidence sources; the two appear nearly identical in this respect. Rather, the intuition is that the evidence required for *rashii* must be of a different order of reliability. This leads to the second option, that two different notions of evidence are at play behind the two evidentials, so that we have a language-internal variation across evidentials in terms of the required concept of justification. I would suggest here that, while the evidence required for *yoo-da* and similar evidentials

such as *mitai* is only weakly externalist in the sense of requiring knowledge-level justification as suggested above, *rashii* is stronger, and manipulates a moderate externalist concept of evidence. Since, by assumption, no reliable external evidence is available in skeptical scenarios, we then have an explanation of the oddity of (18b). If this is correct, we then leave open the possibility of variation in evidential character across whole languages or language groups. I do not know whether this possibility is otherwise justified or realized; empirical work on evidential languages with different behavior than Japanese is needed to determine whether this is so.

One might ask whether this (putative) difference correlates with any other behavior of these two evidential types. The answer is positive. First, unlike e.g. *yoo-da* or *mitai*, *rashii* cannot embed under negation.

(19) * ame-ga furu rashii (to iu) koto-wa nai
rain-Nom fall RASHII (C say) thing-Top Cop.Neg

'It is not the case that it seems it will rain.' (intended)

Secondly, *rashii* behaves differently from the other evidentials in conditional antecedents. Japanese has four distinct types of conditionals with subtly different semantic qualities (see e.g. Masuoka and Takubo 1989 for details). We find the following paradigm with respect to the felicity of the various evidentials in Japanese conditional clauses.

	rashii	mitai, yoo-da
to	×	×
-ba	×	√
-tara	×	√
nara	√	√

I do not know what the source of these differences is, or how exactly they connect up to the facts in (18), but my feeling is that the three are related. Is there an explanation for the impossibility of embedding of *rashii* and its incompatibility with certain kinds of conditionals that follows from stricter requirements on evidential quality? I am not sure; but if there is, it could have interesting repercussions on the proper analysis of evidentials in languages where embedding is impossible, such as the well-studied Cuzco Quechua. I leave this project for future work.

5.2 Causality

The second issue is one that the present paper might appear to have little or nothing to say about. The following dialogue is a translation from the Japanese; 'must_{inf}' picks out an inferential evidential like *mitai*, which, as we have seen, is relatively weak. Still, infelicity arises in cases like this one, where the causal nature of the connection between the evidence itself and what it is supposed to be evidence for is called into question.

(20) a. A: It must_{inf} have rained last night.

b. B: Why do you say so?

c. A: Well, the ground is wet.

d. B: But that's from the sprinkler.

e. A: Then I was wrong./# Still, it must$_{inf}$ have rained.

Note that in this case all requirements I specified are satisfied: the speaker has self-ascribed an increase in probability, which is, prior to learning the actual cause of the ground being wet, completely justified as far as we can tell. Does this mean that more is required for evidence in natural language, perhaps even genuine causality? I do not think so. This issue is closely related to Gettier cases. It is just that here the Gettiered individual becomes 'un-Gettiered' after learning a piece of relevant information. After so learning, he is in a position to evaluate his previous statement as involving a spurious correlation, and consequently no longer self-ascribes a probability increase. Note the similarity to the cases of changes in judgements on epistemic modal statements discussed by e.g. von Fintel and Gillies (2008); again, we have not so much a change in truth-value as a change in truth-value judgement. Note though that cases like these require us to deploy the full Stalnaker system, as we require the ability to model temporal dependence of (self-locating) belief.

6 Conclusion

This paper has proposed an analysis of the evidence required for use of natural language evidentials. According to this proposal, a proposition is evidence for χ for an agent only if, after learning it, the agent self-ascribes an increase in the likelihood of χ. I showed that this view is capable of making sense of complex data involving skeptical scenarios and Gettier cases, both internal and external. Further research is required to determine whether the concept proposed is a universal one or whether it is specific to the particular language studied, and even to the particular evidentials examined within that language. My suspicion is that it is indeed universal, but that certain evidentials (like *rashii*) may have more stringent additional requirements. It is my hope that the present proposal puts the study of evidentiality onto a somewhat firmer footing than previously, and that the results of the paper about subjectivity and evaluation have helped to show something about the nature of evidential knowledge.

References

Asher, Nicholas and Eric McCready. 2007. Were, would, must and a compositional account of counterfactuals. *Journal of Semantics* 24(2):93–129.

Audi, Robert. 2002. The sources of knowledge. In *The Oxford Handbook of Epistemology*, pages 71–95. Oxford University Press.

Bach, Kent and Robert Harnish. 1979. *Linguistic Communication and Speech Acts*. MIT Press.

Bronnikov, Georgui. 2011. *Representation of Inference in the Natural Language*. Ph.D. thesis, University of Texas at Austin.

Brown, Jessica and Herman Cappelen, eds. 2011. *Assertion*. Oxford University Press.

Chung, Kyung-Sook. 2005. *Space in Tense: The Interaction of Tense, Aspect, Evidentiality and Speech Act in Korean*. Ph.D. thesis, Simon Fraser University.

Chung, Kyung-Sook. 2007. Spatial deictic tense and evidentials in korean. *Natural Language Semantics* 15:187–219.

Davis, Christopher, Christopher Potts, and Peggy Speas. 2007. The pragmatic values of evidential sentences. In M. Gibson and T. Friedman, eds., *Proceedings of SALT 17*. CLC Publications.

de Haan, Ferdinand. 1999. Evidentiality and epistemic modality: Setting boundaries. *Southwest Journal of Linguistics* 18:83–101.

DeLancey, Scott. 1997. Mirativity: the grammatical marking of unexpected information. *Linguistic Typology* 1:33–52.

DeRose, Keith. 1992. Contextualism and knowledge attributions. *Philosophy and Phenomenological Research* pages 913–929.

DeRose, Keith and Ted Warfield, eds. 1999. *Skepticism: A Contemporary Reader*. Oxford University Press.

van Ditmarsch, Hans, Wiebe van der Hoek, and Barteld Kooi. 2007. *Dynamic Epistemic Logic*. Berlin: Springer.

Egan, Andy. 2006. Epistemic modals, relativism and assertion. *Philosophical Studies* 133:1–22.

Fagin, Ronald, Joseph Halpern, Yoram Moses, and Moshe Vardi. 1995. *Reasoning about Knowledge*. MIT Press.

Faller, Martina. 2002. *Semantics and Pragmatics of Evidentials in Cuzco Quechua*. Ph.D. thesis, Stanford University.

Faller, Martina. 2004. The deictic core of 'non-experienced past' in Cuzco Quechua. *Journal of Semantics* 21:45–85.

Faller, Martina. 2007. The Cuzco Quechua reportative evidential and rhetorical relations. In A. Simpson and P. Austin, eds., *Linguistische Berichte*, vol. Sonderhafte 14, pages 223–252. Helmut Buske Verlag.

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

Fodor, Jerry. 1987. *Psychosemantics*. Cambridge, MA: MIT Press.

Fodor, Jerry. 2000. *The Mind Doesn't Work That Way*. MIT Press.

Garrett, Edward. 2001. *Evidentiality and Assertion in Tibetan*. Ph.D. thesis, UCLA.

Gettier, Edmund. 1963. Is justified true belief knowledge? *Analysis* 23:121–123.

Hajek, Alan and Lina Eriksson. 2007. What are degrees of belief? *Studia Logica* 86:185–215.

Halpern, Joseph Y. 2003. *Reasoning about Uncertainty*. Cambridge, MA: MIT Press.

Hawthorne, John. 2004. *Knowledge and Lotteries*. Oxford.

Kennedy, Chris. 2007. Vagueness and gradability: The semantics of relative and absolute gradable predicates. *Linguistics and Philosophy* 30(1):1–45.

Krifka, Manfred. 1992. Thematic relations as links between nominal reference and event domains. In I. Sag and A. Szabolcsi, eds., *Lexical Matters*, pages 29–53. Stanford, CA: CSLI Publications.

Kruschke, John. 2010. Doing bayesian data analysis. Manuscript, Indiana University.

Lasersohn, Peter. 2005. Context dependence, disagreement, and predicates of personal taste. *Linguistics and Philosophy* 28:643–686.

Lewis, David. 1979. Attitudes *de dicto* and *de se*. *The Philosophical Review* 88:513–543.

Lewis, David. 1980. A subjectivist's guide to objective chance. In R. Jeffrey, ed., *Studies in Inductive Logic and Probability*, vol. 2. University of California Press. Reprinted in David Lewis, *Philosophical Papers*, volume 2.

Lewis, David. 1996. Elusive knowledge. *Australasian Journal of Philosophy* 74:549–67.

MacFarlane, John. 2011. Epistemic modals are assessment-sensitive. In B. Weatherson and A. Egan, eds., *Epistemic Modality*. Oxford University Press.

Masuoka, Takashi and Yukinori Takubo. 1989. *Kisoo Nihongo Bunpoo [Essential Japanese Grammar]*. Tokyo: Kuroshio Shuppan.

Matthewson, Lisa, Hotze Rullmann, and Henry Davis. 2007. Evidentials as epistemic modals: Evidence from St'át'imcets. In J. van Craenenbroeck, ed., *Linguistic Variation Yearbook 2007*. John Benjamins.

McCready, Eric. 2007. Context shifting in questions and elsewhere. In E. Puig-Waldmüller, ed., *Proceedings of Sinn und Bedeutung 11*, pages 433–447. Barcelona: Universitat Pompeu Fabra.

McCready, Eric. 2008. Evidentials, knowledge and belief. In Y. Nakayama, ed., *Proceedings of LENLS 5*.

McCready, Eric. 2010. Evidential universals. In T. Peterson and U. Sauerland, eds., *Evidentials and Evidence*, vol. 28 of *UBC Working Papers in Linguistics*, pages 105–128. University of British Columbia.

McCready, Eric. 2011a. Emotive equlibria. Manuscript, Aoyama Gakuin University and UT-Austin.

McCready, Eric. 2011b. Testimony, trust, and evidentials. Manuscript, Aoyama Gakuin/University of Texas at Austin.

McCready, Eric and Nicholas Asher. 2006. Modal subordination in Japanese: Dynamics and evidentiality. In A. Eilam, T. Scheffler, and J. Tauberer, eds., *Penn Working Papers in Linguistics 12.1*, pages 237–249.

McCready, Eric and Norry Ogata. 2007. Evidentiality, modality, and probability. *Linguistics and Philosophy* 30(2):147–206.

Millar, Alan. 1991. *Reasons and Experience*. Oxford University Press.

Murray, Sarah. 2010. *Evidentiality and the Structure of Speech Acts*. Ph.D. thesis, Rutgers.

Ogata, Norihiro. 2005. A multimodal dynamic predicate logic of Japanese evidentials. Paper presented at *Language Under Uncertainty* workshop, Kyoto University.

Perry, John. 1979. The problem of the essential indexical. *Noûs* 13:3–21.

Pollock, John and Joseph Cruz. 1999. *Contemporary Theories of Knowledge*. Rowman Littlefield.

Quine, Willard Van Orman. 1960. *Word and Object*. Cambridge: The MIT Press.

Samuelson, Paul. 1938. A note on the pure theory of consumers' behavior. *Economica* 5(17):61–71.

Sapir, Edward. 1922. The Takelma language of southwestern Oregon. In F. Boas, ed., *Handbook of American Indian Languages*, vol. 2, pages 1–296. Government Printing Office.

Stalnaker, Robert. 2008. *Our Knowledge of the Internal World*. Oxford University Press.

Stanley, Jason. 2005. *Knowledge and Practical Interests*. Oxford University Press.

Stephenson, Tamina. 2007. Judge dependence, epistemic modals, and predicates of personal taste. *Linguistics and Philosophy* 30(4):487–525.

Veltman, Frank. 1996. Defaults in update semantics. *Journal of Philosophical Logic* 25:221–261.

von Fintel, Kai and Thony Gillies. 2008. CIA leaks. *The Philosophical Review* 117:77–98.

Williamson, Timothy. 2000. *Knowledge and its Limits*. Oxford.

Wittgenstein, Ludwig. 1991. *On Certainty*. Wiley-Blackwell.