Full Belief and Loose Speech

Sarah Moss

University of Michigan

This paper defends an account of the attitude of belief, including an account of its relationship to credence. The notion of *full* or *outright* belief discussed in this paper is central to traditional epistemology and the philosophy of mind. The notion of credence—also known as degreed belief, or subjective probability—is central to formal epistemology. It is notoriously difficult to spell out the relationship between these notions. To take one familiar question: does believing a proposition require having credence 1 in it? On the one hand, fully believing a proposition seems incompatible with doubting it. On the other hand, it seems that any interesting attitude of belief could not possibly require this maximal credence, since strictly speaking, we are rarely—if ever—maximally confident of any proposition.

      In addition to questions about belief and credence, an account of belief must address a number of other difficult questions. Are rational beliefs closed under entailment? On the one hand, we are generally moved to eradicate inconsistencies that we discover in our beliefs. On the other hand, preface paradoxes seem to show that rational subjects

can indeed believe inconsistent propositions.[1] Does whether you believe a proposition depend partly on your practical interests? On the one hand, it is sometimes argued that belief is interest relative, since you fully believe a proposition just in case you have enough confidence in it to act on it given your interests. But it also seems that rational belief cannot be interest relative, since you should not change what you believe merely because your interests change—at least in cases where your beliefs are evidentially independent of facts about your interests.[2] The account of belief defended in this paper yields answers to these difficult questions.

In addition to addressing familiar questions, my account entails some surprising and controversial claims about belief. For instance, contemporary discussions of belief often start by assuming that there are in fact two very different belief attitudes. It is assumed that there is a strong attitude of belief studied by epistemologists—namely, full or outright belief—and also a second attitude of belief, which the folk use 'believes' to talk about. It is taken to be controversial whether the former attitude of strong belief requires maximal confidence. By contrast, it is taken to be uncontroversial that weak belief does not require maximal confidence, since the folk can truly say sentences such as:

(1)   I *believe* it will rain this afternoon, but I'm not *sure* that it will.

Against this common set of assumptions, I argue that 'belief' is *not* used for distinct attitudes of strong and weak belief. The folk use 'believe' in sentences like (1) to ascribe the very same doxastic attitude studied by epistemologists—namely, the singular attitude of belief.

To sum up for easy reference, my paper contains the following sections:

3. An account of belief
4. Does belief require maximal confidence?
5. Is there a distinction between strong and weak belief?
6. Are rational beliefs closed under entailment?
7. Is belief interest relative?

The central spirit of this paper is one of reconciliation. As I address the above questions, my account aims to resolve the tension between our

1. For sympathetic discussion of this conclusion, see Makinson 1965, Foley 1993, and Christensen 2004.
2. For development of this argument and a survey of its targets, see Ross and Schroeder 2014.

conflicting intuitions about them, and also between conflicting answers to the questions that have been defended in the literature, explaining why each answer seems correct.

Before I state my account of belief, it will be useful to make a detour through some background in the philosophy of language. I begin by introducing *context probabilism*, a compelling thesis about conversational context and the contents of assertion. In section 1, I describe several motivations for context probabilism put forward in recent discussions of epistemic modals and other epistemic vocabulary. In section 2, I develop an account of belief that naturally emerges from this thesis. Adopting context probabilism is not necessary for endorsing my account, though it does carry some additional advantages. In short, the combination of context probabilism and my account of belief is more powerful than the sum of its parts.

1. A Brief Introduction to Context Probabilism

Suppose that Smith and Brown are buying a birthday gift for their friend Jones. Among other gifts, they are thinking about buying her a ticket on a nonsmoking cruise. As they discuss whether to buy the ticket, Smith might express relevant opinions using any of the following sentences:

(2) Jones might smoke.
(3) Jones probably smokes.
(4) It's extremely likely that Jones smokes.
(5) Jones is more likely to smoke than drink.

What opinions do these sentences convey? According to traditional theories of epistemic vocabulary, Smith uses each of these sentences to assert a proposition—namely, a claim about some contextually relevant body of evidence.[3] For example, Smith could use (2) to assert that Jones smoking is compatible with all the evidence at his disposal. The state of the conversation between Smith and Brown can be modeled by a set of worlds—namely, all those worlds compatible with the propositions that Smith and Brown have asserted so far.[4] When Smith utters one of the above sentences, this set of worlds is updated by intersection with the truth conditions of the proposition that Smith asserts.

---

3. See Kratzer 1991 for a canonical defense of this view, and see Swanson 2008 for an opinionated survey of relevant literature.
4. This model of conversational states is developed in Stalnaker 1978 as a formal model for the notion of common ground in Grice 1967.

This traditional picture has recently met with a number of serious objections. For instance, some argue that truth-conditional theories fail to explain our intuitive judgments about eavesdroppers' assessments of assertions made using epistemic modals.[5] Such theories also seem to have trouble explaining our judgments about retractions of assertions made using epistemic modals.[6] According to Yalcin 2007, truth-conditional theories fail to explain our judgments about epistemic modals embedded under supposition operators. And these same objections challenge not only truth-conditional theories of epistemic modals, but also truth-conditional theories of probability operators in sentences like (3) – (5).[7] In short, truth-conditional theories face a variety of challenges when it comes to accounting for ordinary language judgments about epistemic vocabulary.[8]

In addition to ordinary-language arguments against truth-conditional theories, there are also theoretical reasons to prefer an alternative approach. On traditional theories of assertion, speakers directly share their beliefs with each other when they communicate—namely, by asserting the contents of those beliefs—but speakers can only ever come to share their *credences* with each other in a roundabout fashion. Say I have high credence that Jones smokes, and I want you to share my high credence. At best, I can express a belief in some surrogate proposition, where this belief plays a functional role similar to that of my high credence—such as, say, the belief that it is likely, given my evidence, that Jones smokes. As a result, you might end up sharing my high credence that Jones smokes, but only as an indirect effect of updating on the proposition that I assert. This marks an unfortunate and unmotivated distinction between how we convey our beliefs and credences in conversation. As Forrest puts it, "To be able to express a high degree of belief rather than merely express a belief is so useful an ability that we should be most surprised if we had no way of expressing a high degree of

---

5. For instance, see Egan et al. 2005, Hawthorne 2007, and von Fintel and Gillies 2008.

6. For instance, see MacFarlane 2011 and Yalcin and Knobe 2014.

7. For further discussion of probability operators, see section 2.1 of Moss 2018.

8. These ordinary language arguments have been challenged by a number of authors, including Barnett (2009), Sorensen (2009), von Fintel and Gillies (2011), Dowell (2011), and Dorr and Hawthorne (2012). In light of the controversial nature of the relevant ordinary language judgments, one might reasonably prefer to motivate context probabilism on more theoretical grounds, as discussed below.

belief" (1981: 44).[9] In Moss 2018, I develop this line of argument in much greater detail, spelling out three foundational objections to using truth-conditional theories to model the communication of probabilistic beliefs.[10]

Faced with these mounting challenges for truth-conditional theories, it is worth exploring alternative models of communication between subjects with probabilistic beliefs. An especially promising suggestion involves adding probabilistic structure to the conversational common ground. This is the central idea of *context probabilism* as defined in Yalcin 2012b. According to context probabilism, the state of a conversation is not represented by a set of worlds, but rather by a set of *probability spaces*, objects that assign precise probability values to propositions.[11] Similarly, the content of an assertion is not a set of worlds, but a set of probability spaces. At a first pass, Smith may use 'Jones probably smokes' to directly assert the set of probability spaces that assign high probability to Jones smoking. 'Jones is more likely to smoke than drink' may be used to assert the set of probability spaces that assign higher probability to Jones smoking than Jones drinking. This probabilistic theory can be extended to assertions made using epistemic modals and other epistemic expressions.

This probabilistic theory of the contents of assertion can also be extended to the contents of belief. As I argue in Moss 2018, credences should be understood as beliefs in probabilistic contents, where a *probabilistic content* is not a proposition about probabilities, but rather a set of probability spaces. At a first pass, you believe a set of probability spaces just in case your precise credence function is contained in that set. For instance, you might believe the probabilistic content *that Jones probably smokes* in virtue of having high credence that Jones smokes. As a result, my preferred view entails that you can assert the very same contents that you believe. When you have high credence that Jones smokes, you believe a probabilistic content, you can assert it, and your audience can come to believe it. We do not share our credences with each other by expressing full beliefs in functionally equivalent surrogate propositions. The way we communicate our credences is just as direct as the way we traditionally communicate our full beliefs.

9. For a similar line of thought, see the final paragraph of Yalcin 2012a.

10. See section 2.2 of Moss 2018, "Foundational Arguments for Probabilistic Contents of Assertion."

11. Formally, a probability space $S = \langle \Omega_S, \mathcal{F}_S, m_S \rangle$ is an ordered triple consisting of a set Ω_S of possible worlds, an algebra $\mathcal{F}_S$ over Ω_S, and a probability measure m_S on $\mathcal{F}_S$.

There is much more to be said in defense of context probabilism, in addition to the motivations mentioned here. But for the purposes of this paper, I am going to set aside further debate about probabilistic theories of context and content. This paper aims to explore consequences of these theories for our thinking about the attitude of full belief.

2. Simple Sentences as Loose Probabilistic Speech

According to context probabilism, even a simple sentence such as (6) has a set of probability spaces as its content.[12]

    (6)   Jones smokes.

This raises an important question: Exactly what probabilistic content does this simple sentence convey? Here is a straightforward proposal: the content of (6) is just the set of probability spaces according to which it is certain that Jones smokes—that is, probability spaces such that Jones smokes in every world in their domain. Unfortunately, this proposal faces an obvious problem. This set of probability spaces seems much too strong to be the content conveyed by (6). For instance, this set contains only probability spaces that assign probability 1 to the proposition that Jones smokes. But subjects routinely assert and accept simple sentences while having some credence in the negation of their traditional propositional contents.

      In order to answer this problem for the straightforward proposal, it is helpful to compare the problem with another more familiar one. Suppose that we are back in the business of assigning traditional truth-conditional contents to sentences, and we are trying to decide what content is conveyed by the following sentence:

    (7)   Jones arrived at the party at 3:00.

In other words, assuming that the state of our conversation is modeled by a set of possible worlds, which worlds remain in that set after someone says (7)? Here is a straightforward proposal: the content of (7) is just the set of worlds where Jones arrived at the party at *exactly* 3:00. But this set of worlds seems much too strong to be the content conveyed by (7). For instance, this content contains no world at which Jones arrived even one split

---

    12. This claim has been generally endorsed by advocates of context probabilism, including Yalcin (2012a), Moss (2015), Swanson (2016), and Rudin (2018). For further discussion, see section 3.5 of Moss 2018.

second later than 3:00. But subjects routinely assert and accept sentences like (7) without having such strong beliefs about the timing of events.

This familiar problem about the interpretation of sentences like (7) admits of several familiar solutions. According to Lewis 1980, for instance, the semantic content of a sentence like (7) is often weaker than the straightforward proposal suggests. Sentences like (7) are context sensitive. As uttered in many ordinary contexts, (7) denotes a content that is true even at worlds where Jones arrived at 2:58. But in a context where we are setting our watches and every second matters, (7) may denote a content that is false when Jones arrived at 2:58. Alternatively, one might take the interpretation of loose speech to be a pragmatic phenomenon. According to pragmatic accounts, you can use (7) to assert the strong content that Jones arrived at exactly 3:00, while conveying another weaker content to your audience.[13]

For ease of exposition, I am going to assume that Lewis 1980 is correct. Although my account of belief does not ultimately depend on this assumption, I take it that instances of loose speech are context sensitive, and that their literal contents coincide with the contents that they are used to convey. For the purposes of this paper, what matters is the following thesis: *Simple sentences are instances of loose probabilistic speech.* In just the same way that speakers use 'Jones arrived at 3:00' to say that Jones arrived sometime fairly close to 3:00, without saying that Jones arrived at exactly 3:00, speakers use 'Jones smokes' to say that it is fairly close to certain that Jones smokes, without saying that it is absolutely certain that Jones smokes. The content conveyed by 'Jones smokes' is a set of probability spaces, but it is not the set of probability spaces according to which it is absolutely certain that Jones smokes. As uttered in ordinary contexts, the content of 'Jones smokes' contains many probability spaces that merely assign high probability to the proposition that Jones smokes, just as the traditional propositional content of 'Jones arrived at 3:00' contains many worlds at which Jones arrived merely within a few minutes of 3:00.

13. For additional semantic theories of loose speech, see Sauerland and Stateva 2007, Krifka 2007, and Solt 2014. For pragmatic theories, see Lasersohn 1999, Lauer 2012, and Klecha 2018. As noted in section 7 of MacFarlane 2003, Lewis himself seems to change his mind about whether standards of precision affect literal truth conditions; Lewis 1979 can be read as neutral about or even sympathetic with pragmatic theories of loose speech.

This account of simple sentences is supported by extensive similarities between simple sentences and paradigmatic instances of loose speech. For starters, loose speech commonly occurs against a background of more and less precise speech. Sentences that report the time can be arbitrarily precise:

(8)   It's 2:59 / 2:59:59 / 2:59:58.96.

The same goes for sentences with probabilistic contents. For instance, sentences containing probability operators can be arbitrarily precise:

(9)   It's .9 likely / .99 likely / .9883 likely that Jones smokes.

As the sentences in (8) and (9) become more precise, they might become more useful in extremely technical conversations—about the timing of a rocket launch, say, or about potential causes of an apparent lung disease. Conversely, in many ordinary contexts, there is no need for the precision exhibited by any of the above sentences. For practical purposes, one might just as well report the time to the nearest minute or hour. In the case of probabilistic contents, one might just as well round some content to the nearest certainty. In fact, it often violates the maxim of relevance to assert more precise contents when the added precision does not matter for any purposes at hand. To make matters worse, speakers may not have very precise beliefs about the time, in which case they should stick to using expressions that are less precise than those in (8). Similarly, speakers may not have very precise credences when using epistemic vocabulary. It is therefore predictable, and in fact essential, that the practice of speaking loosely should accompany our ability to express precise contents, including precise probabilistic contents.

Another hallmark of loose speech is that it can be modified by *slack regulators* in the sense of Lasersohn 1999—namely, linguistic devices that raise the contextual standards for interpreting imprecise expressions. For example, the following sentences intuitively convey something stronger than the content that it is fairly close to 3:00:

(10)   a.   It's exactly 3:00.
       b.   It's precisely 3:00.
       c.   It's 3:00 on the dot.

Similarly, various operators raise the standards of precision for our interpretation of simple sentences:

> (11) a. Jones absolutely smokes.
>   b. Jones totally smokes.[14]
>   c. It's certain that Jones smokes.

In fact, some expressions have interpretations both as slack regulators of simple sentences and also as slack regulators of more familiar instances of loose speech. For instance, 'absolutely' and 'totally' can raise the standards for our interpretation of gradable adjectives:

> (12) a. That is absolutely flat.
>   b. That is totally flat.

Just as speakers use the sentences in (11) to express an especially high confidence that Jones smokes, speakers use the sentences in (12) to express their belief that an object meets an especially high standard of flatness.

Another shared feature of simple sentences and other loose speech is that in addition to the use of explicit slack regulators, the use of precise speech often raises the standards of precision at a context.[15] This fact helps explain why it often sounds bad to elaborate loose speech using more precise speech, as in the following conjunctions:

> (13) ??It's 3:00 and it's 3:01.
> (14) ??Jones smokes and it's .9 likely that she smokes.

A related phenomenon arises when speakers try to be explicit about the contents of their loose speech. The use of explicit hedging operators can raise the standards of precision, thereby strengthening the contents conveyed by unhedged sentences. This fact helps explain the infelicity of the following conjunctions:

> (15) ??It's 3:00 and it's very close to 3:00.
> (16) ??Jones smokes and it's very close to certain that Jones smokes.

Because the first conjuncts of these sentences are naturally contrasted with their second conjuncts, the former conjuncts are naturally interpreted relative to raised standards of precision. Hence their contents strictly entail the contents of the second conjuncts of these sentences, which

---

14. See Beltrama 2018 for a detailed account of our use of 'totally' as an intensifier that can express "heightened confidence" (234).

15. A referee has encouraged me to emphasize that this is not a universal claim. I am not defending any necessary or sufficient conditions for sentences like (13) and (14) to be felicitous, as the interpretation of loose speech is influenced by a wide variety of contextual features. See section 4 for further discussion.

helps explain why it sounds odd for a speaker to go on to directly express the latter contents.

To sum up, the classification of simple sentences as loose speech is supported by several analogies between simple sentences and paradigmatic instances of loose speech. Having classified simple sentences as loose probabilistic speech, we can answer the question for context probabilism posed at the start of this section. Simple sentences are used to convey different sets of probability spaces in different contexts, depending on the standards of precision relevant for their interpretation.

3. An Account of Belief

The classification of simple sentences as loose probabilistic speech has substantive consequences for the semantics of belief ascriptions, and ultimately for our understanding of the attitude of belief. To start, consider the following sentence:

(17)   Smith believes that Jones smokes.

As you utter (17) in a particular context, exactly what belief are you ascribing to Smith? In order to answer this question, we may compare (17) with belief ascriptions embedding more familiar instances of loose speech. For example:

(18)   Smith believes that it is 3:00.

As you utter (18) in a particular context, exactly what belief are you ascribing to Smith? At a first pass, (18) says roughly that Smith comes close enough to believing that it is exactly 3:00, given some contextually relevant purposes. In just the same sense, (17) says roughly that Smith comes close enough to believing it is certain that Jones smokes, given some contextually relevant purposes.[16]

Here is a more careful answer to our question: just as context determines standards for the interpretation of loose speech, it also determines standards for the interpretation of belief ascriptions embedding

---

16. As a referee points out, it would also be correct to say here that 'Smith believes that Jones smokes' conveys that *it is close enough to certain that* Smith comes close enough to believing that Jones smokes. For sake of clarity, I convey this probabilistic content using a simple belief ascription, rather than its explicit probabilistic counterpart. In other words, I use some loose speech in the context of my paper, rather than prefacing sentences with 'it is close enough to certain that' and raising the standards of precision for the context of my discussion.

loose speech. For any sentence containing loose speech, there is some maximally precise content that the sentence could in principle have. A belief ascription is true as uttered in a context just in case the subject has a belief that is not relevantly different from the maximally precise belief that the embedded sentence could be used to express. For example, (18) could be true as uttered at a context in virtue of Smith believing that it is exactly 2:59, or in virtue of Smith believing that it is within ten minutes of 3:00, as long as there are no contextually relevant differences between these beliefs and the belief that it is exactly 3:00. Similarly, (17) could be true as uttered at a context in virtue of Smith having .9 credence that Jones smokes, for instance, or having imprecise credences ranging from .93 to .98 in the proposition that Jones smokes, or simply having some fairly high credence that Jones smokes.

In order to specify the exact meaning of (17) at a context, then, one must specify the probabilistic beliefs that are not relevantly different from the attitude of believing that it is certain that Jones smokes, according to the standards of that context. At first, it might be tempting to suppose that these standards are determined by a simple error measurement. For instance, one might suppose that you believe a proposition just in case your credence in that proposition is above some contextually determined threshold.[17] But the standards of precision relevant for the interpretation of loose speech are rarely just that simple. By comparison, it might be similarly tempting to suppose that numerals in natural language are always interpreted as being precise to the nearest integer. But in fact, the interpretation of numerals is more complicated. Consider the following example:

(19) The train will arrive in thirty seconds.[18]

This sentence could naturally be used to describe a train that will arrive in 30.6 seconds, even though it would not be accurate to the nearest second in such a context. There is no simple rule determining how precisely we interpret numerals. At a given context, some numerals might be precise to the nearest integer, while others are not. By identifying simple sentences as loose probabilistic speech, my account explains why we should not expect to find any simple rule for defining full belief in terms of thresh-

---

17. This claim and variations on it are sometimes identified as the *Lockean thesis*, following Foley 1992.

18. This example is due to Krifka 2002. For further discussion of scales of granularity and loose speech, see also Krifka 2007 and Sauerland and Stateva 2007.

old credence, even allowing that the relevant threshold may vary accord-
ing to context.

So far I have mainly discussed the truth conditions of belief ascrip-
tions that embed simple sentences. But since belief is just the attitude
ascribed by such belief ascriptions, the foregoing discussion also provides
us with insights about belief. The belief that Jones smokes is the attitude
ascribed by (17), just as the belief that it is 3:00 is the attitude ascribed
by (18):

(17)   Smith believes that Jones smokes.
(18)   Smith believes that it is 3:00.

Asking *what it takes for Smith to believe that Jones smokes* is in many important
respects like asking *what it takes for Smith to believe that it is 3:00.* The remain-
der of this paper derives useful conclusions about the former question
from familiar observations about the latter. These conclusions depend on
a crucial feature that (17) and (18) have in common—namely, they both
depend for their interpretation on contextually determined standards
of precision, where these standards are familiar objects of study in the
literature on loose speech.

At this point, I should pause to note that without accepting context
probabilism, readers may still accept that belief ascriptions embedding
simple sentences are instances of loose speech, thereby enabling me to
use linguistic theories of loose speech to address problems in the epis-
temology of full belief.[19] For the purposes of making at least some illu-
minating connections between theories of loose speech and theories of
full belief, what is essential is the claim that 'believes that Jones smokes'
is loose speech. Fans of context probabilism find it natural to accept
this claim, because they accept that 'Jones smokes' is loose speech. But
opponents of context probabilism can also accept this claim—namely,
by accepting that 'believes' is loose speech. That is, opponents of con-
text probabilism can accept that the verb 'believes' strictly denotes the
property of being certain of a proposition—namely, the propositional
content of the prejacent of the belief ascription—while accepting that
simple belief ascriptions convey that subjects are merely close enough to
certain of such propositions. That being said, context probabilism does
strengthen my account of belief. By identifying simple sentences as loose
probabilistic speech, fans of context probabilism can embrace deeper

19. I am grateful to Ofra Magidor and Julia Staffel for encouraging me to clarify
this point.

analogies and stronger conclusions about belief, as noted throughout this paper. And of course, fans of context probabilism have an independent motivation for accepting my account of belief, since it is a natural consequence of the probabilistic theory of simple sentences developed on their behalf in section 2.

Before turning to discuss applications of my account of belief, it is worth noting one more significant argument in its favor—namely, that my account explains the many striking similarities between existing theories of loose speech and existing theories of full belief. For instance, it is often remarked that the contextual standards for the interpretation of loose speech generally depend on our practical interests. Lauer observes that a speaker will utter 'Mary arrived by 3:00' in situations where "acting as though one believes that 'Mary was here by three' is true is just the same as acting as though one believes that 'Mary was here shortly after three' is true" (2013: 101). As Lasersohn puts it, the contents of loose speech are, by definition, "close enough to the truth for practical purposes" (525). These descriptions of loose speech bear a striking resemblance to the familiar claim that full belief is a state that is close enough to certainty for practical purposes. For example, Wedgwood defines outright belief in p as "the state of being stably disposed to assign a practical credence of 1 to p, for all normal practical purposes" (2012: 321). Weatherson (2016) argues that to believe a proposition is roughly to be disposed to not change any attitude toward salient questions upon updating on that proposition, where salient questions include what bets you should accept.[20]

Another striking parallel between full belief and loose speech is that both help agents manage the cognitive load of reasoning with more precise contents. Just as it is often useful to round to the nearest certainty in conversation with others, it is also useful to round to the nearest certainty in thought. It takes more effort to reason with credences that are precise to several decimal places, just as it takes more effort to reason with temporal beliefs that are precise to the nearest millisecond. Van Der Henst et al. argue that this helps explain why we often speak loosely when asked for the time, saying that we are disposed to give "an answer from which hearers can derive the consequences they care about with minimal effort. A rounded answer is easier to process" (2002: 457). Compare this with the claim that "those who employ credences risk being overwhelmed . . . rather than just discarding the propositions that aren't

---

20. For additional accounts comparing full belief and practical certainty, see Ganson 2008, Fantl and McGrath 2010, and Locke 2014.

believed and focussing on those that are, they will have to keep track of all of them and their associated credences" (Holton 2014: 14).

Both in the case of credences and in the case of temporal beliefs, it might be impossible for ordinary subjects to reason using only precise contents. That is, ordinary subjects might have to engage in at least some reasoning using loose temporal expressions and simple sentences. The complexity of credences is sometimes used to motivate a much stronger conclusion, namely skepticism about whether ordinary subjects have any credences at all. For instance, after describing the complexity of credences, Holton concludes that ordinary subjects only ever reason with full beliefs: "We cannot form credences at all. The Bayesian approach is not an idealization of something we actually do. Instead, it is quite foreign to us" (15). This conclusion seems overly drastic in light of my account of full belief. After all, it is not as if we only ever form or reason with temporal beliefs that are rounded to the nearest hour. The correct conclusion is more modest: just as it is sometimes but not always useful to round to the nearest minute or hour, it is sometimes but not always useful to round to the nearest certainty. Just as with other more familiar forms of loose speech, the use of loose probabilistic speech has significant cognitive benefits, and the same goes for our use of full beliefs in reasoning.

4. Does Full Belief Require Maximal Confidence?

The account of belief in section 3 can answer several difficult questions, the first of which concerns the relationship between full belief, credence, and certainty. To spell out a familiar tension: on the one hand, there is an intuitive sense in which belief requires certainty, i.e., the elimination of every possibility that the belief is false. If you *fully believe* that Jones smokes, then you do not harbor doubts about whether she smokes, which explains why it sounds contradictory to say:

(20)   Jones smokes and Jones might not smoke.

Similarly, it sounds like an indictment of Smith to say:

(21)   Smith believes that Jones smokes and that Jones might not smoke.

These observations suggest that belief requires certainty, or at least that subjects must be maximally confident of what they believe.[21] As Greco

21. In the context of this paper, certainty is strictly stronger than credence 1 or maximal confidence. Suppose that you are throwing darts, and that your next dart is

puts it, "Binary belief is maximal degree of belief—it is the endpoint of the scale of degreed belief" (2015: 179).[22] But on the other hand, it seems that belief could not possibly require anything like credence 1 or certainty. There is an intuitive sense in which ordinary people believe plenty of propositions, while being more confident of some of these propositions than others. In fact, an ordinary person may not have credence 1 in any propositions at all, while still having plenty of beliefs. As Wedgwood puts it, "Non-trivial theories of outright belief face a dilemma . . . having an outright belief in p is a way of having full confidence in p—in other words, a way of treating p as certain. . . . Yet we all seem to have outright beliefs in propositions of which we are *not* maximally confident" (2012: 317).

In addressing this dilemma, it is helpful to consider related questions about other instances of loose speech. Does believing that it is 3:00 require believing that it is exactly 3:00? This question deserves different answers, depending on exactly what it is asking. On the one hand, there is a clear sense in which believing that it is 3:00 does not require believing that it is exactly 3:00. As mentioned in section 3, the belief ascription 'Smith believes that it is 3:00' can be true in virtue of Smith believing that it is exactly 2:58, and more generally, in virtue of Smith coming close enough to believing that it is exactly 3:00. In just this same sense, 'Smith believes that Jones smokes' can be true in virtue of Smith having .98 credence that Jones smokes, or more generally, credences that are relevantly close enough to certainty.

On the other hand, there are other ways of interpreting the question of whether full belief requires certainty. In some contexts, the standards of precision for interpreting (22) are high enough that it has just the same content as (23):

(22)   Jones smokes.
(23)   It is certain that Jones smokes.

In such contexts, (24) may have just the same truth conditions as (25):

---

equally likely to hit uncountably many points on the dartboard, including its point-sized bullseye. Then you can be maximally confident that you will miss the bullseye, while still failing to be certain that you will miss it.

22. See Dodd 2017 for a survey of sympathetic literature, as well as an extended defense of the claim that belief requires credence 1. Clarke 2013 also defends this claim, though only against the background of an unorthodox analysis of credences.

(24)   Smith believes that Jones smokes.
(25)   Smith is certain that Jones smokes.

Furthermore, these high-standards contexts may include the very contexts in which we are asking whether belief requires certainty, or whether the attitude of belief is compatible with doubt, or whether one can fully believe something without being maximally confident of it.

Recall from section 2 that using precise speech often raises the standards of precision for the interpretation of loose speech at a context. In just this sense, discussing possibilities of error may raise the standards of precision for the interpretation of belief ascriptions embedding simple sentences. This explains the striking fact that there are conversations in which one could truly utter either (26) or (27) without being able to truly utter (21):

(26)   Smith believes that Jones smokes.
(27)   Smith believes that Jones might not smoke.
(21)   #Smith believes that Jones smokes and that Jones might not smoke.

This fact deserves the same explanation as the fact that there are conversations in which one could truly utter either (28) or (29), but not (30):

(28)   Smith believes that it is 3:00.
(29)   Smith believes that it is one minute after 3:00.
(30)   #Smith believes that it is 3:00 and that it is one minute after 3:00.

The more precise speech in (27) and (29) naturally raises the standard of precision for the interpretation of (26) and (28). As a result, the latter ascriptions will not be true as uttered in any context in which the former ascriptions are both true and uttered. This observation helps us make sense of claims to the effect that belief is incompatible with doubt. In the high-standards contexts in which such claims are made, belief ascriptions and doubt ascriptions may indeed have incompatible truth conditions.

Another useful observation is that certain expressions for belief are themselves slack regulators, including expressions often used in philosophical contexts:

(31)   a.   Smith has an *outright / all-out / full* belief that Jones smokes.
        b.   Smith *fully* believes that Jones smokes.

These expressions for belief raise the standards of precision for the interpretation of their complements. Here is a comparable example:

(32)   Smith has painstakingly calculated that Jones arrived at 3:00.

As uttered in an ordinary context, the sentence 'Jones arrived at 3:00' might convey merely that Jones arrived fairly close to 3:00. But in the context of (32), it can mean something stronger, since 'painstakingly calculated' encourages us to interpret its complement relative to a higher standard of precision. This same sort of context shifting explains the fact that (33) sounds like a contradiction, even though (1) sounds fine:

> (33) ??I *fully believe* it will rain this afternoon, but I'm not *sure* that it will.
> (1) I *believe* it will rain this afternoon, but I'm not *sure* that it will.

As compared with the first conjunct of (1), the first conjunct of (33) as-cribes belief in a stronger probabilistic content, namely a content that is incompatible with the content of the belief ascribed by the second con-junct of these sentences. To sum up, the term for belief that we use in a belief ascription can itself help determine which probabilistic content we are talking about someone believing. Just as it is true in some but not all contexts to say that believing it is 3:00 requires believing it is exactly 3:00, it is true in some but not all contexts to say that belief requires certainty.

   In order to clarify my account of simple belief ascriptions as loose speech, let me compare it with an alternative account defended by Hawthorne et al. (2015). According to Hawthorne et al., 'believes' and 'is sure that' have different strict semantic contents. On their account, (1) sounds fine for the same reason that (34) sounds fine:

> (1) I believe it will rain this afternoon, but I'm not sure that it will.
> (34) Most of the students failed the exam, but not all of them did.

As uttered in any context, the conjuncts of these sentences have consis-tent contents. By contrast, on my account, 'believes' and 'is sure that' ascriptions have the same strict contents, but they can be used to convey different loose contents. In many contexts, and especially in contexts that contain both sorts of ascriptions, 'is sure that' ascriptions are generally interpreted relative to a higher standard of precision. According to my account, the difference between 'believes' and 'is sure that' is like the difference between the expression 'empty' and 'vacuum' according to theories that classify 'empty' as loose speech. As Kennedy and McNally explain, "We can maintain our claim that the actual denotation of the predicate headed by *empty* is a property that is true only of objects that are completely empty, but that its pragmatic halo includes properties that are true of objects that are just a little bit less than empty" (2005: 357). In other words, 'empty' and 'vacuum' have the same strict content, but the more precise term 'vacuum' is generally interpreted relative to a higher

standard of precision. As a result, the following sentence can be used to convey a consistent content:

(35)   That space is empty, but it's not a vacuum.

According to my loose speech account of belief ascriptions, (1) sounds fine for the same reason that (35) sounds fine—namely, because the second conjunct is naturally interpreted relative to a higher standard of precision than the first.[23]

The conjunctions (1) and (35) are exceptions to the observation that precise expressions tend to raise the standards of precision at a context, even for interpreting loose speech that appears earlier in the same sentence. As mentioned in section 3, our interpretation of loose speech depends on many subtle features of context. In the case of (1) and (35), an important feature is that these conjunctions are naturally read with a particular intonational contour. 'Believe' and 'sure' naturally receive contrastive pitch accents in (1), and 'empty' and 'vacuum' receive the same accents in (35). This encourages the listener to interpret the relevant standard of precision as shifting between their conjuncts, so that the second conjunct merely conveys the negation of the strict content of the first, as opposed to the negation of its loose content. Like the theory defended by Hawthorne et al. (2015), then, my theory can account for the fact that sentences such as (1) are felicitous. At the same time, my theory can account for the existence of high-standards contexts—including some epistemology classrooms—in which the denotations of 'believe' and 'sure' coincide.

5. Is There a Distinction between Strong and Weak Belief?

In recent philosophical papers about belief, it is often assumed that the attitude under discussion is distinct from the subject of many ordinary attitude ascriptions. According to Greco, for instance, "If the claim that belief involves maximal confidence is to be worth taking seriously at all, we cannot be working with a conception of belief closely tied to natural language constructions involving 'belief' and 'believe'" (2015: 180). Ichikawa says that "the ordinary language most people use to talk about doxastic states is messy, and many apparent belief ascriptions discuss something weaker than outright belief" (2017: 180). Leitgeb (2013) dis-

---

23. Similarly, one can say that an object is *still* but not *motionless*, *dried* but not *desiccated*, *quiet* but not *silent*, and so on.

tinguishes the "commonsensical notion of belief" from the "flat-out belief" discussed by epistemologists, and he notes that some theorists distinguish these notions by using 'acceptance' for the latter. According to these authors and many others, there are two very different sorts of belief states. There is a state of weak belief that is discussed in ordinary contexts, which obviously does not require anything like certainty. In addition, there is a state of strong belief discussed by epistemologists, which is such that it is controversial whether it requires anything like certainty. The latter notion of strong belief is often said to be a technical or theoretical notion. As Hawthorne et al. put it, "There may be a theoretical notion of outright or full belief that is strong, and that shares the same evidential norms as assertion, but we argue that this does not correspond to our basic concept of belief" (2015: 1395). To sum up, an increasingly trendy view among contemporary epistemologists is that philosophers and ordinary speakers use 'believes' to express distinct concepts, and that these concepts pick out fundamentally different mental states.

There are several reasons to be suspicious of this trendy view. The philosophical discussion of belief has a long history, and it is far from obvious that this discussion is the product of any coordinated effort to introduce and employ technical vocabulary for an attitude that ordinary speakers do not talk much about. Even among contemporary authors, the philosophical notion of outright belief is sometimes taken to coincide with our ordinary folk notion of belief. For instance, Hawthorne et al. observe that in discussions of whether facts about full belief can be reduced to facts about credence, "It is generally taken for granted that the notion of outright belief is the commonsense one" (2015: 1402). Buchak begins her discussion of full belief and credence by saying: "Full belief (hereafter, just "belief") is a familiar attitude: it is the attitude that the folk talk about, and it has been a subject of epistemology since epistemology began" (2013: 285).

From these observations, some uncomfortable consequences for the trendy view start to emerge. First, when some epistemologists claim that their 'belief' is not a technical term, they are giving us at least some evidence for that claim. Advocates of the trendy view must say either that these epistemologists are wrong about what attitude they are talking about, or that they are talking past epistemologists who use 'belief' in a technical sense. Second, philosophers using the alleged theoretical concept of full belief cannot simply assume that they have succeeded in determining a referent for that concept, especially in the absence of

conventional implicit definitions for it. Finally, the combination of these problems is especially severe. Given that there is not even agreement among philosophers about whether the notion of full belief is a technical notion or a commonsense one, it is harder to imagine that our philosophical discussions of belief still manage to target a unique theoretical notion that is distinct from our ordinary notion of belief.

Fortunately, advocates of context probabilism can reject the trendy view and its uncomfortable consequences. There is no fundamental distinction between different belief concepts, or between different states picked out by those concepts. Belief ascriptions embedding simple sentences are just like belief ascriptions embedding more familiar loose speech. There is not a strong way and a weak way of believing that it is 3:00. The term 'believes' in (36) is not ambiguous between a folk notion and a technical notion of belief:

(36)   Smith believes that it is 3:00.

Rather, there are some contexts in which (36) and (37) have the same truth conditions, and other contexts in which they do not:

(37)   Smith believes that it is exactly 3:00.

For just the same reason, there are contexts where the following sentences have the same truth conditions, and other contexts in which they do not:

(38)   Smith believes that Jones smokes.
(39)   a.   Smith believes that it is certain that Jones smokes.
       b.   Smith is certain that Jones smokes.

What changes from context to context is not the strength of the belief state that is being ascribed, but the strength of the content of that state. Without positing multiple sorts of belief states, we can account for ordinary readings of (36) that are intuitively compatible with Smith believing that it is 2:59, as well as stronger readings that are not compatible with it. Similarly, without positing multiple sorts of belief states, we can account for ordinary readings of (38) that are intuitively compatible with Smith having some doubts about whether Jones smokes, as well as stronger readings that are not compatible with it. There is no special technical notion of belief that is discussed only in philosophical contexts.

In addition to providing an alternative to the trendy view, my account of belief has another surprising advantage—namely, it can reconcile conflicting intuitions about the minimum credence required for

belief. On the one hand, it might seem that to believe something, you must have more than .5 credence in it. This condition on belief is explicitly assumed by several authors, and some recent accounts of belief are tailored to entail it.[24] On the other hand, ordinary speakers sometimes find it perfectly natural to say that subjects believe propositions in which they have less than .5 credence. Suppose that Smith and Jones are examining a jar of jellybeans and guessing how many jellybeans it contains. If Smith guesses that there are 354 jellybeans, there are contexts in which it can sound perfectly fine to say:

> (40)   Smith believes that there are 354 jellybeans in the jar.

It is fair to assume that Smith has less than .5 credence that his guess is correct. Are we therefore required to conclude that (40) is strictly speaking false? At first glance, it appears that if (40) is true, then it must ascribe some mental state that is far from anything of interest to epistemologists as they debate whether belief requires knowledge, say, or whether rational belief is closed under conjunction.

The account of belief in section 3 can respect our intuition that (40) has a true reading, while still acknowledging the unity of folk and theoretical notions of belief. Here again, it is useful to compare simple sentences with more familiar instances of loose speech. Expressions of loose speech often have default standards of interpretation. For instance, sentences containing integers are commonly interpreted as denoting contents that are precise *at least to the nearest integer*. However, as explained in section 3, this default generalization has plenty of exceptions. The same goes for our interpretation of simple sentences. The default standard is that simple sentences express beliefs in contents that are precise *at least to the nearest certainty*. Accordingly, belief ascriptions embedding simple sentences are used to ascribe beliefs in contents that are precise at least to the nearest certainty. For example, (17) is used to say that Smith has a belief that is close enough to certainty that Jones smokes, at least insofar as both entail that it is more than .5 likely that Jones smokes:

> (17)   Smith believes that Jones smokes.

This accounts for our inclination to say that belief requires having more than .5 credence. However, just as for numerals, there are some contexts in which the default standards of interpretation for simple sentences

---

24. For sympathetic discussion of this condition, see Weatherson 2005, Leitgeb 2014, and Staffel 2016. For critical discussion, see section 6.2.4 of Maher 1993.

are not applied. For example, (40) is used to say that Smith has a belief that is relevantly similar to the state of being certain that there are 354 jellybeans:

(40)   Smith believes that there are 354 jellybeans in the jar.

The similarity in question does not guarantee that both beliefs entail that it is more than .5 likely that there are 354 jellybeans. Rather, the relevant similarity is that both beliefs entail that among all possible exact guesses about how many jellybeans are in the jar, the hypothesis that there are 354 jellybeans is the most likely.

To put it another way, it often matters that the subject of a belief ascription has more than .5 credence in the traditional propositional content of the complement of the ascription. Having a belief that is close enough to certainty in a proposition often entails having a belief that recommends guessing that the proposition is true, if asked whether it is true or false.[25] But sometimes, having a belief that is close enough to certainty just means having a belief that recommends guessing that the proposition is true, if it appears as one among many options on a multiple-choice test. Hence the attitude of believing a proposition does not essentially involve having greater than .5 credence in it. It is possible to give a charitable theory of the full range of belief ascriptions found in ordinary language and still accept that ordinary speakers and philosophers are interested in the very same notion of belief.

A brief dialectical note: Some but not all of the arguments in this section are available to advocates of the view that simple belief ascriptions are loose speech because 'believes' is an instance of loose speech. Unlike my context-probabilist account, this alternative view distinguishes the strict denotation of 'believes' from the weaker attitudes that this verb is used to ascribe. Nevertheless, the view does share some advantages with my preferred account. For instance, the view explains the varying strength of belief ascriptions without saying that epistemologists and ordinary speakers have entirely different concepts of belief. In addition, the view provides a compelling account of our conflicting intuitions

25. As explained in Horowitz 2019, your credences can get things right or wrong by licensing true or false educated guesses. This helps address an objection sometimes raised for reductive theories of belief in terms of credence—namely, that such theories fail to account for the fact that false beliefs are in some sense *incorrect* (cf. Fantl and McGrath 2010: 141).

about the minimum credence required for belief, since it also predicts that our interpretation of 'believes' should be subject to default standards of precision.

6. Are Rational Beliefs Closed under Entailment?

The preface paradox introduced in Makinson 1965 presents a classic dilemma for philosophical accounts of belief. Suppose that you have just written a book consisting of many carefully researched claims. If the book is long enough, it seems that you may rationally believe that at least one claim in it is false. Accordingly, you may refrain from believing that every claim in the book is true. In response to this example, many authors have concluded that the attitude of rational belief is not closed under logical consequence. Although you believe each claim in the book, it seems that you may rationally refrain from believing their conjunction.[26] On the other hand, it can seem almost like a truism that you should not go around believing contents that you recognize to be inconsistent. If you believe some contents and you see that they entail some consequence, then there is some intuitive sense in which you should either accept that consequence or revise your initial beliefs. These conflicting intuitions are a recurring theme in discussions of belief. As Frankish puts it, "Flat-out belief does seem to be subject to a tension. . . . We do feel both that we ought to adhere to conjunctive closure and that it is sometimes acceptable to violate it" (2009: 91). Hawthorne (2009) argues that "we are faced with two competing paradigms concerning the nature of rational belief" (50). An account of belief should resolve this tension—first by saying whether rational belief is closed under entailment, and then by explaining why each answer to this question seems compelling.

The preface paradox is traditionally formulated using belief ascriptions: it is rational to believe each sentence of the book is true, but not rational to believe the conjunction of these claims. But fans of context probabilism may usefully consider a first-order variation of the paradox that concerns the preface book itself, as opposed to claims about what rational subjects should believe about the book. Suppose that you have the preface book in front of you. Consider the following claims:

26. For some classic arguments in favor of this response to the preface paradox, see Klein 1985, Foley 1993, Kyburg 1997, Christensen 2004, and Sturgeon 2010.

(41) a. The first sentence of this book is true.
      b. The second sentence of this book is true.

          ...

          The final sentence of this book is true.
          Some sentence of this book is false.

Each of these claims might seem correct when considered in isolation. But together they entail a contradiction. And so it is difficult to say exactly which you should believe. This difficult first-order choice gives rise to a difficult theoretical choice. Suppose you believe each of these compelling claims. As theorists, we appear to confront the difficult choice of either saying that you are irrational, or accepting that rational belief is not closed under entailment.

By endorsing the account of belief in section 3, the context probabilist eliminates these difficult choices. Here again, it is useful to compare our conflicting intuitions about simple sentences with our intuitions about more familiar forms of loose speech. Suppose that you are sitting in a movie theater, and you look around and think to yourself:

(42) a. There are three hundred people here.

Suppose that five minutes later, you see someone enter the theater, and you think:

(43) a. There are three hundred people here.
      b. There is one more person here than before.

Call this the *theater paradox*. Considered in isolation, each of the above claims might seem perfectly correct. But together they appear to entail a contradiction. Suppose you believe each of these claims. Again, we face a theoretical choice: either we must say that you are irrational, or we must say that rational belief is not closed under entailment.

Fortunately, the theater paradox has a clear solution. The foregoing theoretical choice is a false choice. In fact, as uttered in many ordinary contexts, the sentences of (42) and (43) convey perfectly consistent contents. Suppose again that we are back in the business of assigning traditional truth-conditional contents to sentences. There are many worlds where all of the contents loosely conveyed by (42) and (43) are true—namely, any world where there are about three hundred people in the theater both before and after you see someone arrive. It is rational for you to believe these consistent contents. An important feature of the theater paradox is that the paradox itself is not presented in an ordinary context. Rather, when we reflect on the paradox as a whole, it is natural to interpret

each sentence of the paradox relative to a high standard of precision, so that the contents of these sentences are indeed inconsistent. As Cummins observes, the use of relatively precise expressions in dialogues can create "an expectation of precision, or fine granularity, that persists" for the interpretation of other expressions (2015: 26). Cummins gives another similar example:

> A museum visitor . . . asks the curator how old a dinosaur skeleton is. 'It's 75 million and 32 years old', the curator replies. Astounded, the visitor asks how the curator can be so precise. 'Ah, well', the curator replies, 'when I started working here 32 years ago, they told me that it was 75 million years old.' (26)

Reflecting on mathematical or logical relations between contents is just the sort of exercise that naturally raises standards of precision. The museum dialogue and the theater paradox are both amusing in virtue of playfully ignoring the possibility of such context shifting.

Having resolved the theater paradox, we may return to address the first-order variation of the preface paradox introduced above. As uttered individually in ordinary contexts, the sentences of this paradox have consistent probabilistic contents. That is, the contents of these sentences contain many probability spaces in common—namely, any space according to which each sentence in the book is almost certainly true, while it is almost certain that some sentence is false. At the same time, the preface paradox itself is not presented in an ordinary context. In the context of (41), one should interpret each sentence of the paradox relative to a higher standard of precision, and conclude that these sentences are inconsistent.

The difficult theoretical choice presented by the preface paradox is a false choice. Fundamentally speaking, the preface paradox does not challenge the closure of rational belief any more than the theater paradox does. Rational belief requires consistency. Because the ordinary probabilistic contents of the premises of the preface paradox are consistent, you can rationally believe each of these contents. Insofar as the premises have inconsistent contents in the context of the paradox itself, you must reject at least one of these contents. At the same time, it is important to appreciate that contextual standards of precision are not determined by hard and fast rules. It is not hard to imagine someone asserting all three sentences of the theater paradox in tandem, intending to speak loosely and convey three consistent contents. In just the same sense, even theorists discussing the preface paradox may interpret its

premises relative to ordinary standards of precision, and correctly conclude that rational subjects can simultaneously believe each of the resulting contents at once.

7. Is Belief Interest Relative?

There has been a lot of recent debate among epistemologists over the thesis that belief is interest relative. But how exactly should we understand this thesis about belief? Some authors spell out the thesis as a simple dependence claim. According to Ross and Schroeder 2014, for instance, fans of interest relativity argue that "the level of confidence one must have in a proposition to count as believing it depends on pragmatic factors" (260). According to Weatherson, for instance, "Interests affect belief because whether someone believes p depends inter alia on whether their credence in p is high enough that any bet on p they actually face is a good bet" (2011: 592). This dependence claim resembles common characterizations of the interest relativity of knowledge, such as the claim that "whether or not someone knows that p may be determined in part by practical facts about the subject's environment" (Stanley 2005: 85).

Advocates of interest relativity theses sometimes go further in spelling out dependence claims of this sort, identifying the relevant notion of dependence as individual supervenience. For example, Stanley identifies his interest relativity thesis about knowledge as the claim that "there are cases in which two people are similarly situated, but one has knowledge, whereas the other does not, because one has greater practical investment in the truth or falsity of her beliefs" (2005: 2). In a similar spirit, Weatherson says that knowledge is interest relative just in case "there exist at least one pair of cases where the only relevant difference between agents in the two cases concerns their interests, but one knows that p and the other does not" (2011: 594). Analogously, we could say that belief is interest relative just in case there are subjects who have the same credences in every proposition, but where one subject has a belief and the other does not, and where this difference is entirely due to some difference in their practical interests.

In support of the thesis that belief is interest relative, Weatherson puts forward the following example:

> X and Y are parents of a child, Z.

> Y: This salad you bought is very good. Does it have nuts in it?
> X: No. The nuttiness you're tasting is probably from the beans.

> Y: Oh, so we could pack it for Z's lunch tomorrow.
> X: Hang on, I better check about the nuts. Z's pre-school is very fussy about nuts. One of the children there might have an allergy. (2016: 220)

According to Weatherson, *X* starts out this conversation believing that the salad that she bought is nut-free. But by the end of the conversation, she has doubts about whether it is nut-free, as evidenced by the fact that she is unwilling to pack it for *Z*'s lunch without gathering more information. This example is intended to vindicate both versions of the interest-relativity thesis stated above. As Weatherson sees it, whether you believe that a salad is nut-free depends partly on your practical interests, such as the potential cost of your belief being false. Throughout the above conversation, the subject *X* has just the same credences about whether the salad is nut-free. But first she believes that the salad is nut-free, and then she does not believe it, and this change in her belief state is entirely due to her changing practical interests.

In light of this sort of example, should we say that belief is interest relative? In short, it depends. There is an important sense in which belief is interest relative and an important sense in which it isn't. In order to spell out both of these claims, it is useful to consider belief ascriptions embedding other sorts of loose speech. But first, let us consider belief ascriptions embedding even more obvious examples of context-sensitive expressions. Suppose that Jones is on a road trip. Smith wants Jones to drive to Boston this afternoon, while Brown wants her to drive several hours farther. Both Smith and Brown know that Jones is currently about fifteen miles from Boston, and that she has five gallons of gas in her car. Both of the following belief ascriptions can sound fine as uttered in this context:

(44)   Smith believes that Jones has enough gas.
(45)   Brown does not believe that Jones has enough gas.

Accordingly, (46) has a natural reading on which it is true:

(46)   Whether someone believes Jones has enough gas depends on their interests.

In other words, (46) can be used to express the simple truth that whether or not someone believes that Jones has enough gas—that is, enough gas to get to wherever they want her to stop for the night—depends on where they want her to stop.

Similar readings arise for belief ascriptions embedding instances of loose speech. Suppose that Smith and Brown each believe that Jones

makes an annual salary of exactly $149,682. Brown is helping Jones with her income taxes, whereas Smith is helping some local Girl Scouts decide whether Jones could afford to buy a couple of boxes of cookies. Both of the following sentences could be used to describe what Smith and Brown believe:

> (47)   As he advises the Girl Scouts, Smith believes that Jones makes $150K.
>
> (48)   As he reads over tax forms, Brown does not believe that Jones makes $150K.

In the right sort of context, one might accurately describe this situation by saying:

> (49)   Whether someone believes Jones makes $150K depends on their interests.

In other words, (49) can be used to express the simple truth that whether or not someone has a belief state that is close enough to believing that Jones makes exactly $150K—that is, close enough that the difference is irrelevant given their practical interests—depends on what their practical interests are. The crucial observation here is that when it comes to interpreting loose speech embedded in the complement of a belief ascription, the relevant standard of precision can be determined by the derived context of the subject of the ascription, rather than by the global context in which the belief ascription is uttered.[27] Furthermore, in sentences such as (49), the relevant standards of precision can act just like the values of bound pronouns, with different standards corresponding to different believers quantified over by the sentence. The resulting interpretation of (49) is just like familiar bound-variable readings of sentences such as (46), discussed at length by Partee (1989), Stanley (2000), and others.

  These observations about the behavior of loose speech in belief ascriptions reveal a true interpretation of the thesis that belief is interest relative. Just like (46) and (49), the sentence (50) has a natural reading on which it is true:

---

27. This notion of *derived context* is closely related to the notion of *subordinate context* in Roberts 1989, and cousin to the notion of *local context* used in many dynamic semantic theories. For discussion of these notions of context and the relationship between them, see section 4.2 of Stalnaker 2014.

(50)   Whether someone believes the salad is nut-free depends on their interests.

This sentence can be used to say that whether or not someone has a belief state that is close enough to being certain that the salad is nut-free—that is, close enough that the difference is irrelevant given their practical interests—depends on what their practical interests are. This is an available interpretation of (50) because derived contexts can supply the standards of precision relevant for our interpretation of the loose speech in the complement of 'believes', just as in (49). And just as in (49), the relevant standards of precision can vary in tandem with the different subjects of the belief ascription.

In addition to dependence claims such as (50), supervenience claims such as (51) also have natural readings on which they are true:

(51)   There are subjects *A* and *B* such that *A* believes that a certain salad is nut-free, *B* does not, and the only relevant difference between *A* and *B* concerns their practical interests.

For comparison, consider the following claim:

(52)   Smith believes that Jones makes $150K, Brown does not, and the only relevant difference between Smith and Brown concerns their practical interests.

The true reading of this sentence corresponds to something like a sloppy reading of the pronoun that determines the standards of precision for interpreting the loose speech embedded under 'believes'. For just the same reason, (51) has a reading on which it has many true instances. For example, there is a true reading of the claim that the subject *X* starts out believing that a certain salad is nut-free, and then later fails to believe it once her practical interests have changed. To sum up so far, my account of simple sentences as loose speech predicts that belief ascriptions embedding simple sentences should have the same range of interpretations as other belief ascriptions embedding loose speech. As a result, my account predicts that there are true readings of many sentences commonly used to spell out the thesis that belief is interest relative.

At the same time, my account entails that there is another important sense in which belief is not interest relative. Fundamentally speaking, Smith and Brown have exactly the same beliefs about Jones's salary. We could elaborate on this claim by making the following observations:

(53)   Smith and Brown both believe that Jones makes exactly $149,682.

(54)   Smith and Brown both believe that Jones makes around $150K.

(55)   Neither Smith nor Brown believes that Jones makes exactly $150K.

The fact that Smith and Brown share all of these beliefs explains why Smith and Brown would act in exactly the same ways when making decisions related to Jones and her salary. Brown would have advised the Girl Scouts to sell Jones the cookies, for instance, and Smith would have advised Jones to report her exact income on her taxes. Fundamentally speaking, Smith and Brown differ not in their beliefs, but only in the decisions that they actually face.

The same conclusions hold for allegedly interest-relative beliefs. Fundamentally speaking, belief is a relation to probabilistic contents. In the case Weatherson describes, there is no probabilistic content such that whether *X* believes it depends on her practical interests. Hence fundamentally speaking, *X* has the same beliefs about the salad before and after her interests change. Her credence that the salad is nut-free is always close enough to certainty for the purpose of answering a curious question about its contents, and never close enough to certainty for the purpose of packing it for her preschooler. The fact that *X* maintains the same fundamental beliefs throughout the conversation explains why she is disposed to act in just the same ways throughout—why she would consistently be willing to eat the salad herself, for instance, and why she would not at any point be willing to pack it for her preschooler. This critical interpretation of the interest relativity thesis can be sharpened by considering belief ascriptions embedded under change-of-state verbs. Suppose that Smith starts to help Jones with her taxes. Then (47) may start to sound bad, and (56) may start to sound fine:

(47)   Smith believes that Jones makes $150K.

(56)   Smith does not believe that Jones makes $150K.

But it would be misleading to describe this case by saying:

(57)   Helping Jones with her taxes caused Smith to stop believing that Jones makes $150K.

Whether or not he is helping Jones with her taxes, Smith believes that she makes exactly $149,682. This fact is inconsistent with a natural interpretation of (57), on which it conveys that Smith changed his mind about exactly how much Jones makes. In just the same sense, it would be misleading to say that changes in practical interests cause subjects to gain or lose beliefs. For instance, it would be misleading to say:

(58)   Packing the salad for lunch caused *X* to stop believing that it was nut-free.

This sentence would be misleading in just the same way as (57). Whether or not it has a reading on which it is literally true, (58) naturally conveys the false content that packing the salad caused *X* to become less confident that it was nut-free.

The analogy between (57) and (58) vindicates one common objection to the thesis that belief is interest relative. Lutz (2013) develops this objection by observing that I do not lose my belief that the sun will rise tomorrow, simply in virtue of being offered a bet that will pay one penny if it does and cause me eternal torment if it doesn't:

> The problem with views like Weatherson's is that it entails that, in extreme circumstances like the penny-against-torment bet, I will no longer believe that the sun will rise tomorrow. But that seems wrong; this practical circumstance does not destroy my belief (1722–3).

Ross and Schroeder further develop this objection by arguing that interest relative theories of belief are inconsistent with the thesis of *Stability*, which says that a "fully rational agent does not change her beliefs purely in virtue of an evidentially irrelevant change in her credences or preferences" (2014: 277). As Ross and Schroeder explain:

> If a fully rational agent's ice cream preferences provide no evidence for or against the cat's being on the mat, then the agent shouldn't gain or lose the belief that the cat is on the mat purely in virtue of changing her ice cream preferences. (277)

Notice the change-of-state verbs used in these passages— *destroy, gain, lose.* As stated, these claims about belief are intuitively correct, for just the same reason that it is intuitively correct to say that Smith shouldn't lose the belief that Jones makes $150K purely in virtue of starting to help Jones with her taxes. There is no precise content such that Smith should lose his belief in it as a result of his changing practical interests.

At this point, fans of interest relativity might insist that the simple attitude of believing that Jones makes $150K plays an important functional role, where this role cannot be played by any of the precise beliefs that Smith and Brown have in common. After all, at a certain level of description, Smith has some dispositions that Brown does not have. For instance, only Smith is disposed to assertively utter sentences such as:

(59)   Jones makes $150K.

Furthermore, only Smith is disposed to act as if (59) were strictly speaking true. Suppose that as theorists, we want to explain why these very dispositions are shared by all and only certain subjects. It might seem as if our explanation must appeal to the fact that the subjects with these dispositions share a certain belief, namely a loose belief in the content that Jones makes $150K. Analogously, it might seem that full beliefs play an important functional role that cannot be played by credences. For instance, at the start of the salad conversation, *X* is disposed to assertively utter sentences such as:

> (60)   This salad is nut-free.

Furthermore, *X* is disposed to act as if the strict probabilistic content of this sentence were true. In order to explain why these very dispositions are shared by all and only certain subjects, one might be tempted to say that such subjects share a common loose belief in the content that the salad is nut-free.[28]

On reflection, however, we can explain why subjects share these various dispositions without saying that they also share some common loose belief. For instance, we can explain why Smith is disposed to assertively utter (59) by noting that in his circumstances, there is no practical difference between his precise beliefs about Jones and the precise belief that she makes exactly $150K. Analogously, we can explain why *X* initially acts as if it is certain that her salad is nut-free by noting that in her circumstances, there is no practical difference between her credences about the salad and the attitude of being certain that it is nut-free. In just this sense, fundamental explanations of action can proceed without any ascription of the loose belief that Jones makes $150K, and without the ascription of simple beliefs in propositional contents. This conclusion provides a precise and compelling interpretation of the claim that, fundamentally speaking, rational agents do not gain or lose beliefs purely as a result of some change in their practical interests.

As a reminder, this final sympathetic interpretation of *Stability* is compatible with certain interpretations of the thesis that belief is interest relative. Ross and Schroeder present *Stability* as a challenge for the thesis that "believing a proposition simply consists in having a sufficiently high level of confidence in it, but the level of confidence one must have in a proposition to count as believing it depends on pragmatic factors" (260).

---

28. For further development of this argument, see the "functionalist idea" defended in Weatherson 2005: 421.

But on behalf of their opponents, I have aimed to reconcile this thesis with *Stability*. Ross and Schroeder are correct that fundamentally speaking, you should not change your beliefs in virtue of changing your interests. At the same time, advocates of interest relativity can truly say that whether you believe that a certain salad is nut-free depends on your interests.

In this section, I have said that ascriptions of full belief can be true even when subjects believe only probabilistic contents. But earlier, I said that full beliefs help subjects manage the cognitive load of reasoning with more precise contents. Are these claims in tension? At first, it might appear that full beliefs can play their intended role in simplifying cognition only if they are distinct mental states from probabilistic beliefs, where the former mental states are cognitively simpler than the latter.[29]

At this point, it should come as no surprise that we can resolve this tension by reflecting on more general features of loose speech. It is common to use loose contents in thought, as when Smith thinks to himself, "Jones makes $150K." The use of lower standards of precision in thought helps Smith manage the cognitive load of reasoning with more precise contents about money. But the use of 'Jones makes $150K' in thought is not necessary for the truth of a belief ascription embedding this loose speech. 'Rudy the Robot believes that Jones makes $150K' can be true of a robot-like subject whose thoughts are entirely comprised of precise contents about Jones's salary. Similarly, the use of a simple sentence in thought can ground the truth of a simple belief ascription, but it is not necessary for the truth of that ascription. As noted in Weisberg 2018, robots can satisfy simple belief ascriptions in virtue of having precise probabilistic beliefs. Here again, as throughout this paper, our understanding of simple belief ascriptions is improved by reflecting on general facts about loose speech. The common themes of linguistic discussions of loose speech and epistemological discussions of belief are not mere surface similarities, but the result of a substantive connection that provides us with powerful insights about belief.

References

Barnett, David. 2009. "Yalcin on 'Might'." *Mind* 118, no. 471: 771–75.
Beltrama, Andrea. 2018. "*Totally* Between Subjectivity and Discourse: Exploring the Pragmatic Side of Intensification." *Journal of Semantics* 35, no. 2: 219–61.

Buchak, Lara. 2013. "Belief, Credence, and Norms." *Philosophical Studies* 169, no. 2: 285–311.

Christensen, David. 2004. *Putting Logic in Its Place: Formal Constraints on Rational Belief.* Oxford: Oxford University Press.

Clarke, Roger. 2013. "Belief Is Credence One (in Context)." *Philosophers' Imprint* 13, no. 11: 1–18.

Cummins, Chris. 2015. *Constraints on Numerical Expressions.* Oxford: Oxford University Press.

Dodd, Dylan. 2017. "Belief and Certainty." *Synthese* 194, no. 1: 4597–4621.

Dorr, Cian, and John Hawthorne. 2012. "Embedding Epistemic Modals." *Mind* 122, no. 488: 867–913.

Dowell, Janice. 2011. "A Flexible Contextualist Account of Epistemic Modals." *Philosophers' Imprint* 11, no. 14: 1–25.

Egan, Andy, John Hawthorne, and Brian Weatherson. 2005. "Epistemic Modals in Context." In *Contextualism in Philosophy: Knowledge, Meaning, and Truth,* edited by Gerhard Preyer and Georg Peter, 131–70. Oxford: Oxford University Press.

Egan, Andy, and Brian Weatherson, eds. 2011. *Epistemic Modality.* Oxford: Oxford University Press.

Fantl, Jeremy, and Matthew McGrath. 2010. *Knowledge in an Uncertain World.* Oxford: Oxford University Press.

von Fintel, Kai, and Anthony Gillies. 2008. "CIA Leaks." *Philosophical Review* 117, no. 1: 77–98.

von Fintel, Kai, and Anthony Gillies. 2011. "*Might* Made Right." In Egan and Weatherson 2011: 108–30.

Foley, Richard. 1992. "The Epistemology of Belief and the Epistemology of Degrees of Belief." *American Philosophical Quarterly* 29, no. 2: 111–24.

Foley, Richard. 1993. *Working Without a Net.* Oxford: Oxford University Press.

Forrest, Peter. 1981. "Probabilistic Modal Inferences." *Australasian Journal of Philosophy* 59, no. 1: 38–53.

Frankish, Keith. 2009. "Partial Belief and Flat-Out Belief." In Huber and Schmidt-Petri 2009, 75–93.

Ganson, Dorit. 2008. "Evidentialism and Pragmatic Constraints on Outright Belief." *Philosophical Studies* 139, no. 3: 441–58.

Greco, Daniel. 2015. "How I Learned to Stop Worrying and Love Probability 1." *Philosophical Perspectives* 29, no. 1: 179–201.

Grice, Paul. 1967. "Logic and Conversation." In *Studies in the Way of Words,* 22–41. Cambridge, MA: Harvard University Press.

Hawthorne, James. 2009. "The Lockean Thesis and the Logic of Belief." In Huber and Schmidt-Petri 2009, 49–74.

Hawthorne, John. 2007. "Eavesdroppers and Epistemic Modals." *Philosophical Issues* 17, no. 1: 92–101.

Hawthorne, John, Daniel Rothschild, and Levi Spectre. 2015. "Belief is Weak." *Philosophical Studies* 173, no. 5: 1393–404.

Holton, Richard. 2014. "Intention as a Model for Belief." In *Rational and Social Agency: Essays on the Philosophy of Michael Bratman*, edited by Manuel Vargas and Gideon Yaffe, 12–37. New York: Oxford University Press.

Horowitz, Sophie. 2019. "Accuracy and Educated Guesses." In *Oxford Studies in Epistemology* 6, edited by Tamar Szabó Gendler and John Hawthorne, 85–113. Oxford: Oxford University Press.

Huber, Franz, and Christoph Schmidt-Petri, eds. 2009. *Degrees of Belief.* Dordrecht: Springer.

Ichikawa, Jonathan. 2017. *Contextualising Knowledge: Epistemology and Semantics.* Oxford: Oxford University Press.

Kennedy, Christopher, and Louise McNally. 2005. "Scale Structure, Degree Modification, and the Semantics of Gradable Predicates." *Language* 81, no. 2: 345–81.

Klecha, Peter. 2018. "On Unidirectionality in Precisification." *Linguistics and Philosophy* 41, no. 1: 87–124.

Klein, Peter. 1985. "The Virtues of Inconsistency." *The Monist* 68, no. 1: 105–35.

Kratzer, Angelika. 1991. "Modality." In *Semantics: An International Handbook of Contemporary Research*, edited by Arnim von Stechow and Dieter Wunderlich, 639–50. Berlin: De Gruyter.

Krifka, Manfred. 2002. "Be Brief and Vague! And How Bidirectional Optimality Theory Allows for Verbosity and Precision." In *Sounds and Systems: Studies in Structure and Change. A Festschrift for Theo Vennemann*, edited by David Restle and Dietmar Zaefferer, 439–58. Berlin: De Gruyter Mouton.

Krifka, Manfred. 2007. "Approximate Interpretations of Number Words." In *Cognitive Foundations of Interpretation*, edited by G. Bouma, I. Kramer, and J. Zwarts, 111–26. Amsterdam: Koninklijke Nederlandse Akademie van Wetenschapen.

Kyburg, H. 1997. "The Rule of Adjunction and Reasonable Inference." *Journal of Philosophy* 94, no. 3: 109–25.

Lasersohn, Peter. 1999. "Pragmatic Halos." *Language* 75, no. 3: 522–51.

Lauer, Sven. 2012. "On the Pragmatics of Pragmatic Slack." *Proceedings of Sinn und Bedeutung* 16, no. 2: 389–401.

Lauer, Sven. 2013. "Towards a Dynamic Pragmatics." PhD thesis, Department of Linguistics, Stanford University.

Leitgeb, Hannes. 2013. "Reducing Belief Simpliciter to Degrees of Belief." *Annals of Pure and Applied Logic* 164, no. 12: 1338–89.

Leitgeb, Hannes. 2014. "The Stability Theory of Belief." *Philosophical Review* 123, no. 2: 131–71.

Lewis, David K. 1979. "Scorekeeping in a Language Game." In *Philosophical Papers*, vol. 1, 233–49. Oxford: Oxford University Press.

Lewis, David K. 1980. "Index, Context, and Content." In *Philosophy and Grammar,* edited by S. Kanger and S. Öhman, 79–100. Dordrecht: D. Reidel.

Locke, Dustin. 2014. "The Decision-Theoretic Lockean Thesis." *Inquiry* 57, no. 1: 28–54.

Lutz, Matt. 2013. "The Pragmatics of Pragmatic Encroachment." *Synthese* 191, no. 8: 1–24.

MacFarlane, John. 2003. "Three Grades of Truth Relativity." Unpublished manuscript, University of California, Berkeley. philosophy.berkeley.edu/macfarlane/work.shtml.

MacFarlane, John. 2011. "Epistemic Modals Are Assessment-Sensitive." In Egan and Weatherson 2011: 144–78.

Maher, Patrick. 1993. *Betting on Theories.* Cambridge: Cambridge University Press.

Makinson, D. C. 1965. "The Paradox of the Preface." *Analysis* 25, no. 6: 205–207.

Moss, Sarah. 2015. "On the Semantics and Pragmatics of Epistemic Vocabulary." *Semantics and Pragmatics* 8, no. 5: 1–81.

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

Partee, Barbara. 1989. "Binding Implicit Variables in Quantified Contexts." In *Compositionality in Formal Semantics,* 259–81. Oxford: Blackwell.

Roberts, Craige. 1989. "Modal Subordination and Pronominal Anaphora in Discourse." *Linguistics and Philosophy* 12, no. 6: 683–721.

Ross, Jake, and Mark Schroeder. 2014. "Belief, Credence, and Pragmatic Encroachment." *Philosophy and Phenomenological Research* 88, no. 2: 259–88.

Rudin, Deniz. 2018. "Uncertainty and Persistence: A Bayesian Update Semantics for Probabilistic Expressions." *Journal of Philosophical Logic* 47, no. 3: 365–405.

Sauerland, Uli, and Penka Stateva. 2007. "Scalar vs. Epistemic Vagueness: Evidence from Approximators." In *Proceedings of Semantics and Linguistic Theory (SALT) 17,* 228–45. Ithaca, NY: CLC Publications.

Solt, Stephanie. 2014. "An Alternative Theory of Imprecision." In *Proceedings of Semantics and Linguistic Theory (SALT) 24,* edited by Todd Snider, Sarah D'Antonio, and Mia Weigand, 514–33. Ithaca, NY: CLC Publications.

Sorensen, Roy. 2009. "Meta-Agnosticism: Higher Order Epistemic Possibility." *Mind* 118, no. 471: 777–84.

Staffel, Julia. 2016. "Beliefs, Buses, and Lotteries: Why Rational Belief Can't Be Stably High Credence." *Philosophical Studies* 173, no. 7: 1721–34.

Stalnaker, Robert. 1978. "Assertion." In *Context and Content,* 78–95. Oxford: Oxford University Press.

Stalnaker, Robert. 2014. *Context.* Oxford: Oxford University Press.

Stanley, Jason. 2000. "Context and Logical Form." *Linguistics and Philosophy* 23, no. 4: 391–434.

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

Sturgeon, Scott. 2010. "Confidence and Coarse-Grained Attitudes." In *Oxford Studies in Epistemology*, vol. 3, edited by Tamar Szabó Gendler and John Hawthorne, 126–49. Oxford: Oxford University Press.

Swanson, Eric. 2008. "Modality in Language." *Philosophy Compass* 3, no. 6: 1193–1207.

Swanson, Eric. 2016. "The Application of Constraint Semantics to the Language of Subjective Uncertainty." *Journal of Philosophical Logic* 45, no. 2: 121–46.

Van Der Henst, Jean-Baptiste, Laure Carles, and Dan Sperber. 2002. "Truthfulness and Relevance in Telling the Time." *Mind and Language* 17, no. 5: 457–66.

Weatherson, Brian. 2005. "Can We Do Without Pragmatic Encroachment?" *Philosophical Perspectives* 19, no. 1: 417–43.

Weatherson, Brian. 2011. "Defending Interest-Relative Invariantism." *Logos and Episteme* 2, no. 4: 591–609.

Weatherson, Brian. 2016. "Games, Beliefs and Credences." *Philosophy and Phenomenological Research* 92, no. 2: 209–36.

Wedgwood, Ralph. 2012. "Outright Belief." *Dialectica* 66, no. 3: 309–29.

Weisberg, Jonathan. 2018. "Belief in Psyontology." Unpublished manuscript, University of Toronto. Forthcoming in *Philosopher's Imprint*.

Yalcin, Seth. 2007. "Epistemic Modals." *Mind* 116, no. 464: 983–1026.

Yalcin, Seth. 2012a. "Bayesian Expressivism." *Proceedings of the Aristotelian Society* 112, no. 2: 123–60.

Yalcin, Seth. 2012b. "Context Probabilism." In *Logic, Language and Meaning*, edited by Maria Aloni, 12–21. Heidelberg: Springer.

Yalcin, Seth, and Joshua Knobe. 2014. "Epistemic Modals and Context: Experimental Data." *Semantics and Pragmatics* 7, no. 10: 1–21.