

Assertion, Belief, and Context

Roger Clarke

Abstract

This paper argues for a treatment of belief as essentially sensitive to certain features of context. The first part gives an argument that we must take belief to be context-sensitive in the same way that assertion is, if we are to preserve appealing principles tying belief to sincere assertion. In particular, whether an agent counts as believing that p in a context depends on the space of alternative possibilities the agent is considering in that context. One and the same doxastic state may amount to belief that p in one context but not another. The second part of the paper gives a formal treatment of doxastic states, according to which belief is context-sensitive along just these lines. The model is applied to characterize (but not to refute) skeptical arguments.

Introduction

One of the more active areas of epistemological research over the last few decades concerns the context-sensitivity of knowledge, or of “know”. In particular, knowledge (or “know”) is supposed to be sensitive to contextual facts about practical interests and salience of alternatives. Context-sensitivity, in one or another form, is supposed to illuminate such central philosophical issues as skepticism, practical reasoning, and warranted assertion.

I will argue here that belief is similarly context-sensitive. If it turns out that belief is sensitive to practical importance and salience of alternatives, this may affect the case for context-sensitivity of knowledge/“know”. In particular, our intuitions about cases of knowledge-loss across changes of context are often at the heart of the debate over contextualism (see, e.g., the Bank cases of DeRose 1992 and the Airport cases of Cohen 1999). If belief is sensitive to those contextual factors which vary across such cases, it can be expected to affect our intuitions about knowledge-loss across those cases.

A number of authors have explicitly denied that belief is context-sensitive (e.g., Foley 1993, p. 199, Kaplan 1996, p. 105n101, Maher 1986, p. 383n6),¹ but I shall argue in this paper that they are mistaken. In particular, I shall argue that whether an agent counts as believing that p depends, among other things, on the space of alternative possibilities the agent is taking seriously in the context at hand. This space may change without the agent having changed her

¹Each of these works is concerned with building a unified theory of outright belief and degrees of belief; context-insensitivity is in each case simply presented as a desideratum for such a theory. It's not hard to see a motivation for this: if we can only unify outright belief with degrees of belief by supposing belief comes and goes willy-nilly, our unification doesn't clearly achieve much. In this paper, I aim to explain how belief can be context-sensitive without being trivial, but my concern here is outright belief alone. I will not take on the task of giving a unified account of outright belief and degrees of belief, as Foley, Kaplan, and Maher did; but my (2013) does offer such a unified account, consistent with the view of outright belief advanced here.

mind or revised her beliefs at all: for example, a possibility previously ignored may come to be taken seriously once it is raised to conversational salience, say by explicit mention. For convenience, and to avoid confusion with other views about the sensitivity of this or that to one contextual factor or another, I shall call the view advanced here *sensitivism*. In particular, note that sensitivism differs from epistemic contextualism in part by being metaphysical rather than semantic, so to speak: whereas contextualists' central claims have to do with the word "know" and its cognates, sensitivism has to do with belief itself rather than belief reports.²

It might not be obvious what it would mean for belief to be context-sensitive, especially for readers primarily familiar with linguistic context-sensitivity. It's relatively easy to say what context-sensitivity means in linguistic cases: to say that "know" is context sensitive is to say that certain features of context figure in the truth conditions of sentences involving "know". I have already said that we are not concerned with "believe" or any other word that might be used to report belief, so this cannot be the sort of context-sensitivity at issue here. Beliefs, like sentences, do have truth conditions, but the sort of context-sensitivity at issue in this paper does not have to do with truth conditions *per se*. That is, the sensitivist's claim is not that a certain belief counts as true in one context but false in another. Rather, the claim is that the *content of a belief that p* varies across contexts: what it takes for one to count as believing that *p* varies across contexts. Let us make the parallel with linguistic context-sensitivity explicit.

Linguistic context-sensitivity: one and the same sentence may express the proposition *p* in one context but not another.

Sensitivism: one and the same doxastic state may amount to belief that *p* in one context but not another.

The picture that will emerge looks like this. To characterize an agent's doxastic state at a given time, it is not enough to produce an unstructured list of propositions believed. Rather, we need a list of propositions believed for each possible context, where contexts are individuated by the space of possibilities taken seriously. If a change in the agent's circumstances changes this space, then the agent may change from believing/disbelieving *p* to disbelieving/believing *p*, despite there having been no change in her underlying doxastic state. Sensitivism posits context-sensitivity of belief in the sense that there is this additional level of structure in the basic states underwriting facts about which propositions an agent believes at any given time.

In §1, I present an argument for sensitivism based on facts about assertion, in particular the principle that an assertion is sincere if and only if the assertor believes what she asserts. Using a Stalnakerian picture of how assertion works, I argue that we cannot make sense of the above principle without taking belief to be context-sensitive. In §2, I offer a preliminary model³ of

²I should also acknowledge that the argument I'll give for sensitivism is in some ways similar to the arguments for contextualism in DeRose (2002; 2009, ch. 3) and Schaffer (2008). DeRose and Schaffer's arguments, like mine, begin from a link between assertion and knowledge (or, in my case, belief), note a certain context-sensitivity of assertion, and conclude that knowledge (belief) must be similarly context-sensitive. Moreover Schaffer's argument (but not DeRose's) proceeds from a similar Stalnakerian picture of assertion to the one I will rely on. But my argument differs importantly from these previous arguments in that (among other things) their linking principle is the controversial knowledge norm on assertion (cf. Williamson 2000, ch. 11). My argument depends instead on a principle about which assertions are sincere, not which assertions are epistemically permissible.

³I say the model is preliminary because it deals only with outright belief. In Clarke (2013), I provide a model dealing with degrees of belief, which I also take to be context-sensitive; that model is a generalization of the preliminary one offered here.

belief based on the considerations raised in giving the assertion-based argument of §1. Briefly, here are a few features of the formal framework I shall offer: we identify contexts with sets of possibilities; to believe that p in a context is to rule out all not- p -possibilities in that context; we represent an agent's overall doxastic state by indicating which possibilities would be ruled out in which contexts, by using an ordering on possibilities.

But before we begin, it may be helpful to compare the view to be advanced here with some others in the literature. The term “sensitivism” is new, but there are others who have defended sensitivist views. For example, Nozick (1993, p. 96) offers “a (very) tentative suggestion”:

A belief excludes possibilities in a (type of) context. In another type of context, those very possibilities would not be excluded. I believe my new junior colleague is not a child molester. (Asked to list the people in the philosophy building who are not child molesters, I unhesitatingly place his name on the list.) Now the context changes; I need someone to watch my young child for two weeks. A mistake here would be serious—the stakes have escalated. Now I think more carefully. It is not that I did not believe in my colleague's innocence before. In that context, for those purposes, I did believe it; I did not consider, or assign a probability to, the possibility of his being a child molester. In this context, with higher stakes, I consider what probability that might have.

This sounds very close to the picture of belief I will offer in §2: belief is context-sensitive; in particular, belief that p is sensitive to the practical importance of p in the context at hand; and to believe that p in a context is to exclude not- p possibilities. But there are some important differences, too. For one thing, Nozick is concerned here with finding a middle path between “the radical Bayesian,” who thinks there are only degrees of belief and no outright beliefs, and a Levi (1980)-style view on which (rational) belief confusingly implies both certainty and revisability.⁴ So Nozick's suggestion here is not quite that outright belief is context-sensitive, nor that degrees of belief are context-sensitive, but rather that one may shift between having outright beliefs about p and having only degrees of belief about p , depending on contextual factors.⁵ Another important difference between Nozick's suggestion and my view is that he only definitely identifies high stakes as a contextual factor to which belief is sensitive, although he does speculate that there may be other such factors.

Kyburg (1983), Weatherson (2005), Ganson (2008), Sturgeon (2008), and Fantl and McGrath (2009) each argue for a different version of what we might call a “shifting threshold” view of belief.⁶ That is, on such views, to believe p outright is to have a high enough credence in p —but what counts as high enough credence for belief varies with context (and perhaps with the choice of p). In particular, changes in the practical importance of p will generally change the threshold

⁴Giving Levi's view a fair treatment would take us too far afield; suffice it to say that (a) Levi is concerned with rational belief, whereas I am concerned with belief *simpliciter*, and (b) Levi's claims about how rational beliefs change in response to various non-evidential factors do not map straightforwardly onto the present discussion of context-sensitivity.

⁵Nozick (1993, pp. 98–9) does seem to think that belief is context-sensitive in something more like the sense I have in mind—for one thing, he speaks of an agent as being disposed to believe a proposition p in a type of context C —but his discussion leaves it unclear exactly how this is to be spelled out, and what relation this view of belief has to his “(very) tentative suggestion,” quoted above.

⁶Bach (2005, p. 77) also adopts, but does not argue for, such a view.

for belief: if the stakes are higher, then one will need a higher credence in p to count as believing it outright, *ceteris paribus*. I substantially agree with what these authors have to say about outright belief;⁷ the main difference here concerns the arguments for our respective views. These authors are mainly concerned with using belief to explain rational action, and the effects of changes in one's practical interests on one's beliefs; on the other hand, I am mainly concerned with using belief to understand sincere assertion, and the effects of additional sorts of changes in one's conversational context on one's beliefs (e.g., explicit discussion of alternative possibilities). One consequence of this difference: defenders of a shifting threshold view must worry more than I about Ross and Schroeder (2014). In that paper, Ross and Schroeder argue for an alternative, non-sensitivist explanation of the data around rational action that the shifting threshold theorists cited above rely on in arguing for their views. They acknowledge that what I am calling the sensitivist explanation of the relevant data "can provide an elegant explanation", but argue that it is "not alone in doing so" (p. 259). That is, they do not argue that the sensitivist explanation is *bad*, but that there is a *better* non-sensitivist explanation. On the other hand, my argument depends on a completely separate range of phenomena—having to do with sincere assertion rather than rational action—and so I need not respond directly to Ross and Schroeder. In fact, the current paper could be read as an indirect response to Ross and Schroeder: if, as I shall argue, the sensitivist alone can give a satisfactory account of sincere assertion, this is a reason to prefer a sensitivist explanation of the data on rational action.

1 Belief and Sincere Assertion

The link between assertion and belief goes through the notion of sincerity. Here is a platitude about sincere assertion:

An assertion that p is sincere if and only if the assertor believes that p .

I say this is a platitude: it is at the core of our commonsense understanding of sincere assertion.⁸ It provides a desideratum for any theory of sincere assertion. Any theory that tells us the platitude is false has a significant mark against it. This is not to say that the platitude is an adequate analysis of sincere assertion, that there are no counterexamples to the biconditional. (I'm agnostic about this.) See, for instance, Pruss (2012); but note that even Pruss concedes (p. 545) that the platitude does "hold in normal cases, just as in normal cases one knows p if and only if p is true and one justifiably believes p ." Moreover, something close to the quoted platitude is widely defended as giving the correct analysis of sincere assertion,⁹ which suggests the platitude is a common starting-point.

⁷The primary difference in content between shifting threshold views and my own concerns the treatment of credence, rather than belief. Kyburg *et al.* crucially treat credence as context-insensitive. This paper's argument does not rely on any claim about credence, but my (2013) argues that to believe outright that p is to give p credence 1, and outright belief is context-sensitive because credences themselves are context-sensitive.

⁸Thus, Goldberg (2015) takes an even stronger principle as a platitude: "Simply put, when they are performed sincerely, assertions express or manifest one's beliefs. ... *This seems to be a feature of assertion that is common knowledge among all competent speakers*" (p. 8, my emphasis).

⁹Searle (1969, §3.4) accepts the platitude itself. So does Bach and Harnish (1979), and more recently, Bach (2008, p. 69). Williams (2002, p. 74) says that "someone who asserts that P and is sincere says something that is a *direct expression* of his belief that P" (emphasis in original). The official account of sincere assertion in Goldberg (2015, ch. 6) allows the possibility of sincere assertion unhooking from belief, given certain candidates for the epistemic

In what follows, I aim to uphold both the spirit and the letter of the platitude as stated above. This does not commit me to the absence of counterexamples. Rather, I want to see what belief must be like for the platitude to hold in normal cases.^{10,11} Upholding the platitude, I will argue, requires sensitivism.

The spirit of the platitude has it that there is a certain content of an assertion that must be mirrored in the assertor's beliefs if the assertion is sincere. For convenience, call this content the assertion's *sincerity content*. Then sincerity content is, by definition, whatever underwrites the spirit of, or the intuition behind, our platitude: the sincerity content of an assertive utterance is whatever content in the assertion must be matched in the assertor's beliefs for the assertion to be sincere. "Sincerity content" is, thus, a technical term, but it's one whose content is determined by common sense. By reflecting on the notion of sincerity, I will argue in §1.1 that given a Stalnakerian view of assertion, sincerity content must be shaped by the common ground¹² in the context of assertion.

On the other hand, I will argue in §1.2 that whether one has asserted that p (rather than some other proposition) is not sensitive to common ground in this way. As a result, if we are to uphold the platitude as worded—referring as it does to an assertion *that* p and to belief *that* p —as well as its spirit, we must accept that whether one believes that p in a particular context depends in part on the common ground in that context. This argument will be spelled out in full in §1.4.

Here is an overview of the argument to come. First, we get from Stalnaker the following picture of how assertion works. One makes a particular assertion in a particular conversational context. We can characterize a conversational context via the notion of *common ground*. The common ground in a given context is, approximately, what is commonly believed¹³ by the conversational participants. (For more on common ground, see Stalnaker (1978, 2002).) For present purposes, we can represent a conversational context by what Stalnaker calls the "context set": the set of possible worlds compatible with the common ground. The more information in the common ground, the smaller the context set will be. An assertion characteristically aims to add information to the common ground—hence, to reduce the context set. But how a successful assertion that p affects the context set depends on what the context set was to begin with. A

norm of assertion—e.g., in cases where the speaker has a belief she regards as unjustified for her. Nonetheless, on Goldberg's view the platitude holds in normal cases, and may be a correct analysis of sincere assertion, depending again on the correct epistemic norm of assertion. See Pagin (2016) for a survey of prominent accounts of assertion (including sincerity conditions), all of which, it seems to me, take the platitude to hold in normal cases. See footnote 14 for a notable dissenting view.

¹⁰Perhaps *bullshitting* in the sense of Frankfurt (2005) provides a type of counterexample of the platitude: one who asserts what she believes, but is indifferent to whether she believes what she says, sounds insincere. And it may be that there is, as Frankfurt says, "so much bullshit" these days. I'm happy to consider these cases abnormal in the relevant sense, however common bullshit may be. Thanks to an anonymous referee for raising this case.

¹¹Thanks to two anonymous referees for pressing me to say more about the platitude.

¹²But see §1.3 for a qualification here.

¹³Stalnaker (2002) argues that there are some cases where what is in the common ground is determined not by what the conversational participants believe, but by some weaker belief-like attitude: what the participants accept, where "accept" denotes a family of attitudes including believing, presuming, assuming, and accepting for the purposes of an argument or an inquiry (Stalnaker 2002, p. 716). In such a conversation, the sincerity principle in the text will have to be amended: one asserts sincerely iff one accepts what one asserts. In general, the idea is that an assertion is the expression of some belief-like state; the assertion is sincere if one is actually in the relevant state. For simplicity, in the main text, I consider only cases where the relevant state is belief.

successful assertion adds p to the common ground, and so removes all not- p possible worlds from the context set; but just what those worlds are, and indeed whether there are any such worlds to be ruled out, depends on what the context set was.

So much for our Stalnakerian background. Now let us see how we can fit our platitude into the picture so far. That is, we need to locate something in our Stalnakerian framework that can play the role of sincerity content.

Central to the story will be a distinction between an assertive utterance's *expressed content* and its *contextual content*. Expressed content is meant to be something like ordinary speaker-meaning or utterance-meaning. An assertive utterance's contextual content, on the other hand, is a function of both its expressed content and the common ground.

Here, then, is an overview of the picture of assertion I favour. A sentence, uttered in a certain way in a certain context, has a certain *expressed content*—generally, a proposition. This is pretty vague, and intentionally so: I have no insight to offer about which propositions are expressed by which utterances of which sentences; I hope to remain neutral in such debates. For any assertive utterance, pick your favourite story about what proposition was expressed, and that will work for my purposes.¹⁴ In general, expressed content need not be the same as semantic value—in particular, it will be different in cases of idiom, sarcasm, and perhaps some cases of implicature. Intuitively, uttering “Right, Fred’s a good dean and I’m a monkey’s uncle,” with a certain intonation, expresses the proposition that Fred is not a good dean. There are complicated things determining the expressed content of any given utterance, but if you want to know much more about it, you’ll have to find another philosopher, or a linguist.

Here is where I have a story to tell: a proposition, expressed assertively in a certain context, has a certain *contextual content*. A proposition can in general be identified with a set of possible worlds;¹⁵ the intersection of (a) the set of possible worlds corresponding to an assertion's expressed content with (b) the set of possible worlds compatible with the information in the common ground (Stalnaker's context set) determines its contextual content. (Put another way: the contextual content of an assertive utterance is equivalent to the conjunction of its expressed content with the propositions in the common ground.) I shall argue that an utterance's sincerity content is not its expressed content. Contextual content gets us much closer to the mark. (In §1.3 I will define a variant on contextual content, *intended content*, which I claim is identical with sincerity content.) There are thus two stages in getting from an utterance to what matters for sincerity. There are contextual factors involved in both stages (in the first stage, this will

¹⁴There might be some difficulty if you happen to be Herman Cappelen or Ernie Lepore, however. According to Cappelen and Lepore's Speech Act Pluralism, there is no such thing as *the* proposition expressed by a given utterance; and furthermore, there is no algorithm that will give all the propositions expressed (said, asserted, etc.) by an utterance. However, at least as of 2005, Cappelen and Lepore suspect that no one else is a Speech Act Pluralist of this sort, so I shall try not to be too worried (Cappelen and Lepore 2005, 5). Furthermore, Cappelen and Lepore (2005, 202) explicitly reject the sincerity principles I build my argument on: “Speakers need not believe everything they sincerely say.” So Cappelen and Lepore got off the bus well before this note.

I do think that some version of my argument might work even for Speech Act Pluralists, provided we appropriately restrict our sincerity principles—though this would of course be difficult to manage. Roughly, the idea is that if, for a given assertive utterance, we (somehow) pick an appropriate member of the set of propositions expressed by that utterance, the sincerity principles I rely on will hold for that proposition; furthermore, if we plug that proposition in as expressed content in my sense, we can run the argument to follow on the contextual/intended content generated thereby.

¹⁵This identification involves a certain simplification or idealization; see §2.1 for more on this.

include, e.g., determining the referents of indexicals), but the kind of context-sensitivity that interests us here comes in the second stage, the move from expressed content to contextual content.

The following three subsections are devoted to exploring the distinction between expressed and contextual content, and with an eye to locating sincerity content in our framework. Let us begin with some observations about sincerity.

1.1 Sincerity Content is not Expressed Content

A starting observation about sincerity: the sincerity of an assertion depends on what the speaker intends the assertion to do, or what effect the speaker believes the assertion will have if accepted. Insincere assertions involve a mismatch between the asserter’s outward actions and inward beliefs—but the mismatch must be *intentional*. One can assert sincerely and yet convey something one does not believe thanks to miscommunication. The spirit of, or the intuition behind, our platitude calls us to look for communicative intentions true to one’s beliefs when evaluating an assertion’s sincerity. So: what does one intend to do in asserting that p ?

It is widely agreed that what an assertion aims to do is add a proposition to the common ground. But what this amounts to depends on what the common ground was before the assertion. If an assertion of a sentence expressing the proposition p is accepted by its audience, the actual effect on the common ground will be to add p . But what the speaker intends assertion of a sentence expressing p to accomplish is to update the context set (the set of worlds compatible with the common ground) to the intersection of the set of p -worlds with the original context set. Some examples, with pictures, will help illustrate the point.




Figure 1

Let U (the big box) be the universe of all possible worlds. Figure 1a shows the proposition expressed by p : the area of the box not shaded out represents the worlds where p is true. This represents what I call the expressed content.

In Figure 1b, the oval C represents the worlds compatible with the common ground. The area outside the oval is shaded in darkly; the area inside the oval on the not- p side of the universe is shaded in more lightly. The idea is that the darkly shaded area represents worlds that are considered to have been ruled out even before the speaker’s assertion; then the assertion itself aims to update the common ground to include only the unshaded area, which is inside the oval and also on the p side of the universe. This represents what I call the contextual content.

Now we are in a position to understand why we should not take the aim of an assertion that p to be simply adding p to the common ground. If that were the case, we would take the assertion's sincerity content to be the entire section of the universe to the right of the vertical line. Suppose our interlocutors are talking about what kind of bird is sitting in the garden. It's clear to all involved that the bird is some kind of large gull. Amanda (who is about to speak) thinks the only likely large gulls in this area are the Herring Gull and the Lesser Black-Backed Gull (LBB), and she thinks her interlocutors agree with her on this count. Were the conversation closer to sea, Amanda would also think the bird might be a Great Black-Backed Gull (GBB). Amanda only trusts herself to distinguish LBBs from GBBs by their legs, so were the conversation closer to sea, she would suspend judgment about whether the bird is a GBB or an LBB, since its legs aren't visible. Consider Figure 2a. Here, p is the proposition that the bird is a Black-Backed Gull; the circle G represents worlds where the bird is a Great Black-Backed Gull (GBB). Suppose Amanda takes the common ground to rule out everything outside the oval C , including the circle G : that is, she takes the common ground to include the information that the bird in the garden is not a GBB. In that case, aiming to update the common ground to allow only worlds where there is an LBB in the garden, Amanda will say something like "It's a Black-Backed Gull," rather than the wordier "It's a Lesser Black-Backed Gull." Speakers can count on the common ground to do some work for them. Now, if the result of saying "It's a Black-Backed Gull" is the situation depicted in Figure 2a, the speaker has not achieved her aim. The common ground has not been updated as desired. Rather, a successful assertion would result in the situation of Figure 2b, with the GBB-worlds ruled out as well.




Figure 2

On the other hand, if the speaker takes the common ground at the time of her assertion to allow the worlds in both C and G —that is, if she takes her interlocutors *not* already to presume there are no GBBs here, perhaps because they think the shore is closer than Amanda does—and she says "It's a Black-Backed Gull", we can only interpret her as aiming to produce the situation of Figure 2a. Speakers working in different contexts, with different common ground, intend different things by assertively expressing the same proposition. In both contexts, our speaker expresses the proposition that the bird is a black-backed gull, but she aims or intends to produce different updated context sets in each case. My aim in mixing blue paint with yellow would not be to create something blue, but rather to create something green; likewise, a speaker's aim in assertively expressing p in a context is to use p as a tool to yield a certain updated context

set.¹⁶ Just as I would intend different things by mixing the same blue paint into red rather than yellow, I would intend different things by asserting p in one context rather than another.

This is why I claim it is better to identify sincerity content with contextual content than with expressed content. Sincerity content is, by definition, whatever content of an assertion must be mirrored in the assertor's beliefs for the assertion to be sincere. In other words, sincerity content is whatever content underwrites the intuition behind our platitude. Therefore, since the assertor's intentions are what matters in determining whether an assertion is sincere, and since the assertor's intentions are shaped by the common ground in the context of assertion, sincerity content must depend on common ground. Contextual content depends on common ground, whereas expressed content does not. Therefore, it is better to identify sincerity content with contextual content than with expressed content.

1.2 Expressed Content and Asserting that p

So much for sincerity content. What is the role of expressed content in this picture, aside from helping to determine contextual content? For our purposes, what matters about expressed content is the following: we say that one has asserted that p just in case one has made an assertive utterance with the expressed content p . That is, to say that one has asserted that p is to specify the expressed content of one's assertion, not its contextual content. Focussing on contextual content helps us to get at what one intends to do by asserting, with a view to getting at what one must believe in order for one's assertion to be sincere; but our normal way of talking about what people assert does not have this specialized goal. In the bird-in-the-garden case, it is more natural to say that by uttering "It's a Black-Backed Gull," Amanda has asserted that the bird is a Black-Backed Gull (p in Figure 2b) than that she has asserted that the bird is an LBB (the proposition corresponding to the unshaded area in Figure 2b).

A more compelling case for using expressed rather than contextual content to determine when one has asserted that p comes from considering pairs of cases where, intuitively, one asserts the same thing despite the common ground being different (i.e., despite the contextual content being different). Suppose the nation of Freedonia held an election last night. Hugh has never paid much attention to Freedonia or its politics, but he caught a glance at the election results in passing while reading his morning newsfeed. Hugh gleans that Firefly won the election, but the headline he saw didn't say what office was being contested. Now consider two conversations, C_{pres} and C_{pm} : in C_{pres} (Figure 3a) it is common ground that the election determined the new President of Freedonia; in C_{pm} (Figure 3b), on the other hand, it is common ground that the election determined the new Prime Minister of Freedonia. (In Figure 3, p is the proposition that Firefly won the election.)

Now consider what we should say about Hugh in either C_{pres} or C_{pm} , were he to say, "Firefly won the election." It seems to me that, intuitively, in both cases, Hugh has asserted that Firefly won the election. That is, it is intuitive that there is a proposition p such that, in both cases, Hugh has asserted that p . Therefore, we get the wrong result if we say that one asserts that p just in case one utters assertively something with the contextual content p : we would then have to say that Hugh has asserted different things C_{pres} and C_{pm} , since his assertions in each case have different contextual content. We would then say that Hugh has asserted in C_{pres}

¹⁶Thanks to an anonymous referee for this analogy.




Figure 3

that Firefly was elected President of Freedonia last night—or perhaps something even more complicated, spelling out whatever other presuppositions Hugh shares with his interlocutors. But, intuitively, we want to say that Hugh asserts the same thing in both cases—namely, that Firefly won the election. On the other hand, if we say that an assertion that p is one with the expressed content p , we do not run into this problem, since the expressed content of Hugh’s assertion in both cases is the same. This matches our ordinary practice of describing assertions.

1.3 A Qualification

Recall our observation about sincerity from the beginning of §1.1: the sincerity of an assertion should depend on what the speaker intends the assertion to do, or what effect the speaker believes the assertion will have if accepted, rather than on what effect the assertion actually would have if accepted. This suggests a problem with our identification of sincerity content with contextual content, because it is possible for the speaker to be mistaken about what the common ground is. A speaker’s communicative intentions can only be shaped by what she takes the common ground to be, not by what the common ground actually is (in cases where there is a difference). When a speaker asserts that p , she intends to update the context set to the intersection of the set of p -worlds with *what she thought the original context set was*. If the speaker was mistaken about what the original context set was, this will not be the same as the actual effect of successful assertion that p .

Here is an example to illustrate. Moe and Curly work in the same office. Curly is a bit clueless, and Moe is self-absorbed and not very nice. Today is Friday.

Moe and Curly: Moe needs to deposit his paycheque before the start of business on Monday. He has written a big cheque which will be cashed on Monday, and he must make sure there is enough money in his account to cover it. It would be hard for him to make it to the bank after work before the bank closes, but he has a lot of work to do and wants to avoid leaving early if possible. All day, Moe has been complaining about his problem to his co-workers, and asking if anyone knows if the bank is open on Saturdays. He knows that as of two weeks ago, the bank was open on a Saturday, but he still worries that they might have changed their hours since.

Towards the end of the day, Moe runs into Curly. Presuming that Curly knows all about Moe’s problem—surely everyone knows by now—Moe asks if Curly knows if

the bank will be open tomorrow. But Curly hasn't heard anything about Moe's predicament, and presumes that Moe just doesn't know what the bank's normal hours are. Curly was at the bank on a Saturday two weeks ago, so he says, "Yes, they're open Saturdays." This, I take it, is a sincere assertion.

The actual effect of Curly's assertion is to bring Moe to believe that the bank has not deviated from its regular hours; the intended effect of Curly's assertion is to bring Moe to believe that the bank's regular hours include Saturdays. We get the right verdict (that Curly's assertion is sincere) if we say that Curly's assertion is sincere iff he believes what he intends to bring Moe to believe by his assertion. We get the wrong verdict (that Curly's assertion is insincere) if we say that Curly's assertion is sincere iff he believes what he actually brings Moe to believe.

We can use the following simple model to represent Curly's assertion. Let the universe contain just four possible worlds, A , B , C , and D . In worlds A and B (only), the bank was open on Saturdays as of two weeks ago. In worlds A and C (only), the bank's hours now include Saturdays—so the bank will be open tomorrow, in particular. Thus, in worlds B and C (only), the bank has changed its Saturday hours in the last two weeks. Moe thinks it is common knowledge (since he's been going on about it all day) that the bank was open on Saturdays as of two weeks ago; so he thinks the context set is $\{A, B\}$. Curly, however, is oblivious to this. Moreover, Curly hasn't considered, and does not expect Moe to have considered, the rather remote possibility that the bank has changed its hours in the past two weeks, so he thinks the context set is $\{A, D\}$. Because of the mismatch between Moe's and Curly's understanding of the situation, the only presumption Moe and Curly have in common is that the bank has not recently changed its hours from being closed on Saturdays to being open on Saturdays. Therefore, the actual context set is $\{A, B, D\}$.¹⁷ Curly asserts that *the bank is open on Saturdays*; the expressed content of his assertion is true at worlds A and C . Curly aims to update the context set to $\{A\}$, and thinks he succeeds in doing so by eliminating D from the context set. On the contrary, Curly actually updates the context set to $\{A\}$ by eliminating B and D , since Moe no longer thinks the bank might have changed its hours after Curly's assertion, and Curly then takes himself and Moe both to presume D to be false.

Intuitively, Curly's assertion is sincere. We only get this verdict if we judge Curly by the context set as he takes it to be, rather than by the actual context set. Curly has no opinion about whether the bank has changed its hours, but his effect on the actual context set was to exclude this possibility. On the other hand, Curly's effect on the context set as he takes it to be is just to exclude the possibility that the bank's (unchanged) hours do not include Saturday—and this does reflect his opinion.

Therefore, we should take sincerity content to correspond to the intersection of expressed content with (not the actual context set, but) the set of possible worlds compatible with what the speaker thinks is the common ground. Call this the *intended content* of an assertion. In most cases, intended content and contextual content will be identical; they come apart only when the speaker is mistaken about the common ground. However, where contextual content and intended content differ, it is the latter that matters for determining sincerity. Thus, my

¹⁷This depends, of course, on defining the common ground to include all and only what is presumed by all conversational participants, which may be controversial. However, so long as common ground is defined such that it is possible for a speaker to be mistaken about what the common ground is, it should be possible to produce cases like the ones in the text.

official view is that sincerity content should be identified with intended content.

1.4 The Argument

Now we have all the pieces we need to get to the context-sensitivity of belief. Suppose *S* has asserted that *p*. This means that *S* has made an assertive utterance with the expressed content *p*. *S*'s assertion is sincere if and only if *S* believes that *p*. But whether *S*'s assertion is sincere or not depends on *S*'s communicative intentions—what effect *S* intends her assertion to have on the common ground of the conversation—and this is not completely determined by the expressed content of her utterance. To capture *S*'s communicative intentions, we must look at the intended content of her utterance, which depends on what the common ground is (according to *S*). Intended content is identical with sincerity content, and the sincerity content of *S*'s utterance by definition tells us what belief *S* must have for her assertion to be sincere. Therefore, what belief *S* must have for her assertion to count as sincere depends on what the common ground is (according to *S*). But we have already said that *S*'s assertion is sincere if and only if *S* believes that *p*. Therefore, whether *S* believes that *p* or not depends on what the common ground is (according to *S*). This is the conclusion we hoped for.

Here is the argument more explicitly. Suppose that *S* has made an assertive utterance whose expressed content is *p* and whose intended content is *p*'.

1. *S* has asserted that *p*. (By Freedonia cases.)
2. An assertion that *p* is sincere iff the assertor believes that *p*. (Platitude.)
3. Therefore, *S* has asserted sincerely iff *S* believes that *p*. (By 1 and 2.)
4. *S* has asserted sincerely iff *S* has a belief matching the sincerity content of *S*'s assertion. (By definition of sincerity content.)
5. Sincerity content is identical with intended content. (By bird-in-the-garden and Moe/Curly cases.)
6. Therefore, *S* has asserted sincerely iff *S* has a belief with content *p*'. (By 4 and 5.)
7. Therefore, *S* believes that *p* iff *S* has a belief with content *p*'. (By 3 and 6.)

Since *p*' is our assertion's intended content, it is context-dependent in a way that *p* is not: *p*' is a function of *p* and the (perceived) common ground. Therefore, this argument establishes that belief that *p* means different things in different contexts, with different (perceived) common ground: belief is context-sensitive.¹⁸

Our conclusion here establishes a form of context-sensitivity of belief, but does it establish sensitivism? Recall that the sensitivist does not claim merely that people have different beliefs in different contexts. Sensitivism is not, after all, simply the view that people change their minds sometimes. Rather, it is the view that the contents of our beliefs show more structure than is

¹⁸Of course, to get context-sensitivity of belief in general from this argument, we need to rule out the position that only belief in propositions one happens to be asserting is context-sensitive. I have no argument against this position, but it seems entirely unappealing.

commonly thought. Therefore, to establish sensitivism, we need to show that whether S believes that p depends, in part, on what alternative possibilities S is taking seriously. The above argument concludes that whether S believes that p in a context depends, in part, on what S takes the common ground to be in that context. I think this does amount to a dependence on what alternatives S is taking seriously: arguably, what it is for a possibility to be unexcluded by the common ground (i.e., for a possibility to be included in the context set) is for the agents involved to be taking that possibility seriously. If S thinks all parties to the conversation are taking some possibility seriously, then she must think that she herself is taking that possibility seriously as well. S thinking that S is taking a possibility seriously is not quite the same thing as S taking the possibility seriously, but it's pretty darn close, and in most cases will amount to the same thing. Perhaps there are cases where S thinks that she is taking a possibility seriously, but actually is not; however, I cannot think of such a case, nor can I convince myself that a serious objection to the present view based thereon can be offered. If such an objection is possible, I leave it to other philosophers to find it.

2 A Model for Belief

I have presented the foregoing as an argument for sensitivism, but we can draw a stronger conclusion. That is, the considerations of the previous section do not push us to accept just any sensitivist model of belief. In this section, I offer a way of modelling belief states which easily makes sense of the picture of sincere assertion given above.

2.1 Some Simplifications Eliminated

To begin, there are some simplifications I have made up to this point for reasons of clarity, which can now be dispensed with. In particular, I have assumed that successful assertions only increase the common ground, and reduce the context set. This is, of course, not the case. For example, consider again the case where the conversants are wondering what kind of bird is in the garden. An assertion that it might be a squirrel rather than a bird does not aim to eliminate any possibility; rather, it raises a new possibility previously excluded. Prior to the assertion, it was taken for granted by all parties that the creature in the garden is a bird; if the assertion is successful, then all parties will now suppose that it is either a bird or a squirrel. This means an enlargement of the context set, and a reduction in the common ground (as we have been understanding it), since squirrel-possibilities are now under consideration. And there are some assertions which aim only to point out or make salient that certain things are already in the common ground; assertions of the form “As we all know, ...” or “Of course, we’re taking for granted that ...” are often of this sort. (Cf. Abbott 2008.)

Another simplifying idealization is built into my use of possible worlds, and my representing propositions as sets of possible worlds. The trouble with this is that it requires us to represent logically equivalent propositions as identical. The same set of possible worlds corresponds to both p and $(p \wedge q) \vee (p \wedge \neg q)$. Therefore, if we represent an agent as believing the former, we must also represent her as believing the latter, since we represent both by the same object. In other words, representing propositions as sets of possible worlds means building in logical omniscience.

The problem here bears some similarity to a family of problems of logical omniscience encountered by formal epistemologists of various stripes—I have in mind Bayesians and epistemic logicians, in particular—but the difference between my project and theirs makes certain familiar moves unavailable to me. I am interested in understanding and representing belief in general, whether rational or not; but formal epistemologists are usually concerned with representing more ideal agents. This gives them some excuse for building logical omniscience into their models: it is at least debatable whether ideally rational agents must possess perfect logical knowledge. But it is a clear problem if a model of non-ideal, not necessarily rational belief has it that all believers are logically omniscient.

However, the modesty (in a sense) of my goal also opens up a way of dealing with the problem of logical omniscience which is unavailable to Bayesians and epistemic logicians. Let me first explain how I deal with the problem, and then explain why this strategy is unavailable to those others.

We must modify the idea of possible worlds at work in our model. In fact, I shall cease calling them possible worlds, since we will give up any connection with logical or metaphysical possibility; instead, let us speak of *points* in the model. In the model, every proposition is assigned a truth value at each point. Since there is no one-to-one correspondence between points in the model and genuine possible worlds, the truth value a proposition receives at a given point are not determined by its truth value in any possible world. Rather, it is up to the agent, so to speak, to assign a truth value. We can still represent propositions as sets of points, following the standard practice of representing propositions as sets of possible worlds. But, since the distribution of truth values over the universe of points is determined by the agent, rather than by the world, we cannot read facts about individuation of propositions off of the sets of points representing them in the model (or vice versa).

Thus, we must also modify the idea of propositions at work in our model. The pattern of the agent's beliefs will determine the population of points: if the agent believes that¹⁹ p and $(p \wedge q) \vee (p \wedge \neg q)$ are logically equivalent, or indeed that they are the same proposition, then there will be no point where one is true and the other false; otherwise, there will be such points. Thus, the pattern of the agent's beliefs also determines the individuation of propositions in the model: if the agent believes that p and q are equivalent, then both propositions will be represented by the same set of points, so, as far as the model is concerned, they are the same proposition. Once we have this clarification of what the points in our model are, everything works just as it would if we were using possible worlds. To be clear, though, no claim is made here about which worlds are genuinely possible, or about which propositions are genuinely distinct. It might be helpful to think of the “propositions” in the model as propositions under a mode of presentation, or something of the sort. Because the present aim is to model belief, rational or otherwise, the world outside the agent does little to shape the model used to represent the agent's beliefs.²⁰

The above strategy for dealing with the problem of logical omniscience is open to us because we are not giving a logic.²¹ If we were, we would need some recursive rule for determi-

¹⁹More carefully, it would be better to say that two propositions will be represented by the same set of points if the agent believes *as if* they are logically equivalent. That is, what matters is whether one treats p and $(p \wedge q) \vee (p \wedge \neg q)$ as equivalent, not whether one explicitly believes that the two are equivalent. After all, one may lack the concept of equivalence, but still treat various propositions as equivalent.

²⁰Thanks to an anonymous referee for helping to clarify this issue.

²¹Or rather, it is *easily* open to us: people who like the “impossible worlds” program (cf. Berto 2012) can take a

ning how complex propositions are represented in the model. For example, typically, if p and q are represented, respectively, by sets of points P and Q , then $p \wedge q$ would be represented by the intersection of P and Q . Thus, we would not have the option of representing $p \wedge q$ by anything but the set of worlds where both p and q are true; all agents would have to respect conjunction introduction and elimination, so to speak. But I have no need to construct complex propositions from simple ones; for my purposes, there is no problem with regarding all propositions as simple. Therefore, any proposition can be represented by an arbitrary set of points. The set of points where $p \wedge q$ is true need not be the set of points where both p and q are true.

To be sure, representing propositions as sets of points in the model still runs into some problems analogous to those that arise in using sets of possible worlds. In particular, any two propositions which the believer’s logical beliefs treat as necessarily true will be represented as identical. This may make the model unsuitable for representing, say, beliefs about mathematics. I shall not offer a solution of this problem; I concede that the way of modelling belief on offer here is only useful for representing a restricted range of beliefs. However, I think the restriction is not so severe as to remove all usefulness or interest from the models—we can, after all, still represent beliefs about brains in vats, bank hours, and indeed anything else in the physical world.

2.2 The Model

Now we are ready to present the model. §2.2.1 gives a formal statement of the model, and §2.2.2 explains more informally how it works, including an application to a form of skepticism.

2.2.1 Definitions

Let a doxastic state \mathcal{S} for a set of propositions P be a quadruple $(U_{\mathcal{S}}, \mathcal{C}_{\mathcal{S}}, \|\cdot\|_{\mathcal{S}}, \preceq_{\mathcal{S}})$, where $U_{\mathcal{S}}$ is a set of points, $\mathcal{C}_{\mathcal{S}}$ is a set of non-empty subsets of $U_{\mathcal{S}}$ (“contexts”), $\|\cdot\|_{\mathcal{S}}$ is a valuation function taking members of P to subsets of $U_{\mathcal{S}}$, and $\preceq_{\mathcal{S}}$ is a reflexive, transitive, $\mathcal{C}_{\mathcal{S}}$ -well-founded²² relation among the members of $U_{\mathcal{S}}$.²³ “ $x \preceq y$ ” is meant to be read as, “ y is ruled out at least as strongly as x (according to the agent whose beliefs are being represented).” It might also be helpful—though not quite accurate—to think of “ $x \preceq y$ ” as saying that x is more plausible than y according to the agent; I shall say more about this shortly. For a proposition $p \in P$, $\|p\|$ is the set of points in U at which p is true. If x is a point with $x \in \|p\|$, then we will write $x \models p$; if A is a set of points with $A \subseteq \|p\|$, then we will write $A \models p$.

Let a context be represented by a context set $C \in \mathcal{C}$. Because \preceq is well-founded, we can always find, for any C , a minimally ruled-out (maximally plausible) subset $B_C \subseteq C$; that is, there is a $B_C \subseteq C$ whose members are those $x \in C$ such that for all $y \in C$, $x \preceq y$. Then we say that p is believed in context C if and only if $B_C \models p$. That is, p is believed in C just in case it is true at all the least ruled-out (most plausible) C -points.²⁴

similar approach to the problem of logical omniscience, but must face a slew of objections we avoid through the modesty of our aims. That is, we have no need to interpret or justify claims like “There are worlds/points where p and $\neg p$ are both true,” and this makes our life relatively easy.

²²For any $C \in \mathcal{C}_{\mathcal{S}}$, there is a minimal element $m \in C$: for all $x \in C$, $m \preceq x$.

²³The subscripts on U , \mathcal{C} , $\|\cdot\|$, and \preceq will be suppressed where there is no ambiguity which doxastic state is in question.

²⁴The approach of characterizing belief using plausibility rankings on possibilities in this way is not new. In fact,

We can also define a dual operator to belief, which is useful for some applications. Say that p is *doxastically possible* in a context C , and write $\langle B_C \rangle p$, just in case there is a point $x \in B_C$ with $x \models p$. Note that the equivalence $\langle B_C \rangle p \Leftrightarrow B_C \not\models \neg p$ does not hold in general, because we have no assurance that the agent’s valuation function is classical. Informally, $\langle B_C \rangle p$ says that the agent does not rule out p in C ; if the agent is logically coherent, we might say that p is consistent with her beliefs in C .

2.2.2 Discussion

The model defined above is meant to be a generalization of the picture we saw in §1 to cases not involving assertion. The context sets $C \in \mathcal{C}$ in the model intuitively correspond to the Stalnakerian context set generated by the information in the common ground of a conversation: just as the common ground restricts interlocutors’ attention to a subset of the space of possible worlds, the context sets $C \in \mathcal{C}$ in our model restrict a believer’s attention to a subset of the points populating her doxastic state \mathcal{S} . Not all subsets of U clearly correspond to any recognizable context (this is why \mathcal{S} includes the distinguished set of contexts \mathcal{C}). For instance, a singleton set containing just one point, with nothing ruled out (or in) is hard to understand as corresponding to anyone’s beliefs in any real-life context. It seems to me that genuine belief always involves ruling something out; but even if this is not the case, our framework can handle it—if every subset of U corresponds to some possible context, then let \mathcal{C} be the power set of U , the set of all subsets of U . The important thing is that we have the *option* of excluding some subsets of U as not corresponding to any real context.

Note that the inclusion of a given set C in \mathcal{C} doesn’t mean that the agent has already turned her attention to that set. Otherwise, we would have to posit relatively few context sets in the doxastic state of ordinary agents. Rather, a set C should be included in \mathcal{C} just in case the agent already has a robust disposition to rule out certain points in C in favour of certain others. Even if an agent has yet to consider a certain body of presuppositions, she may have definite views about what would be true under those presuppositions; if that is the case, should include in her doxastic state the corresponding context set and plausibility ordering thereon. (Audi’s 1994 distinction between dispositional beliefs and dispositions to believe might affect how one chooses to fill out \mathcal{C} .)²⁵

But wait: one might worry about this appeal to “robust dispositions”: after all, there is by now a large body of empirical work showing that people’s judgments and preferences are highly sensitive to all sorts of incidental factors—framing, ordering, anchoring, and so on. (See, e.g., Slovic and Lichtenstein 1968; Lichtenstein and Slovic 1971, 1973; Tversky and Kahneman 1974; Kahneman et al. 1982.) In short, if one needs to have a robust disposition to rule out possible world x in favour of possible world y for a set containing x and y to be included in \mathcal{C} , then \mathcal{C} will have very few members: the empirical evidence suggests we do not have dispositions like this, but rather dispositions (say) to rule out x in favour of y when judging in a tidy room, but to rule out y in favour of x when the bin is overflowing with pizza boxes (Schnall et al. 2008).

This worry is misplaced: this empirical evidence suggests that there will be *many*, not few,

it is a standard technique in belief revision theory. See, e.g., Baltag and Smets (2006), Board (2004), Stalnaker (1996, §3).

²⁵Thanks to an anonymous referee for pressing me to clarify this, and for raising the objection in the next paragraph.

members of \mathcal{C} . Two clarifications will help to see why this is so. First, recall that we have already distinguished between possible worlds and points in the model. If a believer sometimes treats possible world x as more plausible than possible world y and sometimes *vice versa*, we can represent this by including in the model of her doxastic state multiple points agreeing with x on every proposition. (Cf. the discussion of transitivity failure below.) Second, note that I have not constrained the manifesting conditions of the dispositions-to-rule-out which settle whether a set is to be included in \mathcal{C} : in particular, I have not claimed that for a set containing x and y to be included in \mathcal{C} , one must be disposed (say) to rule out x in favour of y in *all* circumstances, or *whenever* one thinks about the propositions true in x but not y . If one regularly rules out x in favour of y on pleasant days and regularly makes the reverse judgment under gloomy skies, we can represent this in the model by one context set capturing one's fair-weather dispositions and another capturing one's rainy dispositions. Thus, the empirical evidence showing that our judgments and preferences are (not random or erratic but) extremely sensitive to contextual factors like framing etc puts pressure on us to include more, not fewer, sets in \mathcal{C} .

The relation \leq can be seen as generalizing the shading used in §1. Figures 1–2 used at most three colours: dark grey for possibilities²⁶ ruled out by the common ground, light grey for possibilities ruled out by the assertion of p , white for possibilities not ruled out. Context-independent doxastic states, on this model, can be visualized as using arbitrarily many colours: if \mathcal{S} has it that $x \leq y$, then shade y darker than x . In a context C , these shadings on the doxastic state can be used to generate a model in the style of §1: everything outside of C is shaded dark grey; the lightest possibilities within C (i.e., those in B_C) are shaded white; and everything else in C is shaded light grey. Of course, dark grey, light grey, and white will not mean quite the same thing in this model as in that of §1, since we are now working with a model for belief, not assertion. Here, then, dark grey would indicate points unconsidered, or not taken seriously in the context; light grey would indicate points considered but ruled out (disbelieved); and white would indicate points considered, and not ruled out (believed).

This explains why we required \leq to be reflexive, transitive, and \mathcal{C} -well-founded. Reflexivity and transitivity make \leq nicely behaved as a preorder: there are no loops with $x < y$, $y < z$, and $z < x$.²⁷ \mathcal{C} -well-foundedness ensures that any context $C \in \mathcal{C}$ will have a non-empty minimally ruled-out subset B_C . Now, it might seem that requiring transitivity is not just an innocent formal convenience. That is, someone might think that x is more likely than y in a context C_1 where z is not under consideration, and likewise think that y is more likely than z if x is not under consideration, and that z is more likely than x if y is not under consideration (contexts C_2 and C_3 , respectively); and it seems that to represent such an agent's doxastic state, we must violate transitivity. But this is not the case, for we can always introduce a fourth point, z' , which agrees with z on every proposition $p \in P$, but with $z' < x < y < z$ (see Figure 4, where an arrow from a point v to a point w indicates that $w < v$).²⁸ Then, if we take $C_1 = \{x, y\}$, $C_2 = \{y, z\}$, and

²⁶I promised to stop referring to points in the model as “possible worlds,” but “possibilities” is just too convenient a term to give up completely.

²⁷Here, of course, “ $a < b$ ” means $a \leq b$ and $b \not\leq a$.

²⁸Incidentally, note that so far, I have tacitly assumed that there is no context in which the agent takes all three possibilities x, y and z into consideration; if there is such a context, then whichever possibility is (or possibilities are) not ruled out will be the one that the new point (or points) agrees with on all propositions. When there is no such synoptic context, the three original possibilities are in a symmetrical situation, and it makes no difference which one is “copied” by the new point(s).

$C_3 = \{x, z'\}$, we get a representation of the agent as described, but without having to violate transitivity. This is possible because there is no requirement that distinct points of U must differ in the propositions they make true.

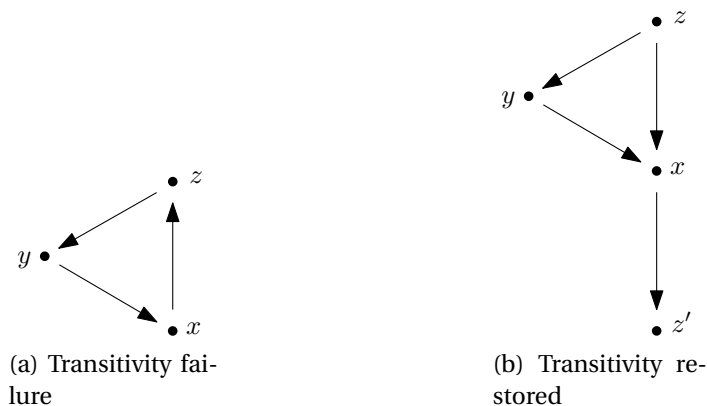


Figure 4

Here is another way of putting the point. A doxastic state \mathcal{S} for the propositions P is a good model of an agent S if and only if \mathcal{S} gets S 's pattern of belief and non-belief in the propositions P right across the contexts represented by the members of $\mathcal{C}_{\mathcal{S}}$. The points $x \in U_{\mathcal{S}}$ are useful for representing those propositions and contexts, but do not have any independent content; outside the context of the model \mathcal{S} , it does not make sense to ask about whether the agent S believes, say, that x is actual. Thus, without violating transitivity, we can build a doxastic state which accurately models an agent who believes p rather than q in a context where r is not a serious alternative, believes q rather than r when p is not a serious alternative, and believes r rather than p in a context where q is not a serious alternative. Just take the doxastic state of Figure 4b, with $\|p\| = \{x\}$, $\|q\| = \{y\}$, and $\|r\| = \{z, z'\}$. Then the context $C_1 = \{x, y\}$ will be one where the agent believes p rather than q , $C_2 = \{y, z\}$ one where the agent believes q rather than r , and $C_3 = \{x, z'\}$ one where the agent believes r rather than p . That is to say that in C_1 , the agent believes p and rules out the only serious q -possibility, and takes for granted that r is false, or in other words, does not take any r -possibility seriously. Thus, the transitivity constraint does not reduce the expressive power of doxastic states.

Note that the different levels or shadings of the points in U induced by \leq should not be thought of as degrees of belief, at least in the usual sense. Rather, they should be thought of as *dispositions to believe*. What it means for some point x to be shaded more lightly than some other point y is not that the agent believes x is the case more strongly or with more certainty than y ; rather, what it means is that in a context where attention is restricted to the points x and y , the agent will believe those things true in x , regardless of what is true in y . This is part of why I said above that it is better to think of “ $x \leq y$ ” as saying that y is ruled out at least as strongly as x , rather than as saying that x is at least as plausible as y (according to the agent). To illustrate, here is a plausible (but highly simplified) way of modelling many people's attitudes to skeptical scenarios; I have in mind people who feel the force of skeptical arguments.

Let $U = \{2H, 1H, 2H_{BIV}, 1H_{BIV}\}$. These points correspond, respectively, to: the world where I have two hands and nothing else is unusual ($2H$); the world where I have only one hand,

having lost one in an accident ($1H$); the world where I am a handless brain in a vat having experiences as of having two hands ($2H_{BIV}$); and the world where I am a handless brain in a vat having experiences as of having lost one hand in an accident ($1H_{BIV}$). The worlds $2H$ and $2H_{BIV}$ are indistinguishable to me, as are $1H$ and $1H_{BIV}$: for each pair, I have exactly the same experiences in both members of the pair. For convenience, let us write “ $x \simeq y$ ” when $x \leq y$ and $y \leq x$. Thus, we will have $2H \simeq 2H_{BIV}$ and $1H \simeq 1H_{BIV}$. Furthermore, let us suppose that I do seem to have two hands—that is, I have an experience as of having two hands. I see two hands before me, I have no memory of losing a hand in an accident, and so on. Then I shall be inclined to rule out any point where I seem to have one hand rather than any world where I seem to have two hands: we will have $2H \leq 1H$, $2H_{BIV} \leq 1H_{BIV}$, $2H \leq 1H_{BIV}$, and $2H_{BIV} \leq 1H$. This gives us the situation of Figure 5.




Figure 5: Points and their ordering.

Now, in most ordinary contexts, the skeptical possibilities $2H_{BIV}$ and $1H_{BIV}$ will not arise. Thus, an ordinary context set $C_{ordinary}$ will only include $2H$ and $1H$. Since $2H$ is the unique minimal element of such a $C_{ordinary}$, we will have $B_{C_{ordinary}} = \{2H\}$; so the agent will, in ordinary contexts, believe whatever is true in $2H$ —in particular, that she has two hands. This is depicted in Figure 6a. However, by raising skeptical possibilities to salience, an interlocutor can expand the context set to a new $C_{skeptical}$ which includes $2H_{BIV}$. Since $2H_{BIV} \simeq 2H$, we will have $B_{C_{skeptical}} = \{2H, 2H_{BIV}\}$, so that the agent only believes whatever is true in *both* $2H$ and $2H_{BIV}$ —in particular, she will not believe that she has two hands, since this is false at $2H_{BIV}$, though neither will she believe that she has no hands, since this is false at $2H$. This is the situation depicted in Figure 6b.²⁹ To recast these comments in terms of doxastic possibility, it is doxastically possible for the agent that she does not have hands in $C_{skeptical}$, but not in $C_{ordinary}$: if p is the proposition that the subject has hands, then we have $\langle B_{C_{skeptical}} \rangle \neg p$, but $\neg \langle B_{C_{ordinary}} \rangle \neg p$.




Figure 6: (a) Ordinary and (b) skeptical contexts.

²⁹In Figure 6b, $1H_{BIV} \in C_{skeptical}$. This is somewhat arbitrary; it is not clear to me whether, in general, skeptical arguments bring to salience skeptical versions of possibilities already ruled out. In any case, nothing here turns on the inclusion or exclusion of $1H_{BIV}$ in the skeptical context set.

What I mean when I say that the ordering given by \leq does not give an ordering of degrees of belief should now be clear. The effect of skeptical conversational manoeuvre like saying “But how do you know you’re not a brain in a vat, stimulated to have experiences as of having hands?” is not to bring the skeptical brain-in-a-vat possibility to equal plausibility as the ordinary alternative; rather, it is to destroy belief in the ordinary alternative by bringing into consideration an alternative possibility one cannot rule out.

Let’s take a step back for a moment, and ask why we need an ordering \leq at all. That is, why generalize the shading system of Figures 1–2? Suppose one thinks the actual world is a member of some set A , though one is unsure which member of A is actual. Then it’s plausible that one ought always to consider all members of A to be live possibilities. That is, one ought not exclude the members of A from B_C , for any context C . Conversational salience and raised stakes might cause one to consider previously ignored possibilities one is unwilling to rule out, such as skeptical scenarios; but (so this line of thought might go) one ought never allow mere circumstance to cause one to rule out non-crazy ways one genuinely thinks the world might be. To characterize such an agent’s beliefs across contexts, we need only two colours, so to speak: those possibilities the agent would rule out, and those she would not; and she would not rule out the members of A .

I shall not dispute the claim that one ought to have beliefs of this sort for two reasons. First, there are a lot of “ought”s in the previous paragraph; this should be a red flag. We are concerned here with belief *simpliciter*, not rational belief, or the beliefs on ought to have. We want a way of modelling belief that can characterize irrational agents as well as rational ones, and there is no reason to think that irrational agents might not be caused by changes in context to rule out, or ignore, or forget ways one thinks the world might be. Therefore, in general, we will need more than two colours to characterize agents in general. Second, I think there is value in having a multicoloured model even for the beliefs of agents who do satisfy the “ought”s above. This is partly because belief is not far distant from other belief-like attitudes, such as supposing, hypothesizing, etc.—and there is a lot one could say about using the present account of belief to deal with, e.g., sincere assertion in suppositional contexts, but that would take us too far afield—but it is also partly because even ideally rational agents sometimes speak (and think) loosely. That is, there are contexts where it makes sense to ignore possibilities one would, on reflection, refuse to rule out. For one example, contexts where one is reasoning about idealized scenarios—contexts where one chooses to ignore certain genuine but irrelevant possibilities to facilitate productive reasoning—will be of this sort. For another example, see my discussion of the preface paradox in Clarke (forthcoming).

In this connection, it is also worth pointing out a certain limitation of the present account: I have nothing to say here about how an agent’s beliefs should or will be updated in the face of new evidence. But the present approach to belief is inspired by plausibility models used in certain theories of belief revision (e.g., Baltag and Smets 2006). On this sort of approach, roughly, an agent’s beliefs put a plausibility ordering on worlds, with the agent believing those propositions true in the most plausible worlds; but in the face of various kinds of events, the plausibility ordering may change, so that erstwhile implausible worlds become the most plausible, and so the agent’s beliefs change. On such a framework, the whole ordering matters, not just which worlds are maximally plausible. If two agents agree on which worlds are most plausible, but disagree on the rest of the ordering, then they will have the same beliefs for the moment, but by performing the “same” revision on the two orderings (that is, if the two agents respond the

same way to the same event), the two agents will come to have different beliefs. So, by taking the present approach and keeping track of the agent's ordering \preceq , we allow the possibility of giving a Baltag-Smets-style account of belief revision—though it is of course not trivial to do so.³⁰

Before we move on to generalize from the particular skeptical case described above, a word about ruling out is in order. One often hears about ruling out possibilities from proponents of relevant alternatives (RA) approaches to knowledge. The sort of ruling out invoked to explain knowledge is different from the sort involved here. We might say that RA theorists are primarily concerned with an epistemic sort of ruling-out, whereas I am primarily concerned with a psychological sort of ruling-out. Thus, an RA theorist of a certain (contextualist) stripe might diagnose skeptical arguments as being effective because they bring to relevance remote possibilities about which one has no evidence, and so which one cannot rule out; since one cannot rule out a relevant alternative to one's having hands, one cannot know that one has hands. On the other hand, it certainly is possible to rule out in my sense the possibility that one is a brain in a vat. I suspect the skeptic would have little success destroying, say, my grandmother's belief that she has hands (to say nothing of what the skeptic might do to my grandmother's *knowledge* that she has hands); it is a fact about her psychology that it would be very difficult to make her take seriously the possibility that she is a brain in a vat. Thus, the model of Figure 6b simply would not describe her, even when the skeptic gets to work on her. On the other hand, a certain proportion of undergraduates in introductory epistemology classes are very easy to put in the situation described by Figure 6b. The short of it is: my models say nothing about one's reasons for ruling out a possibility, or for ignoring it.

Now, back to our example. It is, of course, highly simplified. Here are the important parts, though, for our diagnosis of effective skeptical arguments: we have a population of ordinary worlds ($1H$ and $2H$); a population of skeptical worlds where a target proposition (that I have hands) is false ($1H_{BIV}$ and $2H_{BIV}$); the skeptical worlds have ordinary counterparts from which they are indistinguishable ($1H_{BIV} \approx 1H$ and $2H_{BIV} \approx 2H$); ordinary context sets include only ordinary worlds, but bringing skeptical possibilities to salience can expand the context set to include some skeptical worlds; in particular, some skeptical worlds will be included in the new belief set, $B_{C_{skeptical}}$. This pattern will generalize. Effective skeptical arguments are generally effective precisely because they point to worlds which are indistinguishable from those we would otherwise believe are actual; we have no reason to rule out the skeptical worlds without also ruling out the ordinary worlds, and so we wind up accepting the skeptical possibilities as legitimate and un-ruled-out. Thus, belief in a target proposition is destroyed.

But this does not require that we have equal confidence in the skeptical and ordinary possibilities. It only requires that we do not rule out the skeptical possibilities. This is why the ordering given by \preceq is not an ordering on degrees of belief, but rather indicates conditional dispositions to believe. I give a model for degrees of belief in Clarke (2013).

The preceding example is meant primarily to serve as an illustration of how my doxastic states are supposed to work, but it also has some intrinsic interest of its own. In particular, it is supposed to be an advantage of contextualism that it explains the effectiveness of skeptical arguments. (Briefly: the skeptic puts her victim into a context in which skeptical scenarios are salient, and this raises the standards for knowledge; in such a context, ordinary people do not

³⁰Ranking theory (Spohn 1988, 1990) also deals with (im)plausibility orderings on worlds, but is less obviously similar to the present approach.

know even ordinary facts about their own hands; the skeptic then leads us to conclude, illegitimately, that we never know ordinary facts about our own hands; but on the contrary, in ordinary contexts, the standards for knowledge are not so high, and ordinary people do, after all, ordinarily know ordinary things about their hands.) The example we have just worked through gives a similarly appealing diagnosis of the effectiveness of skeptical arguments without committing us to any controversial theory of knowledge or of “know”. Just as the contextualists say, skeptical arguments serve to make salient remote possibilities, e.g., that one is a brain in a vat; this destroys belief in ordinary alternatives to those possibilities, e.g., that one has hands. Lacking belief, one also lacks knowledge, and the skeptic has succeeded.

However—and this may be a strength, as well—unlike the contextualist response to skepticism, the present explanation says nothing about why or whether skeptical arguments are bad. That is, what I offer here is not so much a response to skepticism as a diagnosis of its effectiveness. I say this may be a strength because there have been a number of objections posed in the literature to the contextualist response to skepticism (e.g., Bach 2005, Feldman 2001, 2004, Klein 2000); it seems to me that what is appealing about the contextualist response is the diagnostic part—it makes sense to think that skeptics undermine our knowledge of ordinary things by bringing to light remote possibilities.³¹ But this diagnosis is compatible with multiple senses of “undermining”: maybe the skeptic reveals to us that we never know very much, or maybe the skeptic tricks us into thinking so by causing us to lose knowledge temporarily through loss of belief. One might plausibly think that the former diagnosis is correct in cases of what we might call “healthy” skepticism (e.g., explicitly raising the possibility that this used-car salesperson is pushing one to buy a lemon reveals that one does not know him to be as trustworthy as he seems), but the latter diagnosis is correct in cases of “extreme” skepticism (e.g., it is difficult for some to ignore the possibility of a Cartesian demon, but one need not have evidence ruling this possibility out in order to know that one has hands). Put slightly differently, any successful skeptical argument results in destroying belief through raising alternative possibilities that the believer is disposed not to rule out. With extreme skeptical arguments, these are possibilities the believer need not rule out in order to know the targeted proposition; with healthy skeptical arguments, these are possibilities the believer must rule out in order to know the targeted proposition. (Of course, this leaves open the question of exactly which skeptical arguments are healthy and which are extreme.) Thus, there is a desirable flexibility in the sensitivist explanation of how skeptical arguments work when they are successful.³²

Finally, now that we have our model, we can give a precise explication of the sincerity principle appealed to in §1. That principle was: An assertion that p is sincere if and only if the assertor believes that p . Let $\mathcal{S} = (U, \mathcal{C}, \|\cdot\|, \preceq)$ be the doxastic state of an agent who has just made an assertive utterance with the expressed content p . Then, the agent has asserted sincerely that p if and only if $B_C \models p$, where $C \in \mathcal{C}$ is the context set corresponding to what the agent thinks is the common ground (i.e., C is the set of all points where the information the agent thinks is in the common ground is true). Recall that “ $B_C \models p$ ” is to be read as “[the agent] believes that

³¹This should not be confused with the claim that the contextualist response correctly diagnoses the disagreement between the skeptic and the dogmatist; on the contrary, the contextualist response is notably criticized for getting the latter diagnosis wrong in Feldman (2004) and Klein (2000). Rather, the claim is that the contextualist response correctly diagnoses the skeptic’s effect on her victim (i.e., on someone who is persuaded by, or at least feels the force of, the skeptical argument).

³²Thanks to an anonymous referee for pressing me on this point.

p ,” and it is easy to see that this is a formalization of the principle in question. Furthermore, note that since C is determined by what the assertor thinks the common ground is, there is a clear sense in which the condition $B_C \models p$ means that the assertor’s beliefs mirror the intended content of her utterance. After all, the intended content corresponds to $C \cap \|p\|$; so we might say that our formalized principle requires that the agent rule out all points outside the intended content—i.e., that the assertor’s beliefs rule out all possibilities she intends her assertion to rule out.

2.2.3 Assessment

To conclude, let us assess the strengths and limitations of the model offered here. As I just mentioned, one limitation is that it is entirely qualitative, in the sense that it does not deal with degrees of belief at all. I address this limitation in Clarke (2013). Furthermore, the model is entirely static: it is designed to represent an agent’s beliefs at a particular time, but has nothing to say about how an agent’s doxastic state might change over time. This limitation I shall not address in the present work, because it derives from a third limitation. This is a peculiar quirk of my project: I aim to give an account of belief which will be of use and interest to epistemologists, but I am primarily concerned with the nature of belief, not just rational belief. Without the restriction to rational belief, it is hard to put constraints on how a doxastic state might evolve over time; hence, my model is merely static.

On the other hand, the model does provide a framework for an account of rational belief, and so should already be of interest to epistemologists. We might require, say, that certain beliefs be held in all contexts, perhaps via some requirements on the makeup of the universe U . Or we might put some additional constraints on \leq . Or we might have some requirement on how one responds to context-shifting manoeuvres, or to new evidence. I hope to have gotten clearer on what belief is like partly as a means to getting clearer on what rational belief is like.

The model has other, more specific strengths too: it is designed to allow an intuitive treatment of sincere assertion, of course; and I have argued that it allows a nice explanation of what skeptical arguments can do. Furthermore, as I argue in Clarke (forthcoming), it allows a surprising solution to the preface paradox, according to which preface writers may be consistent.³³

References

- Abbott, B.: 2008, ‘Presuppositions and Common Ground’, *Linguistics and Philosophy* **21**, 523–38.
- Audi, R.: 1994, ‘Dispositional Beliefs and Dispositions to Believe’, *Noûs* **28**(4), 419–34.
- Bach, K.: 2005, ‘The Emperor’s New “Knows”’, in G. Preyer and G. Peter (eds.), *Contextualism in Philosophy: On Epistemology, Language and Truth*, pp. 51–89. Oxford, Oxford.

³³Versions of this paper were presented at the 2010 Western Canadian Philosophical Association, the 2013 Society for Exact Philosophy, the 2015 Seminario Interuniversitario de Ciencia Cognitiva (Valladolid); my thanks to all three audiences. Thanks also for helpful discussion and comments to Paul Bartha, Dan Greco, Jonathan Jenkins Ichikawa, Carrie Ichikawa Jenkins, Masashi Kasaki, Joe Morrison, Adam Morton, Patrick Rysiew, Max Weiss, a slew of anonymous referees, and especially to Chris Stephens.

- Bach, K.: 2008, 'Applying Pragmatics to Epistemology', *Philosophical Issues* **18**, 68–88.
- Bach, K. and R. M. Harnish: 1979, *Linguistic Communication and Speech Acts*. MIT Press, Cambridge, MA.
- Baltag, A. and S. Smets: 2006, 'Dynamic Belief Revision over Multi-Agent Plausibility Models', in G. Bonanno, W. van der Hoek, and M. Woolridge (eds.), *Proceedings of the 7th Conference on Logic and the Foundations of Game and Decision Theory (LOFT 2006)*, pp. 11–24. University of Liverpool.
- Berto, E.: 2012, 'Impossible Worlds', in E. N. Zalta (ed.), *Stanford Encyclopedia of Philosophy (Winter 2012 Edition)*. URL = <<http://plato.stanford.edu/archives/win2012/entries/impossible-worlds/>>.
- Board, O.: 2004, 'Dynamic Interactive Epistemology', *Games and Economic Behavior* **49**, 49–80.
- Cappelen, H. and E. Lepore: 2005, *Insensitive Semantics: A Defense of Semantic Minimalism and Speech Act Pluralism*. Blackwell, Oxford.
- Clarke, R.: 2013, 'Belief Is Credence One (In Context)', *Philosophers' Imprint* **13**(11), 1–18.
- Clarke, R.: forthcoming, 'Preface Writers Are Consistent', *Pacific Philosophical Quarterly*.
- Cohen, S.: 1999, 'Contextualism, Skepticism, and the Structure of Reasons', *Philosophical Perspectives* **13**, 57–89.
- DeRose, K.: 1992, 'Contextualism and Knowledge Attributions', *Philosophy and Phenomenological Research* **52**(4), 913–29.
- DeRose, K.: 2002, 'Assertion, Knowledge, and Context', *Philosophical Review* **111**, 167–203.
- DeRose, K.: 2009, *The Case for Contextualism*. Oxford, Oxford.
- Fantl, J. and M. McGrath: 2009, *Knowledge in an Uncertain World*. Oxford, Oxford.
- Feldman, R.: 2001, 'Skeptical Problems, Contextualist Solutions', *Philosophical Studies* **103**, 61–85.
- Feldman, R.: 2004, 'Comments on DeRose's "Single Scoreboard Semantics"', *Philosophical Studies* **119**(1–2), 23–33.
- Foley, R.: 1993, *Working Without a Net: A Study of Egocentric Epistemology*. Oxford, Oxford.
- Frankfurt, H. G.: 2005, *On Bullshit*. Princeton, Princeton.
- Ganson, D.: 2008, 'Evidentialism and Pragmatic Constraints on Outright Belief', *Philosophical Studies* **139**(3), 441–458.
- Goldberg, S. C.: 2015, *Assertion: On the Philosophical Significance of Assertoric Speech*. Oxford, Oxford.

- Kahneman, D., P. Slovic, and A. Tversky: 1982, *Judgment Under Uncertainty: Heuristics and Biases*. Cambridge, New York.
- Kaplan, M.: 1996, *Decision Theory as Philosophy*. Cambridge, Cambridge.
- Klein, P.: 2000, 'Contextualism and the Real Nature of Academic Skepticism', *Philosophical Issues* **10**, 108–16.
- Kyburg, H.: 1983, 'Rational Belief', *The Behavioral and Brain Sciences* **6**, 231–45.
- Levi, I.: 1980, *The Enterprise of Knowledge: An Essay on Knowledge, Credal Probability, and Chance*. MIT Press, Cambridge, MA.
- Lichtenstein, S. and P. Slovic: 1971, 'Reversals of Preference between Bids and Choices in Gambling Decisions', *Journal of Experimental Psychology* **89**, 46–55.
- Lichtenstein, S. and P. Slovic: 1973, 'Response-Induced Reversals of Preference in Gambling: An Extended Replication in Las Vegas', *Journal of Experimental Psychology* **101**, 16–20.
- Maher, P.: 1986, 'The Irrelevance of Belief to Rational Action', *Erkenntnis* **24**(3), 363–84.
- Nozick, R.: 1993, *The Nature of Rationality*. Princeton, Princeton.
- Pagin, P.: 2016, 'Assertion', in E. N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*, winter 2016 edition. Metaphysics Research Lab, Stanford University.
- Pruss, A. R.: 2012, 'Sincerely Asserting what You Do not Believe', *Australasian Journal of Philosophy* **90**(3), 541–6.
- Ross, J. and M. Schroeder: 2014, 'Belief, Credence, and Pragmatic Encroachment', *Philosophy and Phenomenological Research* **88**(2), 259–88.
- Schaffer, J.: 2008, 'Knowledge in the Image of Assertion', *Philosophical Issues* **18**, 1–19.
- Schnall, S., J. Haidt, G. L. Clore, and A. H. Jordan: 2008, 'Disgust as Embodied Moral Judgment', *Personality and Social Psychology Bulletin* **34**(8), 1096–1109.
- Searle, J. R.: 1969, *Speech Acts: An Essay in the Philosophy of Language*. Cambridge, Cambridge.
- Slovic, P. and S. Lichtenstein: 1968, 'The Relative Importance of Probabilities and Payoffs in Risk-Taking', *Journal of Experimental Psychology Monograph Supplement* **78**(3, Part 2), 1–18.
- Spohn, W.: 1988, 'Ordinal Conditional Function: A Dynamic Theory of Epistemic States', in W. Harper and B. Skyrms (eds.), *Causation in Decision, Belief Change, and Statistics*, Vol. II, pp. 105–34. Kluwer, Dordrecht.
- Spohn, W.: 1990, 'A General Non-Probabilistic Theory of Inductive Reasoning', in R. Schachter, T. Levitt, J. Lemmer, and L. Kanal (eds.), *Uncertainty in Artificial Intelligence*, Vol. 4, pp. 149–58. North-Holland, Amsterdam.

- Stalnaker, R.: 1978, 'Assertion', *Syntax and Semantics* **9**, 315–32. Reprinted in Stalnaker (1999, 78–95). Page references are to the reprinted version.
- Stalnaker, R.: 1996, 'Knowledge, Belief and Counterfactual Reasoning in Games', *Economics and Philosophy* **12**(2), 133–63.
- Stalnaker, R.: 1999, *Context and Content: Essays on Intentionality in Speech and Thought*. Oxford, Oxford.
- Stalnaker, R.: 2002, 'Common Ground', *Linguistics and Philosophy* **25**, 701–21.
- Sturgeon, S.: 2008, 'Reason and the Grain of Belief', *Noûs* **42**(1), 139–65.
- Tversky, A. and D. Kahneman: 1974, 'Judgment Under Uncertainty: Heuristics and Biases', *Science* **185**, 1124–31.
- Weatherson, B.: 2005, 'Can We Do Without Pragmatic Encroachment?', *Philosophical Perspectives* **19**(1), 417–43.
- Williams, B.: 2002, *Truth and Truthfulness: An Essay in Genealogy*. Princeton, Princeton.
- Williamson, T.: 2000, *Knowledge and Its Limits*. Oxford, Oxford.