

## Assertibility and Sensitivity

Geoff Pynn

Received: 11 December 2012 / Accepted: 3 September 2013  
© Springer Science+Business Media Dordrecht 2013

**Abstract** Epistemologists have proposed various norms of assertion to explain when a speaker is in an epistemic position to assert a proposition. In this article I propose a distinct necessary condition on assertibility: that a speaker should assert only what she sensitively believes, where a subject's belief is sensitive just in case the subject would not hold it if it were false. I argue that the Sensitivity Rule underwrites simple explanations for three key features of assertibility that pose explanatory challenges to other prominent proposals: the fact that assertibility is open under known entailment, the general impropriety of assertions that a lottery ticket has lost made purely on the basis of the speaker's knowledge of the odds, and the fact that assertibility varies widely with features of the conversational context. I close by considering three distinct roles the Sensitivity Rule might play in the overall theory of assertibility.

**Keywords** Assertion · Closure · Norms of assertion · Sensitive belief · Subjunctive conditionals

### 1 Epistemic Assertibility

Assertions can be criticized on many grounds. One kind of complaint is epistemic. A speaker who asserts without the right sort of evidence, justification, or warrant for what she asserts can be legitimately criticized on that basis. Theorists have proposed various 'norms of assertion' to explain this phenomenon. The dominant idea is that there is an epistemic rule governing the practice of assertion; the challenge has been to say what the rule is. The candidates specify some epistemic relation between the speaker and the content she asserts, and then link the propriety of her assertion to whether she stands in the specified relation with respect to that content. The most prominent proposal is the Knowledge Rule:

---

G. Pynn (✉)  
Department of Philosophy, Northern Illinois University, Zulauf Hall 915, DeKalb, IL 60115, USA  
e-mail: gpynn@niu.edu

(KR) One should assert that P only if one knows that P.<sup>1</sup>

KR has many competitors, for example, the Rational Credibility Rule (Douven 2006), the Reasonable-To-Believe Rule (Lackey 2007), the Certainty Rule (Stanley 2008), and the Safety Rule (Pelling 2012).

No epistemic rule tells the whole story about proper assertion. One should not, generally speaking, make rude assertions; one could satisfy any epistemic rule of assertion but assert improperly because rudely. Even ignoring rudeness and other non-epistemic factors, KR states only a necessary condition for proper assertion. KR leaves open the question whether one is in an epistemic position to assert anything one knows. Most participants in the debate between KR and its competitors have echoed Williamson in treating their favored epistemic rule as *constitutive* of assertion; the debate concerns not only how to understand assertibility, but also how to differentiate assertion from other speech acts. If the epistemic rule of assertion is treated as its constitutive norm, then it is natural to think that there are no other rules specifically governing assertion; as Williamson sees it, 'all other norms for assertion are the joint outcomes of [KR] and considerations not specific to assertion' (Williamson 2000, 241). Even so, the supposition that KR is the constitutive norm of assertion does not imply that knowledge is sufficient for epistemically proper assertion, since other considerations not specific to assertion may imply that any speaker who knows that P but fails to meet some distinct condition would do something epistemically improper by asserting that P.

In this article, I'll propose and defend a distinct necessary condition for epistemically proper assertion. You *sensitively believe* P just in case you believe that P, and would not believe that P if P were false. I propose that speakers should assert only what they sensitively believe. The simplest statement of the Sensitivity Rule is:

(SR) One should assert that P only if one sensitively believes that P.<sup>2</sup>

Further considerations will require elaborating the proposal in various ways, but none represents a significant departure from this initial statement of the idea. I won't argue directly against KR or any of its competitors. But SR easily explains key aspects of epistemic assertibility that these other rules struggle to accommodate.

The Sensitivity Rule links assertibility to a condition familiar to epistemologists. Sensitive belief is a central component of Nozick 1981's theory of knowledge, and closely related to the notion of a 'conclusive reason' that is the building block of the account of knowledge developed by Dretske 1971. It also plays a key role in DeRose 1995's contextualist account of knowledge attributions. Nozick's proposal that sensitive belief is necessary for knowledge continues to attract attention, though the

<sup>1</sup> See Williamson 2000, Chap. 11.

<sup>2</sup> It is important that the rule require that the speaker believe that P, and not just that she satisfy the condition that she would not believe that P if P were false. Suppose I have no belief about the number of stars in the sky, but choosing a number  $n$  at random, I assert that there are  $n$  stars in the sky. Clearly, my assertion is epistemically improper, even if by a stroke of luck it happens to be true. Yet it's very likely true that if there were not  $n$  stars in the sky, I would (still) not believe there to be. It would be a decisive objection to SR if it did not deliver the verdict that such an assertion was improper, including a requirement that the speaker believe the proposition asserted enables it to do so. Thank you to an anonymous reviewer from saving me from this error.

attention is just as often critical as sympathetic.<sup>3</sup> As we'll see, some of the considerations that make sensitivity problematic as a condition for knowledge make it well suited to explain assertibility.<sup>4</sup>

I'll discuss three arguments for the Sensitivity Rule. First, epistemic assertibility is open under known entailment; the Sensitivity Rule is uniquely well positioned to explain why (Sect. 2). Second, echoing a point first made by DeRose 1996, the Sensitivity Rule explains why one typically can't assert, purely on the basis of one's knowledge of the odds, that someone's lottery ticket has lost (Sect. 3). Third, the Sensitivity Rule explains why epistemic assertibility is context sensitive, and in particular why it tends to vanish as error possibilities become conversationally salient (Sect. 4).

A few terminological stipulations: I'll say that a speaker is *in a position to assert* that P, or that P is *assertible* by her, or that her assertion of P is *proper* or *appropriate* or *warranted* just in case she would not be subject to any legitimate epistemic criticism were she to assert that P.<sup>5</sup> Henceforth, I'll drop the qualifiers 'epistemic' and 'epistemically', leaving them understood. The idea that assertion is governed by epistemic rules is one way of explaining why some propositions are assertible by a speaker but others aren't. We could instead eschew rule-talk in favor of the idea that assertion involves representing oneself as standing in a certain epistemic position, or committing oneself to being in that position. Then epistemic criticism of a speaker's

<sup>3</sup> Kripke's famous 'Nozick-bashing' lectures, reprinted as Chap. 7 in Kripke 2011, are an influential source of skepticism about the sensitivity requirement; see also Chap. 7 of Williamson 2000. Becker and Black (2012) collect recent work on the link between sensitivity and knowledge. The majority of the papers in that volume do not endorse the idea that sensitive belief is required for knowledge; see in particular the contribution by Vogel for vigorous criticism. On the other side of the ledger, in addition to the sympathetic contributions to the Becker and Black volume, see Ichikawa 2011 for a defense of something close to the sensitivity condition, as well as Cross 2010 and DeRose 2010, who both argue that *insensitivity* of belief is sufficient to falsify a knowledge attribution.

<sup>4</sup> Sensitivity has a close cousin known as *safety*. Sosa 1999 introduces the notion of safety as the contrapositive of sensitivity: a safe belief is one that would be held only if it were true. Other writers characterize a safe belief as one that could not easily have been falsely held. Pelling (2012) defends a Safety Account of assertion, on which one must assert that P on basis B only if one could not easily have falsely asserted a proposition similar to P on a basis similar to B. There are interesting similarities between this proposal and my own. One of Pelling's arguments for the Safety Account concerns its ability to explain the impropriety of lottery assertions; in Sect. 4 I argue for the Sensitivity Rule on the same grounds. In general I find much to like about the Safety Account, in part because I think that sensitivity and (certain characterizations of) safety are less independent than epistemologists sometimes assume.

However, I think the Safety Account is less well positioned than the Sensitivity Rule to explain the highly context-sensitive nature of assertibility (discussed in Sect. 4 below) and doesn't explain the openness of assertibility (discussed in Sect. 2 below). Moreover, I disagree with Pelling's judgments about cases he calls 'existentially known assertions' and with both him and Lackey about cases the latter calls 'selfless assertions' (Lackey 2007). It seems to me that in both cases, the speaker does something epistemically wrong by asserting, and so I count it as a strike against the Safety Account that it treats them as epistemically proper. Still, this may not be a deep disagreement; it reflects the somewhat orthogonal question of whether epistemically proper assertion requires belief. It may be possible to tweak both the Sensitivity and Safety Rules to yield versions that do and do not require belief. Thanks to an anonymous referee for bringing Pelling's paper to my attention

<sup>5</sup> This is not a definition of assertibility or warranted assertion, just a characterization of the particular phenomenon I'm interested in. In a subtle and careful paper, Coffman 2011 distinguished between an assertion's being epistemically improper and its being improper on epistemic grounds: in the latter case, the assertion does not violate an epistemic rule but suffers some other impropriety, the explanation of which involves a defect in the speaker's epistemic position. I adopt a usage on which assertion is epistemically proper only if it is proper in both of Coffman's senses.

assertion would be legitimate whenever she fails to stand in the represented position or takes on a commitment that she is not in an epistemic position to take on. It is plausible to think that we can translate between talk of representation, commitment, and rules without risking anything of significance, at least so far as the assertibility data go. In this article, I will employ rule-talk for the sake of consistency and brevity.

Finally, SR links assertibility to a subjunctive conditional. It is common, in discussions of sensitivity in epistemology, to treat the sensitivity conditional as a claim about possible worlds; in particular, as being equivalent to the claim that at the possible world or worlds most similar to the actual world where P is false, the subject does not believe that P. I will resist this impulse. The rule I am concerned with has to do with the satisfaction of a condition involving a subjunctive conditional. We arrive at the related claim about possible worlds by treating the sensitivity conditional as a counterfactual, and presupposing a widely accepted semantic theory of counterfactuals. I don't wish to call into question the semantic theory; indeed, in Sect. 4 I will make use of it. But the shift to talk of similarity invites errors. As Lewis and Stalnaker, the inventors of the semantic theory, as well as subsequent writers have agreed, the relevant notion of similarity cannot be understood simply as overall similarity between possibilities, and its precise characterization remains a matter of controversy.<sup>6</sup> Moreover, as Lewis and Stalnaker also agreed, which possibilities count as closest is a highly context-sensitive matter. Finally, it is hasty to assume that the subjunctive conditional relevant to sensitive belief always expresses a counterfactual. For example, a subjunctive conditional that begins 'If P were false...' could naturally be used as a premise in an abductive argument that P is, in fact, false. SR says that P is assertible by S only if what is expressed by 'If P were false, S would not believe that P' is true, regardless of whether what is expressed is a counterfactual conditional.

## 2 Assertibility and Closure

A speaker can be in a position to assert P without being in a position to assert Q, even if she knows that P entails Q: assertibility is open under known entailment. Cases are easy to construct from familiar examples. Donna is looking at a new apartment, talking on her cell phone with her mother. When she walks into the bedroom, she notices that the walls are bright red. Somewhat taken aback, she asserts to her mother that the bedroom is red. The assertion is epistemically flawless, or would be in many easily imaginable circumstances. In particular, its epistemic propriety does not require that Donna have any special reason or evidence for thinking that the bedroom isn't white and lit by red lights. She needn't, for example, first check to make sure there are no red light bulbs in the lamps illuminating the walls before asserting that the walls are red. Yet without so checking she is liable to do something epistemically improper by asserting that the bedroom isn't white and deceptively lit by red lights to appear red. So her epistemic position upon looking around the bedroom and noticing the red walls can be sufficient for

<sup>6</sup> 'He would look better if his beard were shorter': that could easily be true, even though at the possible worlds most like the actual world but where his beard is shorter, it is only microscopically shorter, and so close enough to its actual length to have no noticeable effect on his appearance.

properly asserting that the bedroom is red, but not for properly asserting something she knows to be entailed by that proposition.

The phenomenon is ubiquitous. For each of the statements in group A, there are easily imaginable circumstances in which a speaker would be in a position to assert it but not in a position to assert the corresponding statement in group B:

A

This is a bottle of merlot.

We won't be able to afford the trip to Cardiff this summer.

Martha just left for work.

B

This isn't a mislabeled bottle of cabernet.

None of us will win a large lottery this spring.

Martha didn't just tell me a very convincing lie about where she was going.

These examples all indicate that assertibility is open:

*Openness.* For some P and Q, a subject can be in a position to make an epistemically proper assertion that P but not in a position to make an epistemically proper assertion that Q, despite knowing that P entails Q.

What explains the openness of assertibility?

The Sensitivity Rule underwrites a simple explanation for Openness. Even knowing that P entails Q, your belief that P may be sensitive though your belief that Q is not. Given Donna's basis for believing that the bedroom is red (i.e., that it looks red), she would not believe it was if it were not red (for then the bedroom would not look red). But if she has no special reason for believing that the walls aren't white and deceptively lit by red lights beyond its low probability given her background knowledge, then it is not the case that if the bedroom were so lit she would not believe that it wasn't. In the unlikely event that clever red lights were installed she would still believe that the room was not white and deceptively lit by red lights. Thus, as far as SR is concerned, she is in a position to assert the former but not the latter. Moreover, if we alter the case in such a way that Donna's belief about the lighting becomes sensitive, it appears that she is in a position to assert that the room isn't deceptively lit. For example, if she has just checked the bulbs and verified that they're indeed white, then her belief that the room isn't deceptively lit is sensitive, and she is in a position to assert as much. The alteration to the case that makes her belief sensitive also makes it assertible; SR is confirmed.

The explanation works for the other pairs in A and B as well. Given the ordinary circumstances in which the A proposition would be assertible, the speaker's belief in the B proposition would not be sensitive. If we alter the circumstances in such a way as to render the belief in B sensitive, we find that it also becomes assertible. SR explains why the B proposition was not assertible though the A proposition was and suggests the kind of evidence the speaker would need for the B proposition to be in a position to assert it; to wit, evidence that makes her belief sensitive.

One standard objection to the claim that sensitive belief is necessary for knowledge is that knowledge satisfies something like closure under known entailment. If we assume that a subject can know a proposition even though she knows it to entail propositions that she does not sensitively believe, then either knowledge is not closed under known entailment or sensitive belief is unnecessary for knowledge. Nozick himself happily chose the first horn of the dilemma, but few have followed him in this. Yet the fact that sensitive belief is open under known entailment means that the Sensitivity Rule explains a key aspect of assertibility. Moreover, if knowledge *is* closed under known entailment, then the openness of assertibility implies that KR doesn't provide a complete explanation of assertibility. Suppose that knowledge is necessary and sufficient for assertibility, and that Donna knows that if the bedroom is red then it's not white and deceptively lit by red lights. Since knowledge is closed under known entailment, she knows the former only if she knows the latter. Given the supposition, then, she is in a position to assert the former only if she is in a position to assert the latter. Since she can be in a position to assert the former without being in a position to assert the latter, the supposition must be rejected.<sup>7</sup>

Rather than endorsing a condition for assertibility that is not closed, could we instead deny Openness? Perhaps the proposition about the lights is assertible by Donna, but by asserting it she would act improperly along some non-epistemic dimension. For example, her assertion might be improper or misleading because irrelevant. There is no denying that given the hitherto fairly mundane direction of the conversation, the assertion would be a surprising move. But it is implausible to suggest that the assertion's irrelevance creates an appearance of *epistemic* impropriety. Suppose she stops the conversation mid-course to assert that her cousin is concerned about the disappearing bees. Though at least as irrelevant and perhaps subject to censure, this assertion would generate no appearance of epistemic impropriety. She could change the course of the conversation in many surprising ways without ever risking that appearance.

Perhaps Donna's assertion would pragmatically convey that she has some special evidence about the lights. Since she does not, her assertion would be misleading, and hence improper. Though not epistemically improper in the narrower sense relevant to assertibility, the impropriety might still count as

---

<sup>7</sup> Few epistemologists would endorse the straightforward principle that knowledge is closed under known entailment. The widely accepted 'closure' principle is rather: if a subject who knows that P competently deduces Q from P, thereby coming to believe Q while retaining her knowledge of P, then she knows that Q; see, for example, Hawthorne 2005 and Kvanvig 2006, as well as David and Warfield 2008 for a detailed discussion of various Chisholmings. But moving to this weaker principle does little to mitigate the challenge posed by cases like Donna's to the KR-based explanation of assertibility. Suppose that Donna knows that the bedroom is red, but before asserting it deduces and comes to believe that it is not white and cleverly lit by red lights. The competent deduction principle implies that she now either knows both propositions or knows neither. Given the supposition that knowledge is necessary and sufficient for assertibility, she's now either in a position to assert both propositions or neither. But performing that mental exercise would neither put her in a position to assert that the walls are not cleverly lit nor undermine her ability to assert that the room is red. Moreover, independent of any particular articulation of the closure principle, it is quite unattractive to suppose that the reason Donna is in a position to assert that the bedroom is red, but not that it's not white and cleverly lit is that she knows the former but not the latter. Yet this is just what KR requires if it is to explain the divergence in assertibility.

epistemic in a broader sense; it would mislead her interlocutor about her overall epistemic position. But what explains how this additional information would be conveyed? One obvious answer involves an appeal to conversational implicature. Yet if the conveyed material were conversationally implicated, it ought to be cancellable.<sup>8</sup> Her assertion about the lights would not be helped at all by an attempt at cancellation. If anything, such an attempt seems to make things even worse: 'The bedroom isn't white and cleverly lit by red lights – though I haven't actually checked the color of the lights; I'm only saying that because I think it's highly unlikely.' Far from restoring the assertion's propriety, the cancellation attempt highlights the speaker's epistemic misdeed. Moreover, how would the implicature be generated? One very natural answer is that an epistemic rule governing assertion underwrites an inference from the assumption that Donna is speaking properly to the false conclusion that she has a certain kind of evidence for what she has asserted; in particular, to the kind of evidence that puts her in a position to assert properly. But that answer is unavailable here, since the proposal is that the assertion is improper for some reason other than its conflicting with an epistemic rule for assertion.

A different strategy for denying Openness is to say that while assertion is governed by the Knowledge Rule, in cases such as Donna's, the speaker would lose her belief in the proposition that the wall is red in the course of drawing the inference to the proposition that the wall isn't white and lit by red lights.<sup>9</sup> The trouble with this strategy is that it does not explain what is wrong with an assertion made by a speaker in Donna's epistemic position who *does* retain and form the relevant beliefs. A speaker in Donna's epistemic position who confidently believes that the wall is red, and that it is not white and lit by red lights, would still do something epistemically improper by asserting that it's not white and lit by red lights, given that she has no evidence for this proposition beyond her ordinary background assumptions.<sup>10</sup> Moreover, while it does seem that speakers who assert propositions they do not believe do something improper, what such speakers do wrong does not seem to be what Donna would do wrong by asserting that the wall is not white and lit by red lights. She would make an assertion not warranted by her epistemic position; she would not (or, not merely) go wrong by asserting a proposition she does not believe.

It is not only the Knowledge Rule that struggles to explain Openness. Any account on which assertibility is wholly identified with an epistemic condition that is closed under known entailment (either literally or in the weaker sense characterized by something like a competent deduction principle) will be similarly flummoxed. It is plausible to think that rational credibility, justification, certainty, and safety are all relations that satisfy a closure principle of some variety. This implies that these competitors to KR will face analogous explanatory challenges. If the epistemic

<sup>8</sup> Not all conversational implicatures are cancellable, so this is not a decisive objection. Davis 1998, for example, points out that some conversational implicatures are entailed by what is said. But that is not a plausible understanding of this case. The cancellability test for conversational implicature may not be foolproof, but it does figure centrally in our understanding of the phenomenon.

<sup>9</sup> Thanks to an anonymous reviewer for this suggestion.

<sup>10</sup> Donna may be able to get away with the assertion with special pleading: 'Oh, come *on*—*surely* it's not white and lit by red lights!' But a flat-out, unhedged assertion would be improper even if Donna believed what she asserted.

relation specified by the candidate rule is closed, then the rule fails to specify a necessary and sufficient condition for assertibility. Since sensitive belief is precisely as open as assertibility, the Sensitivity Rule is immune to this challenge.<sup>11</sup>

### 3 Lottery Assertions

One of the standard arguments for the Knowledge Rule begins with the observation that a speaker who asserts that someone's lottery ticket has lost just on the basis of her knowledge of the odds has done something epistemically improper.<sup>12</sup> As Williamson puts the point, one's listener 'will be entitled to feel some resentment when you later discover the merely probabilistic grounds for my assertion' (Williamson 2000, 246). On the assumption that I cannot know on the basis of the odds alone that your ticket has not lost, KR explains why this resentment is justified. And indeed it is commonly thought that I can't know, on this basis, that your ticket has lost. But this thought is not unproblematic. When I know the odds, my evidence that your ticket has lost is very strong indeed. Or, putting matters somewhat differently, the chance that I'm wrong given my evidence is exceedingly small. If such a small chance of error can prevent me from knowing that your ticket has lost, then we should treat similarly small chances of error as blocks to knowing in other cases. Fallibilist theories of knowledge commonly include a proviso that small chances of error are compatible with knowing.<sup>13</sup> My evidence for the proposition that my wife is now at work is that I just spoke with her on the phone and she told me that she was. That evidence is intuitively sufficient for me to know that she's there right now. But there's a chance that I'm wrong, given that evidence. Though it is not as easy or natural to quantify the chance

<sup>11</sup> What about an assertion of the conjunction: 'The wall is red, and not white and lit by red lights?' It seems to me that such an assertion would be just as bad as one of the second conjunct alone, but there is concern that SR gives the opposite verdict. The trouble is best approached in terms of the possible-worlds analysis of the conditional: if the closest possible worlds where the wall is not red are worlds where it's not white and lit by red lights, which they must be if they're to be worlds where I don't believe the wall is red, then the sensitivity conditional relevant to the conjunctive assertion looks to be true. Hence, SR does not predict that the conjunctive assertion would be improper.

In response to this concern, note that in asserting a conjunction, you thereby assert each conjunct. So a speaker is in a position to assert a conjunction only if she is in a position to assert each conjunct. (Thanks to [blind] for this point.) Hence, to properly assert that the wall is red and not white and lit by red lights, one must be in a position to assert that the wall is not white and lit by red lights. And for the reasons given above, one can sensitively believe the conjunction without sensitively believing that the wall is not white and lit by red lights, so you may satisfy the condition in SR with respect to the conjunction without being in a position to assert the conjunction.

Another possible response is to rely on the considerations raised in Sect. 4 below. Asserting the conjunction is liable to shift the salient similarity relation or expand the set of conversationally relevant possibilities, thus rendering the sensitivity conditional untrue. So even if, in any conversation where 'The wall is red' is assertible, the sensitivity conditional relevant to the conjunctive assertion is true, were a speaker to assert the conjunction she'd make the conversation one in which the sensitivity conditional relevant to the 'The wall is red' untrue.

<sup>12</sup> See Williamson 2000, Chap. 11. Lackey 2007, 618, gives cases where, she claims, lottery assertions are not improper. I think that in the cases she gives, urgent practical circumstances imply either that the ordinary rules governing assertion are suspended, or else violations are excused.

<sup>13</sup> For example, Fantl and McGrath defend a version of fallibilism on which 'you can know that P even if there is a non-zero epistemic chance for you that not-P' (Fantl and McGrath 2009, 11). See Dougherty and Rysiew 2009 for a similar characterization of fallibilism.



as it is in the lottery case, it seems at least as great as that your lottery ticket has won. The assumption that the chance of error in the lottery case robs me of knowledge may have pernicious skeptical consequences.

This is not an argument that I can know that the lottery ticket has lost on the basis of the odds, nor that fallibilists are committed to this claim; there may be relevant differences between the chance of error in the lottery case and the chance of error in other cases. The point is just that for fallibilists there is reason to be open-minded about the possibility that I can know that a lottery ticket has lost on the basis of the odds. Yet we needn't let open-mindedness about lottery knowledge lead us to think that it is epistemically permissible for me to assert that your ticket has lost on the basis of my knowledge of the odds. An argument that probability-based lottery knowledge is, in fact, possible would leave KR incapable of explaining the impropriety of lottery assertions; it would not show them to be epistemically proper. A fallibilist willing to accept the existence of lottery knowledge needs a different explanation for the epistemic impropriety of lottery assertions.

Prominent competitors to KR cannot easily explain the impropriety. Douven, for example, argues that rational credibility is the epistemic norm of assertion. It is plausible that when a subject's evidence makes a proposition as probable for her as the odds make the lottery proposition, the proposition is rationally credible for her. Thus, one might think that the rational credibility norm is in for counterexamining by the lottery case.<sup>14</sup> Douven instead denies that lottery propositions are rationally credible (Douven 2006, 258-259). The reason he gives is that conjoining the view that lottery propositions are rationally credible with the claim that rational credibility is closed under known entailment leads to a well-known paradox.<sup>15</sup> Yet, as he and Williamson have shown in other work, attempts to offer formal constraints on rational credibility that permit escape from paradox without denying closure have not been successful (Douven and Williamson 2006). Without an explanation for why lottery propositions are not rationally credible, the rational credibility theory of assertibility faces a serious challenge. So do accounts that link assertibility with justified or reasonable belief, since those relations are also plausibly thought to obtain between a subject and anything as probable on her evidence as the relevant lottery propositions are.<sup>16</sup>

The Sensitivity Rule, by contrast, yields a straightforward explanation for the impropriety of lottery assertions. When your belief that my ticket lost is based just on your knowledge of the odds, it's clear that if my ticket were not to have lost, you would still believe that it had. Your belief is not sensitive, and so by SR the proposition is not assertible. SR also explains the discrepancy between the lottery case and assertions that are manifestly proper, though about which the speaker has, given her evidence, a comparable or even greater chance of being wrong. When I believe P on the basis of evidence that does not entail P, there is a chance that I'm wrong about P, but not that my

<sup>14</sup> In this connection, see Dudman 1992 for a lottery-based attack on the idea that high subjective probability warrants assertion.

<sup>15</sup> See Kyburg 1961, 197.

<sup>16</sup> Stanley 2008's suggestion that either epistemic or subjective certainty is required for proper assertion doesn't face this challenge: whereas it is rather a stretch to say that a lottery belief is unjustified or unreasonable, it is no stretch at all to say that one is not certain that the ticket has lost. The worry for either version of the certainty rule is that it is too demanding. For there are many beliefs we find perfectly assertible that we would under even mild pressure concede are uncertain in both senses; for example, my belief about my bike's location. Stanley responds to this concern by giving a contextualist semantics for 'certain,' where what counts as the 'highest' degree of confidence of justification shifts with context.

belief is not sensitive. My wife's telling me that she's now at work does not entail that she is at work; nonetheless, if she weren't now at work, she would not have told me that she was. Given that my belief that she is at work is based on what she told me, if she were not at work I would not believe that she was. So that proposition is assertible, as far as SR is concerned, despite the fact that the chance that I'm wrong given my evidence is no less and possibly greater than the chance that I'm wrong about the lottery ticket.<sup>17</sup>

Unlike KR and its competitors, SR predicts the impropriety of lottery assertions without broaching any controversial epistemological questions or appealing to any sophisticated pragmatic machinery.<sup>18</sup> It is consistent with the Sensitivity Rule explanation to affirm that one can know that one's ticket has lost on the basis of the odds and that this is rationally credible and reasonable to believe. (And of course it is also consistent with the Sensitivity Rule explanation to deny these things as well.) The lottery proposition's unassertibility is due simply to the fact that the relevant subjunctive conditional is untrue.

#### 4 Assertibility and Context

Whether a speaker's epistemic position is sufficient for asserting that P fluctuates with various conversational factors. Suppose that you and I are strangers who've just met at the airport terminal, waiting to board our delayed flight. Though I didn't see my cats before leaving the house this morning, I did see them before bed and heard one of them howling in the night. It's not unusual for me not to see them in the morning, since they are usually napping in the attic. I confidently believe that they're at home, and (as a matter of fact) they are. Noticing cat hair on my coat, you strike up a conversation. Here are two directions our chat might go:

*Kennel.* You ask whether I kennel my cats while traveling. No, I tell you, I've tried in the past, but wrangling them into their carriers is such a pain, and they're always unhappy when they return home. Plus, I've got a cat-loving friend in town who's willing to come by and check on them every day when I travel. So, I conclude, 'My cats are at home.'

*Escape.* You tell me a story about the last time you had to catch an early morning flight: while you were bringing your bags to the car, one of your cats slipped out the back door without your notice. Luckily one of the neighbors discovered her mewling outside the next morning and was able to let her in with his spare key. But now you're always sure to check to make sure they're inside before leaving on a trip. What about me? Are my cats safe at home? 'Well,' I say, 'That's never happened to me before, and my cats are usually pretty lazy

<sup>17</sup> DeRose 1996 makes an analogous point about how the subjunctive conditional involved in the sensitivity condition can explain the difference between intuitive cases of 'knowledge' and 'no-knowledge,' even when there is no appreciable difference in the size of the chance of error in each. That paper anticipates a form of SR. DeRose's explanation presupposes that assertibility and knowledge are tightly linked—a link he there treats as a 'working hypothesis' and endorses in later work. One might also explore the link between assertibility and sensitivity without presupposing any connection between either and knowledge. That is the orientation of the present article.

<sup>18</sup> As an example of the latter, see Weiner 2005, who offers a Gricean explanation for the impropriety of lottery assertions, and Lackey 2007, who claims that improper lottery assertions might be understood as violations of a pragmatic norm against making assertions one has reason to believe will be misleading.

about getting out. So I'm confident they didn't escape this morning. But I didn't see them before I left, so I can't say for sure that they're at home.'

In the kennel conversation I am in a position to assert that my cats are at home. In the escape conversation, I am not. But my epistemic position with respect to that proposition is constant throughout: my evidence about the cats' whereabouts is the same in each case, as is my confidence in its truth. Merely entertaining the question of whether they have escaped without my notice does not alter my epistemic position. Rather, what shifts is what epistemic position is required for proper assertion. My position licenses an assertion in the first conversation, but not in the second. Assertibility is context-sensitive.

Friends of KR can accommodate the context-sensitivity of assertibility in three ways. First, they can accept that KR specifies only one necessary condition for assertibility; perhaps there is a distinct condition satisfied in the kennel conversation but not in the escape conversation. Second, they can wed KR to a contextualist or relativist theory of knowledge attributions: if the truth-conditions for 'S knows that P' fluctuate with context, then given KR we should expect assertibility to fluctuate with context as well.<sup>19</sup> And third, they can defend a revisionary theory of knowledge on which factors independent from a subject's epistemic position determine whether she is in a position to know.<sup>20</sup> If my epistemic position is sufficient for knowing that the cats are at home in the kennel conversation but not in the escape conversation, then given KR it is no surprise that the proposition should be assertible in the first conversation but not the second. Here I shall remain neutral on the prospects for any of these routes, except to point out that epistemologists skeptical about contextualism, relativism, and impurism about knowledge will also be skeptical about taking the second or third routes. Epistemological conservatives should hope for an account of assertibility that does not require endorsing a revisionary thesis about knowledge or the truth-conditions of knowledge attributions.<sup>21</sup>

There is very strong evidence that subjunctive conditionals are context-sensitive. If subjunctive conditionals are context-sensitive, then whether a speaker's belief satisfies the sensitivity condition will also vary with context. The Sensitivity Rule, then, should lead us to expect assertibility to be context-sensitive. In the rest of this section I'll discuss two different ways of thinking about the context-sensitivity of subjunctive conditionals and show how each implies that the variability of assertibility illustrated by the cat cases above provides indirect confirmation of the Sensitivity Rule.

<sup>19</sup> See, for example, DeRose 2002, 2009 and Schaffer 2008 for contextualist versions of the Knowledge Rule.

<sup>20</sup> Hawthorne 2004; Stanley 2005, and Fantl and McGrath 2009 are prominent works defending 'impurist' theories of knowledge on which the strength of position required to know varies with the subject's situation. Stanley does not use impurism about knowledge to explain the variability in assertibility, though. He instead proposes a certainty norm of assertion and suggests accounting for the variability in assertibility with the context-sensitivity of 'certain' (Stanley 2008).

<sup>21</sup> Those who reject the Knowledge Rule in favor of another epistemic norm face a similar trilemma. If the alternative does not explicitly specify a context-sensitive factor, then the contextual variability of assertibility implies that either the rule does not give a sufficient condition for assertibility, the epistemic relation specified in the rule requires a contextualist or relativist semantics, or the correct understanding of the epistemic relation is somehow analogous to the view of knowledge defended by impurists.

I do not wish to argue against any of these possibilities. The point is just to show that the Sensitivity Rule yields a simple and straightforward explanation for each of the elements of assertibility discussed, including (as the rest of this section will show) its contextual variability. That is evidence that the Sensitivity Rule is the only epistemic norm governing assertion. But it is compatible with more ecumenical views, too.

Supposing that subjunctive conditionals are context-sensitive, whose context matters when assessing whether the speaker satisfies the condition specified in the Sensitivity Rule? The natural answer is that it's the speaker's context. The alternative is that it's the context in which the propriety of her assertion is being assessed. But this alternative suggests a form of relativism about assertibility that seems unpromising. So the Sensitivity Rule, phrased to make the context-sensitivity of the relevant subjunctive conditional explicit, is:

(SR) S should assert that P in C only if S believes that P and 'S would not believe that P if P were false' is true with respect to C.<sup>22</sup>

What features of the context are subjunctive conditionals sensitive to?

I'll use  $P \rightarrow Q$  to abbreviate 'If P were the case, Q would be the case.' We begin with the core of Lewis's familiar analysis of counterfactuals: that  $P \rightarrow Q$  is true at  $w$  just in case at all the closest possibilities to  $w$  where P is the case, Q is the case.<sup>23</sup> The 'closeness' of a possibility is understood in terms of similarity. Making that understanding explicit the Lewisian core becomes:  $P \rightarrow Q$  is true at  $w$  just in case at the possibilities most similar to  $w$  where P is the case, Q is the case. One obvious way to build context-sensitivity into the framework is to let context fix and rank the relevant respects of similarity. Lewis clearly endorsed such context-sensitivity in his original presentation of the theory; there, 'closeness' was a function of overall similarity, and Lewis happily acknowledged that what factors are relevant to 'overall similarity' is a matter that varies with context.<sup>24</sup> In later work, Lewis gave a more specific articulation of the similarity relation relevant to the assessment of counterfactuals (Lewis 1986). But the terms involved in Lewis's later proposal are themselves deeply context-sensitive, e.g., he there requires that the most similar possibilities contain no 'big, widespread, diverse violations of law.' Moreover, merely ranking various considerations doesn't settle the question of how much weight to give to each. So on Lewis's later view as well, different possibilities will count as most similar in different contexts.<sup>25</sup>

<sup>22</sup> Some theorists, notably Dorothy Edgington, deny that subjunctive conditionals have truth-conditions (see Edgington 2008). I have not yet completely thought through the consequences of this denial for the present proposal, but there is a serious *prima facie* worry. The traditional move for those who deny that conditionals have truth conditions is to account for their workings in terms of assertibility conditions. But if conditionals have only assertibility conditions, then the condition on the right-hand side of the Sensitivity Rule must be understood as being satisfied just when a particular conditional is assertible. One might reasonably complain that a general account of assertibility in terms of the assertibility of a particular conditional is problematically circular. For the rest of this section I'll side with the majority in assuming that subjunctives are truth-conditional.

<sup>23</sup> Lewis 1972.

<sup>24</sup> See, for example, Lewis 1972, pp. 91-92, where he writes, 'Overall similarity consists of innumerable similarities and differences in innumerable respects of comparison, balanced against each other according to the relative importances we attach to those respects of comparison. [...] The truth conditions for counterfactuals are fixed only within rough limits; like the relative importances of respects of comparison that underlie the comparative similarity of worlds, they are a highly volatile matter, varying with every shift of context and interest.'

<sup>25</sup> Lewis was untroubled by this residual context-sensitivity: 'Plenty of unresolved vagueness remains, of course, even after we have distinguished the four sorts of respect of comparison and ranked them in decreasing order of importance' (Lewis 1986, 48). And earlier in the paper: 'Counterfactuals are infected with vagueness, as everyone agrees. Different ways of (partly) resolving the vagueness are appropriate in different contexts. Remember the case of Caesar in Korea: had he been in command, would he have used the atom bomb? Or would he have used catapults? It is right to say either, though not to say both together. Each is true under a resolution of the vagueness appropriate to some contexts' (Lewis 1986, 34).

Adopting a Lewisian view, then, we should expect the truth-conditions of the sensitivity conditional to vary with the respects of similarity most significant in the speaker's context; given the Sensitivity Rule, a proposition's assertibility by a speaker will vary in the same way. We can account for the variability in the cases above by positing a shift in the salient similarity relation. The respects of similarity most significant in the kennel conversation are not those most significant in the escape conversation. In the kennel conversation, we are focused on our plans for ensuring that our cats are cared for while we travel. The similarity relation made salient is thus one in which at the closest possible worlds to actuality, our cats are cared for in the way we plan. Privileging this aspect of similarity, the closest possibilities where my cats are not at home are those where I have made a different plan for them. Since, in those possibilities, I don't falsely believe that they're at home, 'I would not believe the cats were at home if they weren't' is true with respect to the context of the kennel conversation. In the escape conversation, by contrast, my interlocutor's question leads us to consider ways in which our plans may be foiled, and hence ways in which we could be wrong about how things are with our cats. The similarity relation thus made salient weights factors involving my evidence and method of belief formation more heavily. Given the nature of my evidence, that weighting rates among the possibilities most similar to actuality where my cats are not at home some where I mistakenly believe that they are. So with respect to the context of the escape conversation 'I would not believe the cats were at home if they weren't' is not true.

Variability in the similarity relation is not the only possible source of context-sensitivity in subjunctive conditionals. Suppose that there is a single determinate similarity relation relevant to the evaluation of a subjunctive conditional that orders every possibility in terms of its closeness to actuality. Their truth-conditions may also vary with which close antecedent possibilities are *relevant*. Here is an example. As we're eating dinner, you smell gas; you go to check and realize that you've left the stove top on and the kitchen has been filling with gas for hours. After you turn it off, open the windows, and clear the room, you say to the rest of us, 'It's a good thing no one struck a match! If we had, there would have been an explosion!' What you've said seems clearly true, and you can reasonably expect your dinner guests to nod in mutual relief. Still, you might have a nit-picker among them: 'Well, suppose the match had been a dud,' she says, '*then* there wouldn't have been an explosion.' If you accommodate her suggestion, the context shifts in such a way that it becomes difficult to truly utter your initial conditional again. The following conjunction is infelicitous: 'You're right; if we'd struck a match that was a dud, there wouldn't have been an explosion. But if we'd struck a match, there would have been an explosion.'<sup>26</sup>

The infelicity suggests a shift in truth-conditions.<sup>27</sup> We can account for this shift without supposing that the similarity relation has shifted. Your guest's

<sup>26</sup> The conjunction can be forced through, but only by making it clear that the possibility that the match is a dud is irrelevant or somehow ruled out: 'That's true, if we'd struck a dud, there wouldn't have been an explosion. But there are no duds among my matches, so if we'd struck a match, there would have been an explosion.' So long as the dud possibility, once mentioned, isn't treated as irrelevant or ruled out, another utterance of the original conditional will be intuitively untrue.

<sup>27</sup> Some will wish to say that your guest has pointed out a consideration that shows the original utterance to have been false. This is a rather severe judgment. More importantly, it suggests a path that, if consistently taken, would lead to the conclusion that the mundane subjunctive conditionals we utter and believe with confidence are typically false, since most can be rendered intuitively false by the mention of unusual P&~Q possibilities just like the dud scenario.

remark has the effect of expanding the set of relevant match-striking possibilities. Before she mentioned it, the dud possibility was not relevant. In accommodating her remark, we made that possibility relevant. Since the closet dud possibility is not one where striking leads to an explosion, once it becomes relevant not all of the closest striking possibilities contain explosions. So her intervention shifts the truth-value of the conditional by making a new close possibility relevant.

The effect of your guest's remark suggests a modification of the Lewisian core to say that  $P \rightarrow Q$  is true at  $w$  just in case at all the closest *relevant* possibilities where  $P$  is true,  $Q$  is true.<sup>28</sup> Of course, much needs to be said about the relevant notion of relevance. But even before anything more is said, it should be clear that this modification predicts the needed variation in the truth-value of the sensitivity conditional between the cat conversations. In the kennel conversation, the possibility that my cats have escaped unnoticed is not relevant, and so does not stand in the way of my belief's satisfying the sensitivity condition. In the escape conversation, though, that possibility is relevant, and so does stand in the way. Given SR, then, we have the needed explanation for the variation in assertibility. Speaking more generally, the modified Lewisian view together with SR predicts that as possibilities in which a speaker falsely believes that  $P$  become conversationally relevant, she will tend to lose her ability to assert  $P$  with epistemic propriety. This seems to be just how assertibility does vary. Given the independent plausibility of the modified Lewisian view, then, SR is again confirmed.<sup>29</sup>

<sup>28</sup> This echoes the account of subjunctives given in Lycan 2001, as well as Jonathan Ichikawa's recent defense of a theory of counterfactuals along similar lines (Ichikawa 2011). Ichikawa's proposal is stronger, since he drops talk of closeness altogether. I worry about cases where a  $P \& \sim Q$  possibility is relevant but not as close as any  $P \& Q$  possibility, since Ichikawa's proposal would seem to render  $P \rightarrow Q$  false in such cases. You can truly assert things like: 'In our high-crime neighborhood the possibility of car theft is always relevant, which is why we've taken extra steps to ensure that our car is safe. If it weren't in the garage we wouldn't believe that it was.' But this worry is not very deep, and Ichikawa's view is elegant. For other proposals in this vein, see Brogaard and Salerno 2008 and Gillies 2007.

<sup>29</sup> A reviewer points out that if subjunctives are context-sensitive, certain inference patterns typically thought to be fallacious may, in fact, be formally valid; in particular, antecedent strengthening, contraposition, and hypothetical syllogism (see Brogaard and Salerno 2008). I am not convinced that antecedent strengthening can be saved: Brogaard and Salerno's defense of its validity appears to depend upon the dubious assumption that whenever there are no  $Q$ -worlds among the contextually relevant closest  $P$ -worlds, there are no contextually relevant closest  $P \& Q$ -worlds (see their discussion on p. 43). I see no reason to accept this assumption. That the closest relevant worlds where  $P$  is true are all worlds where  $Q$  is false does not imply that there are no closest relevant worlds where  $P$  and  $Q$  are both true; all it implies is that no  $P \& Q$ -worlds are among the closest relevant  $P$ -worlds.

But suppose Brogaard and Salerno are correct. Then, with respect to any context where 'If the cats were not at home, I would not believe that they were' is true, so is 'If the cats were not at home and had escaped unnoticed, I would not believe that they were at home.' Nonetheless, SR still explains why 'The cats are at home' is unassertible in the escape conversation despite being assertible in the kennel conversation. In the escape conversation, a different similarity relation is salient (or different possibilities are conversationally relevant), and so with respect to the context of *that* conversation, 'If the cats were not at home, I would not believe that they were at home' is not true. Moreover, the Brogaard and Salerno-inspired reason for thinking that 'If the cats were not at home and had escaped unnoticed, I would not believe that they were at home' is true with respect to the kennel conversation is that there are no conversationally relevant possibilities where the antecedent is true, i.e., that it is vacuously true. To avoid concerning ourselves with such cases, we could rephrase SR to require that the conditional be *non-vacuously* true.

## 5 Two Challenges for the Sensitivity Rule

There are standard objections to the sensitive belief requirement on knowledge that could be used to make analogous complaints against SR. In this section, I'll discuss the two strongest objections.

In some cases a belief that P appears not to be sensitive because if the belief were false, the subject would have a different and misleading basis for believing that P, and would falsely believe P on that basis. But in such cases, insensitivity seems compatible with both knowledge and assertibility. The Nozick's grandmother case is the most famous example: seeing the grandson looking healthy, the grandmother believes that he is well. Yet if he had been ill, her children would have kept him home and told her on the phone that he was well. Nozick said that the grandmother's belief was not sensitive and modified the sensitivity requirement for knowledge in response to the counterexample. If Nozick is right that the grandmother's belief is not sensitive, then we have an objection to SR as well, since she seems perfectly well positioned to assert.

One response is to hold that the grandmother's belief is, indeed, sensitive. Focusing our attention on the fact that she's looking right at her healthy grandson, it does not seem incorrect to say that if the grandson were not well she would not believe that he was.<sup>30</sup> This can be explained by supposing that the conditional's truth-conditions shift with the salient similarity relation. When epistemic factors such as the grandmother's evidence and method of belief formation are weighted heavily in assessing similarity, at the closest possibilities where the grandson is ill the grandmother's belief concerning his health has a similar etiology; if he were not well he would not look well, and so she would not then believe him to be well. On the other hand, when the salient similarity relation is one on which grandmother's epistemic situation is of little significance, what we know about her children's plans leads us to treat some of the closest possibilities where he is ill as possibilities where she is misinformed.

This appeal to context-sensitivity does not fully disarm the objection, however. Suppose that the grandmother finds out about her children's plans and shares them with her friends. In telling the story, she could truthfully and appropriately utter, 'Since my children are schemers, if he were ill, I'd still believe he was well. Nonetheless, he is well.' SR appears to have a counterexample. We need to explain how the similarity relation salient in the context can render the sensitivity conditional false without undermining the grandmother's ability to assert that her grandson is well.

Candidate explanations are easy to come by. One involves saying that the similarity relation needed in the assessment of the sensitivity conditional is more or less fixed across contexts, and always puts heavy weight on considerations concerning evidence and method of belief formation. Then SR would need supplementation with a more precise account of the similarity relation relevant to assessing the epistemic propriety of the speaker's assertion. A closely related strategy is to explicitly mention the considerations relevant to sensitivity in the condition itself. We could, for example, follow Nozick by incorporating the speaker's method of belief formation into the sensitivity condition.<sup>31</sup> This would involve replacing the

<sup>30</sup> See DeRose 1995 for a similar suggestion for handling grandmother-style cases.

<sup>31</sup> See Nozick 1981, 197ff.

requirement in SR with a requirement that the speaker's belief be based on a sensitive method, something like: 'If P were false and S were to use her actual method of belief formation to arrive at a belief about whether or not P, then S wouldn't believe, via that method, that P.' Or we could adopt a more recent proposal by Troy Cross and instead of sensitive belief, require sensitive grounds for the belief; something like: 'S's belief that P is grounded in some Q such that if P were false, S would not believe that Q.'<sup>32</sup> The grandmother's belief satisfies both the sensitive method and sensitive ground requirements, and so such modifications would neutralize the force of counterexamples where the speaker would have a different basis for her belief if what she asserts were false.<sup>33</sup>

The second standard objection to a sensitivity requirement on knowledge begins with the claim that our inductive beliefs are often not sensitive. Ernest Sosa's influential example involves dumping trash down a chute (Sosa 1999, 145). After I dump the trash down the chute, I believe that it's in the dumpster at the bottom of the chute. This belief isn't based on my having seen the trash enter the dumpster but on my inductive evidence concerning what happens when the trash is dumped down the chute, together with my memory of so dumping it. Sosa says that my belief is not sensitive: if the trash were not in the dumpster, that would likely be because it was stuck in the chute, and so I would likely still believe that it was in the dumpster. If Sosa is correct, then this is an counterexample to SR, since I seem well positioned to assert that it's in the dumpster. Appeals to methods and grounds don't help here; if the dumpster belief is not sensitive, that's because my method and grounds for belief are not sensitive.

But such cases can be handled, provided we bear in mind the context-sensitivity of subjunctive conditionals. Indeed, a more careful look at Sosa's case reveals that it confirms SR. First, note that circumstances can easily evolve in which someone who believes on inductive grounds that the trash is in the dumpster is not in a position to assert that the trash is in the dumpster. Consider the following conversations:

*Hallway.*

You: 'Is the trash in the dumpster? Or is it still out in the hallway?'

Me: 'The trash is in the dumpster.'

*Chute.*

You: 'Is the trash in the dumpster? Or is it caught in the chute?'

Me: 'The trash is in the dumpster.'

My recent memory of dumping the trash down the chute, together with my evidence about what typically happens to trash so dumped, puts me in a position to

<sup>32</sup> See Cross 2010, p. 49. Ichikawa 2011 gives a closely related gloss on sensitivity, where a belief that P based on evidence E is sensitive just in case the subject would not believe that P on the basis of E if P were false (Ichikawa 2011, 302). Both proposals are foreshadowed by Dretske's notion of a conclusive reason (Dretske 1971), though see Cross 2010, fn. 8, for an important difference.

<sup>33</sup> Either modifications would also block other standard 'counterexamples' to sensitivity conditions such as Williamson's 'I'm not a brain in a vat who's hallucinating that he's a mountain climber' (Williamson 2000, 158) and Hawthorne's no-ticket lottery case (Hawthorne 2004).



assert that the trash is in the dumpster in the hallway conversation but not in the chute conversation. This is just another example of the context-sensitivity of assertibility. And it can be diagnosed in the same way that the example of context-sensitivity discussed in the previous section was: in the second conversation, but not the first, a particular possibility where I falsely believe that the trash is in the dumpster is relevant. With respect to the first conversation, where that possibility is not relevant, the sensitivity conditional is intuitively true: given my basis for belief, I could rightly insist that of *course* I wouldn't believe that the trash was in the dumpster if it weren't there. After all, I just dumped it down the chute! But the relevance of the chute possibility, together with the nature of my evidence, renders the sensitivity conditional false with respect to the second conversation. And in that conversation I can't assert that the trash is in the dumpster. Assertibility fades along with the truth of the relevant subjunctive conditional. This strategy promises to handle most potential counterexamples involving assertions based on inductive evidence. As long as possibilities where  $p$  is false but the speaker has her actual inductive evidence for  $p$  are not relevant, the path is open for the sensitivity conditional to come out true and for the assertion to be epistemically permissible. If some such possibilities come to be treated as conversationally open, the sensitivity conditional is unlikely to be true and the speaker unlikely to be in a position to assert.<sup>34</sup>

## 6 An Ecumenical Conclusion

Unlike the Knowledge Rule and other popular candidate norms of assertion, the Sensitivity Rule underwrites straightforward explanations of three key aspects of assertibility: why assertibility is open under known entailment, why lottery assertions (and other assertions based on manifestly probabilistic evidence) are epistemically improper, and why assertibility fluctuates with the focus of conversation. In this conclusion, I'll briefly describe three ways of understanding the role SR plays in the full account of assertibility.

The first, most conservative role for SR is as an additional necessary condition on assertibility over and above whatever other epistemic norm or norms govern the practice. A proponent of another rule of assertion may endorse SR as well. There are

<sup>34</sup> This strategy may have limits. Consider Vogel's Heartbreaker case (Vogel 1990). Sixty pros are in the golf tournament; the Heartbreaker is a particularly tough par five hole. Knowing about golf and statistics, you believe, before the tournament, that not all 60 players will get a hole-in-one on the Heartbreaker. Is your belief sensitive? There is pressure to say no, but that the belief is nonetheless assertible. I think we should resist the pressure to deny that the belief is sensitive. One way to do so is to say that at the closest possibilities where all 60 players get a hole-in-one, golf is quite a bit different than it actually is; given how incredibly unlikely the antecedent is, the conditional invites a backtracking interpretation. (It's not easy to conceptualize just how unlikely the antecedent is: according to *Golf Digest*, the odds of a tour player making a hole-in-one on just an ordinary par three hole are around 1 in 3,000. Given those chances, the odds of all 60 players getting a hole-in-one are approximately one in  $2.4 \times 10^{209}$ ; for comparison, the universe is approximately  $4.3 \times 10^{26}$  nanoseconds old.) Moreover claims like 'If all 60 players were going to get a hole-in-one on the Heartbreaker, golf would be very different than it actually is—perhaps golf balls would have homing devices, or the holes a meter in diameter' hardly seem false, and can quite easily seem true. Since your basis for belief includes your knowledge of the way golf actually is, this may be enough to make your belief satisfy the condition in the fully articulated Sensitivity Rule. (For a different and more concessive way of handling the Heartbreaker, see Cross 2007.)

aspects of assertibility that SR does not obviously explain. For example, KR easily accounts for the impropriety of 'Moorean' assertions such as 'P but I don't know that P' and neatly explains our practice of challenging assertions by asking how a speaker knows the proposition she has asserted.<sup>35</sup> If SR cannot account for either of these aspects of assertibility, then we have reason to think that it is not the only rule relevant to assertibility.<sup>36</sup>

SR may instead play a less direct role in the account of assertibility. Suppose that a proponent of KR endorses a contextualist semantics for knowledge attributions. A contextualist version of KR could explain both the openness and contextual variability of assertibility. For example, one may say that with respect to an ordinary context Donna's epistemic position is sufficient for making it true to say that she knows that the walls are red, but that with respect to a context where the possibility of the wall's being white and deceptively lit by red lights is salient, her epistemic position is not sufficient to make this true. A similar story could explain the effects of conversational focus on assertibility. Such an explanation needs supplementation with an account of the mechanisms by which the truth-conditions of a knowledge attribution vary with context. SR would then provide an important constraint on how such an account should go: the mechanisms governing the semantics of knowledge attributions, whatever they are, must imply that a belief will count as 'known' with respect to a particular context only if it satisfies the relevant sensitivity conditional with respect to that context.<sup>37</sup>

Finally, SR might articulate the fundamental condition for assertibility. This is a bolder approach. I have not argued for this bold claim here. But the payoff for such boldness is theoretical simplicity. If SR can help us understand not just the aspects of assertibility discussed here but explain why any assertion has the epistemic status that it does, then we would have reason to regard it as the only rule governing assertion, and perhaps even the constitutive norm of assertion.

**Acknowledgements** Earlier versions of this paper were given at Northwestern University, Wayne State University, and the Central States Philosophical Association. Thanks to audiences and commentators on those occasions, especially Jamin Asay, Larry Powers, Baron Reed, and Bruce Russell. The idea for the paper grew out of my dissertation, and I thank Troy Cross and Keith DeRose for their help. Justin Dealy and Peter van Elswyk read drafts and gave helpful comments, as did two anonymous referees for this journal. Finally, I owe Mylan Engel a great debt of gratitude for insightful discussions about these issues and for reading through several complete drafts of the paper.

<sup>35</sup> See Williamson 2000 for further discussion, as well as recent work by Benton (Benton 2011, 2012).

<sup>36</sup> Though I don't have space to adequately defend the claim here, I think that both phenomena can be explained by SR. First: SR implies that by asserting P you represent yourself as sensitively believing that P; by asserting that you don't know that P you make conversationally relevant any close possibilities where your belief that P is false. When a close possibility that you falsely believe that P is conversationally relevant, 'If P were false, I wouldn't believe that P' is liable to be untrue. So in asserting 'P but I don't know that P' you represent yourself as having a sensitive belief, but then go on to say something that undermines that representation. Hence, there is reason to think that SR itself accounts for the impropriety of 'Moorean' assertions. Second: The question 'How do you know that P?' prompts the addressee to furnish her grounds for belief, or to describe her method of belief formation. A speaker's grounds and methods for believing P are directly relevant to whether her assertion that P satisfies the sensitivity condition, and so given SR it is no surprise that assertions may be challenged in this way.

<sup>37</sup> Taken in this direction, the Sensitivity Rule would mesh nicely with the 'Rule of Sensitivity' proposed by DeRose 1995, whereby the standards for 'knowledge' in a context should rise, if need be, so that a subject's belief meets them only if her belief is sensitive.

## References

- Becker, K., & Black, T. (2012). *The sensitivity principle in epistemology*. Cambridge: Cambridge University Press.
- Benton, M. (2011). Two more for the knowledge account of assertion. *Analysis*, 71, 684–687.
- Benton, M. (2012). Assertion, knowledge, and predictions. *Analysis*, 72, 102–105.
- Brogaard, B., & Salerno, J. (2008). Counterfactuals and context. *Analysis*, 68, 39–46.
- Coffman, E. J. (2011). Does knowledge secure warrant to assert? *Philosophical Studies*, 154(2), 285–300.
- Cross, T. (2007). Comments on Vogel. *Philosophical Studies*, 134(1), 89–98.
- Cross, T. (2010). Skeptical success. *Oxford Studies in Epistemology*, 3, 35–62.
- David, M., & Warfield, F. (2008). Knowledge-closure and skepticism. In Q. Smith (Ed.), *Epistemology: New essays*. Oxford: Oxford University Press.
- Davis, W. (1998). *Implicature: Intention, convention, and principles in the failure of Gricean theory*. Cambridge: Cambridge University Press.
- DeRose, K. (1995). Solving the skeptical problem. *The Philosophical Review*, 104(1), 1–52.
- DeRose, K. (1996). Knowledge, assertion, and lotteries. *Australasian Journal of Philosophy*, 74(4), 568–580.
- DeRose, K. (2002). Assertion, knowledge, and context. *The Philosophical Review*, 111, 167–203.
- DeRose, K. (2009). *The case for contextualism*. Oxford: Oxford University Press.
- DeRose, K. (2010). Insensitivity is back, baby! *Philosophical Perspectives*, 24, 161–187.
- Dougherty, T., & Rysiew, P. (2009). Fallibilism, epistemic possibility, and concessive knowledge attributions. *Philosophy and Phenomenological Research*, 78(1), 123–132.
- Douven, I. (2006). Assertion, knowledge, and rational credibility. *The Philosophical Review*, 115(4), 449–485.
- Douven, I., & Williamson, T. (2006). Generalizing the lottery paradox. *British Journal for the Philosophy of Science*, 57, 755–779.
- Dretske, F. (1971). Conclusive reasons. *Australasian Journal of Philosophy*, 81, 411–430.
- Dudman, V. H. (1992). Probability and assertion. *Analysis*, 52, 204–211.
- Edgington, D. (2008). Counterfactuals. *Proceedings of the Aristotelian Society*, 108, 1–21.
- Fantl, J., & McGrath, M. (2009). *Knowledge in an uncertain world*. Oxford: Oxford University Press.
- Gillies, A. (2007). 'Counterfactual scorekeeping'. *Linguistics and Philosophy*, 30, 329–360.
- Hawthorne, J. (2004). *Knowledge and lotteries*. Oxford: Oxford University Press.
- Hawthorne, J. (2005). The case for closure. In M. Steup & E. Sosa (Eds.), *Contemporary debates in epistemology*. Malden: Blackwell.
- Ichikawa, J. (2011). Quantifiers, knowledge, and counterfactuals. *Philosophy and Phenomenological Research*, 82(2), 287–313.
- Kripke, S. (2011). *Philosophical troubles: Collected papers, vol. 1*. Oxford: Oxford University Press.
- Kvanvig, J. (2006). Closure principles. *Philosophy Compass*, 1(3), 256–267.
- Kyburg, H. (1961). *Probability and the logic of rational belief*. Middletown: Wesleyan University Press.
- Lackey, J. (2007). Norms of assertion. *Noûs*, 41(4), 594–626.
- Lewis, D. (1972). *Counterfactuals*. Malden: Blackwell.
- Lewis, D. (1986). Counterfactual dependence and time's arrow. *Noûs*, 13(4), 455–476.
- Lycan, W. (2001). *Real conditionals*. Oxford: Oxford University Press.
- Nozick, R. (1981). *Philosophical explanations*. Cambridge: Harvard University Press.
- Pelling, C. (2012). Assertion and safety. *Synthese*. doi:10.1007/S11229-012-0223-7.
- Schaffer, J. (2008). Knowledge in the image of assertion. *Philosophical Issues*, 18(1), 1–19.
- Sosa, E. (1999). How to defeat opposition to Moore. *Philosophical Perspectives*, 13, 141–153.
- Stanley, J. (2005). *Knowledge and practical interests*. Oxford: Oxford University Press.
- Stanley, J. (2008). Knowledge and certainty. *Philosophical Issues*, 18, 35–57.
- Vogel, J. (1990). Are there counterexamples to the closure principle? In M. Roth & G. Ross (Eds.), *Doubting: Contemporary perspectives on skepticism*. Dordrecht: Kluwer.
- Weiner, M. (2005). Must we know what we say? *The Philosophical Review*, 114(2), 227–251.
- Williamson, T. (2000). *Knowledge and its limits*. Oxford: Oxford University Press.