

*Assertion, Sincerity, and Knowledge*

EDWARD S. HINCHMAN

University of Wisconsin-Milwaukee

**Introduction**

Lottery cases (“Your ticket’s a loser,” asserted with evidence only of its low probability of winning) and Moore’s paradox (e.g. “It’s raining but I don’t believe it”) appear to support the knowledge account of assertion, according to which one should assert only what one knows.<sup>1</sup> Since Susan does not *know* that Andrew’s ticket has lost, she should not assert that it has lost. Since asserting that it’s raining commits one to knowing that it’s raining, Sam should not go on to assert that he fails to believe it, since if that is true then he fails to satisfy a necessary condition on knowing it.<sup>2</sup> The knowledge account thus seems well positioned to explain the error or incoherence in these assertions.

This paper preserves an emphasis on knowledge but nonetheless presents grounds for an alternative explanation. My alternative approach divides the explanandum, explaining the *error* in lottery and Moorean assertions with one move (sections II and V) and the deeper *incoherence* in lottery assertions with another (sections III and IV). By ‘error,’ I mean the species of speech-active infelicity that J. L. Austin called an illocutionary ‘abuse.’<sup>3</sup> By ‘incoherence,’ I mean not a directly normative but an expressive feature of the speech act: when an assertion is incoherent, in this respect, it necessarily gives prima facie evidence of confusion. In sections III and IV we’ll explore how lottery assertions distinctively represent the speaker as manifesting a confused state of mind. Of course, the speaker may not be actually confused, and the assertion may not, given special background conditions, actually manifest confusion. The point is that lottery assertions necessarily give prima facie evidence that the speaker is confused in a way that bears on her understanding of what she’s doing in making the assertion, even if that evidence is defeated by special background conditions. As we’ll see in section V, Moorean assertions are not necessarily incoherent in this respect: they do not necessarily give prima facie evidence of such confusion. To the good methodological question why work on assertion should pay special attention to lottery cases and Moore’s paradox, this is my answer: the cases highlight different aspects of assertion’s epistemic nature. Moorean assertions distinctively reveal the epistemic nature of the illocutionary norm informing assertion, whereas lottery assertions distinctively reveal a fundamentally epistemic species of incoherence that goes beyond mere illocutionary error.

Here, in outline, is how the two sides of my argument fit together. The species of *error* at issue derives from a respect in which lottery and Moorean assertions are

uninformative: the speaker is not being appropriately responsive to her addressee's epistemic needs, thereby failing to meet the core obligation that she incurs in addressing her assertion to this hearer. My argument for this claim pays special attention to how what the addressee needs from the speaker depends on the context in which the addressee would come to believe what the speaker asserts—I'll call this the addressee's 'doxastic context.' I'll emphasize how the addressee's context may differ from the speaker's doxastic context (that is, from the context in which the speaker would believe what she asserts). Beyond the error, the species of *incoherence* at issue derives from a deeper respect in which a lottery assertion is uninformative: though the speaker does not aim to deceive, it is difficult to see how her assertion could express any judgment she has made—or, as we'll see, *would* relevantly make. In each respect, error and incoherence, it is difficult to see how she appropriately 'stands behind' her assertion—both what she asserts and her act of asserting it—though, again, there is no reason to suspect her of lying. But the uninformativeness of a lottery assertion creates an appearance not merely of error but of confusion: the speaker judges that *p* while quite transparently lacking epistemic authority to inform anyone whether *p*.

My diagnosis of the incoherence in lottery assertions yields a knowledge account not of assertion but of judgment. (Strictly speaking, as we'll see when we turn to Moorean assertions, it yields a 'conclusive warrant' account of judgment, but for now there's no harm in using the pithier phrase—as long as we remember that I'm going to emphasize the conclusiveness of the warrant, not the mental state of knowing.)<sup>4</sup> If the speaker transparently lacks authority to inform *anyone* whether *p*—to give anyone her word that *p*—then it's hard to see how she could without confusion count as judging that *p* or as being disposed to judge that *p*.<sup>5</sup> And without judging that *p* or at least being disposed to judge that *p* from her addressee's perspective (on her own evidence)—as we'll see, the latter is the interesting case—it's hard to see how she could without confusion intend to inform her addressee that *p*.

Rival accounts, including the knowledge account of assertion, make a misdiagnosis because they misconceive the illocutionary nature of assertion. In particular, they misconstrue the illocutionary norm requiring that an assertion be appropriately informative. The trick is to see how lottery cases and Moore's paradox provide support for a thesis that I'll defend independently: that making an informative assertion is not simply a matter of asserting what you believe, or even what you could reasonably believe. The thesis undermines a presupposition of rival accounts: that the ideal of sincerity imposes a norm on assertion, and that the norm requires that you assert only what you believe, or at least what you could reasonably believe.<sup>6</sup> We can easily construct cases that force us to abandon one of those conjuncts. The cases reveal that sincerity either does not impose a norm on assertion, or does not require asserting only what you believe or could reasonably believe.

I'll begin by developing that puzzle about sincerity, without taking a position on which disjunct we should adopt. However we resolve to use the term 'sincere,' we can thereby see that the illocutionary norm governing informative assertion does not require either that you believe or that it be reasonable for you to believe

what you assert but instead requires that you help your audience believe what they are by-*your*-lights warranted believing in *their* doxastic context, a context that may differ from your own. It is this moral that explains the error in lottery and Moorean assertions, since what's distinctive of them is that this dimension of interlocutory assistance has gone missing. The moral also reveals the incoherence in a lottery assertion, wherein the presumption that you might assist anyone in forming a warranted belief represents you as confused about what such assistance would require of you.

The emerging emphasis on the epistemic needs of your addressee, as opposed to others who might overhear your assertion, will bring us back to the important difference between lottery assertions and Moorean assertions that I've already emphasized. My key claim will be that a lottery assertion necessarily gives prima facie evidence that you're confused to anyone—whether addressee or overhearer—who observes you make the assertion, whereas a Moorean assertion does not give such evidence to an overhearer (even if it does to the addressee). Every observer, addressee and overhearer alike, will have grounds for accusing you of illocutionary error in making a Moorean assertion, but the assertion does not necessarily give prima facie evidence of confusion to the overhearer. Why not? Simply because, as we'll see in the final section, a Moorean assertion, unlike a lottery assertion, may manifest a perfectly unconfused set of attitudes toward the epistemic needs of those who may hear or overhear your assertion.

### I. A Puzzle about Sincerity

My argument thus begins from a puzzle about the bearing of sincerity on a core illocutionary norm governing assertion. I'll argue that two difficult-to-shake assumptions cannot be jointly sustained when we consider several rather ordinary kinds of case that I'll sketch. My novel approach to lottery and Moorean assertions applies to those much-discussed cases a codification of the moral that I'll draw from these new cases.

The two assumptions concern the nature of sincerity and its bearing on the illocutionary nature of assertion. On the one hand, an ideal of sincerity figures in how we tend to conceptualize the ideals of truthfulness and informativeness that lend force to the illocutionary norms governing the practice of assertion. I mean truthfulness and informativeness in the aretaic sense: not merely *happening* to speak the truth or to inform, but doing so in a way that manifests requisite virtues.<sup>7</sup> On the other hand, we tend to assume that a sincere assertion necessarily gives voice to what the speaker believes (or, at least, to what she believes it would be reasonable for her to believe, if for some non-epistemic reason she doesn't believe it). Putting the assumptions together, we get the result that a speaker who asserts what she does not believe (or believe she could reasonably believe) must be asserting untruthfully or uninformatively.<sup>8</sup> Either way, we assume, such an assertion violates a core norm governing the practice of assertion. Such a speaker should, we assume, be criticized *qua* asserter—even if her assertion was in some other respect permissible or praiseworthy.

I put the point in terms of the twin ideals of truthfulness and informativeness because the cases I'm about to sketch might appear to reveal a tension between those ideals and thus between the cognate norms. The apparent tension marks two ways we might put the ideals or norms into relation. We might formulate the norms of truthfulness and informativeness loosely as follows:

(T) Say what you think.

(I) Say what will be helpful to your addressee.

From the speaker's point of view, (T) looks inward, while (I) looks outward. I'm going to argue that we should let (I) guide our interpretation of (T). That is, I'm going to treat (I) as a guide to what (T) means—rather than vice versa.

Here's an orienting piece of that argument. As I've said, (T) seems to articulate part of the ideal of truthfulness.<sup>9</sup> But why should we care that speakers be thus truthful? Truthfulness in asserting appears to matter primarily because we often need to *rely* on speakers for the truth. What seems most important to proper assertion is not merely that the speaker be asserting the truth but that she be informing us of what we need to know. The cases that we'll be exploring emphasize this link between being truthful and being informative. We are concerned to enforce (T) because we regard that as a way of enforcing (I).

But what if we viewed matters the other way round, regarding truthfulness as the key to informativeness? From that perspective—treating (T) as a guide to what (I) means—we may assume that the principal way a speaker can be helpful in asserting that *p* is by revealing that she believes that *p* (or that she would be reasonable in believing that *p*, if she fails to believe it for some non-epistemic reason).<sup>10</sup> I do not believe that the assumption is confirmed by the practice of assertion that our lives actually sustain. Though a speaker can often help her addressee by revealing what she believes, the cases that I'll sketch show how revealing what she believes can prevent a speaker from doing what she illocutionarily ought to do when she asserts. Again, I'm focusing on illocutionary, not on perlocutionary, norms. My point will not be that revealing what you believe can prevent you from satisfying some further—that is, perlocutionary—aim or norm. My point will be that revealing that you believe that *p* can prevent you from satisfying the illocutionary aim and norm inherent in the speech act of asserting that *p*. In Austin's term, it can count as an illocutionary 'abuse.'

As one further act of stage-setting before we turn to the cases, let me define a term that will figure prominently in my treatment, the verb 'to judge.' For ease of formulation, I'll employ a deliberative model of belief formation. By 'judging' that *p* I'll mean the act of doxastically concluding that *p*, where 'doxastically' concluding means concluding in the way characteristic of belief formation.<sup>11</sup> The crucial feature of judgment, for our purposes, is this: when you have drawn a doxastic conclusion that *p*, the question whether *p* is no longer open for you—though you may of course re-open it. When I say that the question whether *p* is no longer 'open' for you—otherwise put, that you have 'concluded' that *p*—I don't mean merely that you've ceased deliberating, or merely that you are no longer disposed to deliberate. I mean that you've undertaken a doxastic commitment. For a reason that will emerge in

section V, I do not want to assume that such a conclusion or commitment must take the form of a belief; as we'll see, believing appears to add to judging an attitude of appropriate self-trust.<sup>12</sup> But in a context of normal self-trust, judging that *p* amounts to believing that *p*. (Though we could dispense with the assumption, I'll for simplicity assume that believing does not come in degrees.)<sup>13</sup> Judging should in this respect be contrasted with hypothesizing or supposing, wherein the deliberative question remains open.<sup>14</sup>

It will prove worthwhile to ask just what is involved in undertaking the doxastic commitment at the core of judgment. Though I lack space for a full defense of the idea, I'll offer a dialectical argument for this proposal: that the core of your doxastic commitment to *p* lies in the presumption that you are in position to inform others that *p*. The dialectical argument draws a comparison with the knowledge account of assertion. The knowledge account of assertion emphasizes the epistemic dimension of a speaker's aim in asserting that *p*: to present herself as knowing and as thereby enabling her audience to know that *p*. That emphasis is correct, but—as I'll argue at length—it does not entail or motivate a knowledge account of assertion. The cases that help us see this also help us see that the epistemic dimension of a speaker's aim in asserting that *p* derives from her more fundamental aim in judging that *p*: to close deliberation by putting herself in position to inform possible interlocutors that *p*, thereby enabling them to know that *p*. I'll thus recast my opponents' emphasis on an epistemic dimension of assertion as an emphasis on an epistemic dimension of the more fundamental act of judgment. Again, I cannot offer a full defense of this move. But I do aim to make it attractive to anyone sympathetic to the knowledge account of assertion.

We're at last ready for the cases that generate the puzzle about sincerity.<sup>15</sup> Consider first a non-standard lottery case. (We'll discuss the standard species of lottery case in section II.) Imagine that Alex has won the lottery and has lots of good evidence that he has won the lottery. But he doesn't believe that he has won, because he takes the doxastic standard in his context to be very high—for example, he thinks he has to rule out an elaborate and thus improbable hoax. 'Surely I've not really won,' he sincerely asserts. 'But you have,' his friend Sheryl replies, continuing: 'Face it, Alex, you've won the lottery.' Here Sheryl has no evidential advantage over Alex, nor does she believe that he's misweighing the evidence. Imagine they both know both facts about Sheryl. Sheryl's assertion may nonetheless be perfectly informative. She does believe what she asserts, but her principal aim in making the assertion is to assure Alex that the doxastic standard is not as high as he is assuming. She thus aims at a species of doxastic-deliberative assistance that is guided by a sensitivity to how her addressee should by-her-lights deliberate in *his* context.

Of course, we may assume that Alex's doxastic context does not differ from Sheryl's. But cases with the feature that I'm emphasizing are easy to concoct along natural lines and needn't involve the assumption that interlocutors are operating in a shared doxastic context, with a single standard for closing deliberation.<sup>16</sup> Consider this case. Ann has lots of good evidence that her spouse is having an affair, but she nonetheless doesn't believe it, feeling a need to rule out far-fetched alternatives. Her friend Siegfried has access to no evidence beyond what's already vividly before

her, but he feels an obligation to tell her, ‘Ann, he’s having an affair.’ Does Siegfried thereby simply instruct her to reweigh the evidence? Quite conceivably not; perhaps he knows she has been weighing and reweighing the evidence, by his lights properly, for weeks. Imagine he views Ann as possessing all easily available evidence, rational in how she weighs it, but nonetheless mistaken in her failure to draw a conclusion. He thus assures her that the doxastic standard isn’t as high as she seems to assume it is.<sup>17</sup>

Notice how the intervention presupposes that Siegfried is, or at least represents himself as, judging in Ann’s context, not merely in his own. Let’s assume that the standard for believing that Ann’s spouse is having an affair is higher in Ann’s context than it is in Siegfried’s own context narrowly described. Assume that apart from the present conversation—say, in contexts of mere gossip with others—Siegfried takes it to be clear enough that Ann’s spouse is having an affair. We may assume that the standard for belief in a context of pure gossip is typically lower than it would be in the context of confronting one of the gossip’s targets.<sup>18</sup> But in addressing Ann he judges, as it were, in her context—with its higher standard—not in that other sort of context. This is because the assertion would not be properly informative if Siegfried intended to be speaking from his own context and not from Ann’s. Since speaking merely from his own context would provide a basis for Ann to mistrust him, he must take care to avoid giving her that impression. This dynamic was also present in the previous example, though there it was easier to assume that the interlocutors simply shared a context. Here, Siegfried and Ann do not share a context; that is, it would not be against the same standard that each would count as knowing that Ann’s spouse is having an affair. Siegfried must therefore represent himself as helping Ann to know in the context of her epistemic needs, not merely in the context of his own.

In what follows I’ll focus on the variant of this case in which Siegfried does not feel entitled to tell Ann flat-out that her spouse is having an affair, though he continues to regard the affair as gossip-worthy. Here we find one species of counterexample to the idea that informative assertion is simply asserting what you believe. Siegfried may take himself to believe that Ann’s spouse is having an affair, and the self-ascription may seem unproblematic right up to the moment at which he addresses her. But then he hesitates as he realizes that his belief derives from that context of gossip. Of course, the hesitation may involve a thought that the gossip has been unfair to her predicament, or other thoughts that articulate a moral or otherwise non-epistemic dimension of his relation to her. But it may also involve a more fundamental recognition that he is not entitled to draw this conclusion in *her* context, with its standard of evidence commensurate with how much is at stake for her in the question. Does he thereby count as changing his mind and unforming the belief that leads him to gossip? I don’t see why he must. Though we’d expect him to be a bit chastened in his disposition to spread this gossip, now that he has confronted the victim and found that he doesn’t feel entitled to tell her, he may without akratic tension regard himself as having settled the matter in his own very different context of epistemic needs.

Looking ahead, we should bear in mind both sub-variants of the case: (a) when he addresses Ann, Siegfried believes that her spouse is having an affair but does not feel entitled to tell her because he regards her context as more demanding than his own; (b) when he addresses Ann, Siegfried abandons his belief that her spouse is having an affair because he regards himself as thereby entering her more demanding context. Sub-variant (b) would rest with the traditional interpretation of the sincerity norm on assertion: Siegfried does not feel entitled to tell Ann that her spouse is having an affair because he does not himself believe that her spouse is having an affair. Again, I don't see why that need be the only intelligible reading of the case. Sub-variant (a) seems equally intelligible—intelligible, that is, as a case in which the speaker does nothing wrong. One reason to find sub-variant (a) equally intelligible will emerge from our next case, in which the asymmetry in context works the other way round: the speaker's context is more demanding than the addressee's.

Whenever the doxastic contexts of speaker and addressee differ, we confront a general question about the illocutionary norm governing assertion. Is the non-'abusive' assertion one that expresses the speaker's beliefs—that is, what she would judge in *her* context—or one that expresses what she thinks her addressee(s) should believe—that is, what she would judge in *his* (or their) context? Our next species of case provides a good reason to think the latter. We can easily imagine how a speaker might informatively tell someone that *p* without having concluded that *p* in her own context—that is, without having formed the belief that *p*.

Here are some details we can work from. Sheila is allergic to nuts and is wondering whether the muffin that she has just split with Aaron contains any nuts. Aaron also wonders whether it contains nuts, but merely because he dislikes their taste. Neither has yet taken a bite, and both are wondering if they should. Imagine that he asks Sheila whether the muffin contains nuts, and she, understanding his doxastic context and regarding the probability that it contains nuts as very low, replies that it does not. But also imagine that she does not feel entitled to the conclusion, drawn in her own doxastic context, that the muffin does not contain nuts. Imagine that she applies to both her and Aaron's circumstances a univocal conception of what it would take for the muffin to count as nut-free. The difference isn't that Aaron couldn't taste the slight trace of nut that could kill her. For all she knows, he could.<sup>19</sup> The difference instead involves the costs to each of being wrong: she wants to avoid the Emergency Room, whereas Aaron merely wants to avoid a mildly unpleasant taste. So she does not judge that the muffin is nut-free, though she feels entitled to tell Aaron that it is nut-free.

Other things equal, of course, Sheila ought to explain the discrepancy in doxastic context to Aaron. But imagine that other things are not equal. Imagine, for example, that the door to the subway car is about to close, and this assertion is all she has time for. Or imagine that she has a good reason not to divulge her nut allergy. Against that backdrop, her assertion is properly informative. After all, not to answer in the way she answers would, against that backdrop, give the impression that she regards it as an open question in *his* context whether the muffin is nut-free. And that, by her own lights, would amount to misleading him. It's hard to see how the



norm governing assertion could on its own ensure that she owes Aaron anything more.

On my reading of the case, Sheila's assertion does not violate the norm governing assertion and is therefore not an illocutionary abuse. On an alternative reading, her assertion does violate the norm—but, in light of broader considerations, permissibly. On this alternative reading, Sheila should not be criticized for the illocutionary abuse because it was more important for her to give what assistance she could before the subway car's door closed, or because she was more fundamentally entitled to keep her nut allergy undisclosed to Aaron, even if she wasn't strictly entitled to give the assistance in the form of that assertion. Perhaps our intuitions that Sheila should not be criticized for making the assertion derive from these broader considerations and are compatible with the assumption that her assertion is nonetheless an illocutionary abuse. Why should we prefer my reading to this alternative?

We can derive a reason to prefer my reading by flipping the case around so that it resembles sub-variant (b) of the Siegfried-Ann case. In that sub-variant, Siegfried does not feel entitled to tell Ann that her spouse is having an affair because he doesn't believe it, having abandoned the belief upon confronting Ann's more demanding context of epistemic needs. Let's accordingly imagine that Sheila refrains from telling Aaron that the muffin is nut-free and that her reason for the omission is simply that she does not herself, in her more demanding context, believe that the muffin is nut-free. There remains, of course, this difference: the standard governing Siegfried's context rises because of his engagement with Ann's context, whereas the standard governing Sheila's context is high despite her engagement with Aaron's context. And that difference makes a key difference in how the cases fall under the sincerity norm on assertion—assuming, for now, that there is a sincerity norm on assertion. (Recall that I'm not committed to that assumption.) Again imagine a background in which Sheila has good reason not to try to explain the context-asymmetrical situation to Aaron. She thus has, schematically put, two options: (i) do justice to A's epistemic needs by telling A that p, through she does not herself believe that p, or (ii) refrain from telling A that p on the grounds that she does not herself believe that p. The traditional interpretation of the sincerity norm on assertion would obligate her to choose the second option over the first. But that interpretation strikes me as simply mistaken in this instance. Whereas Siegfried makes no error by refraining from telling Ann that her husband is having an affair, choosing (ii) over (i), Sheila does make an error when she refrains from telling Aaron that the muffin is nut-free, choosing (ii) over (i), given that she could easily tell him. The difference derives from how the speaker's act engages the addressee's context in the two cases. What matters is how the speaker would judge in the addressee's context.

To see Sheila's error more clearly, imagine the scenario from Aaron's perspective. If Aaron came to understand his context-asymmetrical situation, he would feel entitled to criticize Sheila for failing to inform him of what she was in position to inform him. After all, he would naturally think, she did regard her evidence as easily sufficing for knowledge in *his* context: there was no doubt in her mind whether she could, resting with that evidence, let *him* know that the muffin is



nut-free. Of course, she could have let him know. Her only basis for thinking that she was not entitled to the assertion derived from something that from this angle seems irrelevant: that she, given her very different epistemic needs, did not feel entitled to the judgment herself. ‘But what about my needs?’ he might ask with a tinge of resentment. ‘Why didn’t you tell me what you were clearly, by your own lights, in position to let me easily know?’ There was no speech act, of course, so he could not accuse her of illocutionary error. But he seems right to regard her restraint as revealing a misunderstanding of the sincerity norm on assertion.

The result vindicates my reading of the original case, wherein Sheila does tell Aaron that the muffin is nut-free. If she has only the two options, and if the sincerity norm on assertion does not justify refraining from telling him, then the norm on assertion must at least permit telling him—despite her lack of belief. Again, I don’t claim that this is what sincerity itself requires or permits. While I’m attracted to the idea that sincerity is more a matter of doing justice to what your addressee needs from you than of self-revelation, a reader could accommodate the traditional view of sincerity by abandoning the thesis that sincerity perfectly codifies this norm on assertion. And we might use similar reasoning to justify treating the Siegfried-Ann case along the lines of sub-variant (a), wherein Siegfried does not abandon his belief while feeling that he is not entitled to manifest it in an assertion to Ann. While sub-variant (b) is perfectly possible, of course, we can now see that there is no reason to insist that his hesitation must manifest a tendency to abandon his belief when confronting her more demanding context. If we complicate the case by imagining that Siegfried knows that someone in a low-stakes context is overhearing his assertions, we get a further reason to resist the idea that the speaker’s felt epistemic needs must mirror those of the addressee. (This parallels a complication we’ll add to the Sheila case in section V.) As we’re now beginning to see, what the speaker believes does not directly bear on the obligations that the speaker bears when contemplating whether to assert.

We’re also beginning to see how questions of assertion bear on questions of judgment. One might try to argue that Sheila does judge that the muffin is nut-free when she tells Aaron that it is, because that assertion determines a new doxastic context for her. But this argument encounters the problem that Sheila self-consciously refrains from drawing the conclusion that the muffin is nut-free ‘for herself’ even as she tells Aaron that it is. We therefore need to distinguish the context in which she would judge ‘for herself’ from the context in which she would tell Aaron. Moreover, if we say that Sheila judges ‘for Aaron’ that the muffin is nut-free, we do not mean by ‘judges’ *concludes doxastic deliberation*. By hypothesis, Sheila does not actually draw this doxastic conclusion.

If it makes any sense to say that Sheila ‘judges’ that the muffin is nut-free in this case, it is only because she recognizes that that is what she *would* judge if she were in Aaron’s doxastic context. Were it not for her allergy, she would take herself to be sufficiently warranted—relative to that different standard—that the muffin is nut-free and thus would judge that the muffin is nut-free. In light of this counterfactual, Sheila regards herself as capable of informing Aaron that the muffin

is nut-free, where ‘informing’ him means providing him with what in context would count as a sufficiently warranted belief, should he believe on her say-so—sufficiently warranted, that is, by the epistemic standard governing his actual doxastic context. Looking ahead to the polemic against the knowledge account of assertion, we can put the conclusion this way. You can be entitled to assert what you do not know—indeed, what you do not even judge or believe—as long as you are entitled to view your interlocutor as thereby able to come to know it.

This result forces us to reconceive the nature of assertion—abandoning one or both of the assumptions from which we began—by reconceiving how sincerity figures in the norms governing informative assertion.<sup>20</sup> When Sheila tells Aaron that the muffin is nut-free, she is indeed properly informative—despite the fact that she does not at all actually judge what she says to be the truth. To be sure, she does not believe that the muffin is not nut-free. In fact, she believes that it is probably nut-free. But when she looks for a plausible excuse not to share the muffin with Aaron—assuming, as we might intelligibly imagine, that she does not want to reveal her allergy—she acts on the hypothesis that the muffin is *not* nut-free. Of course, insincerities or other illocutionary abuses may transpire as she extricates herself from the shared project, but we can easily imagine that they do not concern the speech act at issue. (Indeed, we need not have imagined that there was any shared project.) Sheila can thus be properly informative in telling Aaron that the muffin is nut-free, not only while not judging or believing that the muffin is nut-free but while acting on the hypothesis that it is *not* nut-free.<sup>21</sup>

The intelligibility of this sort of case forces us to choose: should we abandon the assumption that sincerity requires believing what you assert, or should we abandon the assumption that insincerity must amount to an illocutionary abuse? Once again, my argument in this paper permits us to go either way—though I am myself inclined to adopt the former resolution. I find it odd to call Sheila ‘insincere’ when she is doing exactly what is required by the context in which she needs to make her assertion informative. But if one finds it too difficult to abandon the assumption that sincerity requires asserting what you believe, one may instead conclude that Sheila’s insincerity is not in this case an abuse. All that matters to the argument that I’ll pursue is that one concede that the illocutionary norm governing assertion requires that a speaker’s truthfulness serve her informativeness—where the ideal of informativeness is as these test cases depict it: the assertion must meet the epistemic standard governing the *addressee’s* doxastic context. (Again, the speaker assumes *her* conception of what evidence is available. The addressee’s doxastic context is the context in which his practical exigencies determine whether that evidence—the speaker’s evidence—suffices to warrant him in concluding deliberation with a belief.) To reflect the manner of engagement required of the speaker, I’ll call this illocutionary norm a ‘second-personal’ norm of informativeness on assertion.<sup>22</sup> We’ll be inquiring further into the norm as we proceed.

## II. Lottery Assertions

How might we apply this lesson to the knowledge account of assertion? My argument will be threefold. First (in this section), I’ll argue that the *error* in the

standard lottery case—not my new lottery case in the previous section, but the case of Susan and Andrew described at the outset—derives not from the speaker’s failure to know what she asserts but from her failure to be properly informative in the way we’ve seen illustrated by the cases in section I. Second (in sections III and IV), I’ll argue that the *incoherence* in the standard lottery case derives from the difficulty in understanding how the speaker could without confusion count as judging what she asserts—since her assertion, on her evidence, proves uninformative not merely to her actual but to any possible interlocutor. Third (in section V), I’ll argue that the error in Moore’s paradox lies not in the assertion itself—“It’s raining, though I don’t believe it’s raining”—but in the assumption that its two conjuncts are addressed to a single interlocutor occupying a single doxastic context. In each of these three instances, the error or incoherence derives from how the assertion purports to serve, but fails actually to serve, the needs that define addressees’ doxastic contexts—not from its relation to the speaker’s context considered on its own. There is a connection to an epistemic standard, then, but that standard is defined from addressees’ perspectives, not from the speaker’s. The upshot vindicates what we might call (somewhat misleadingly, as we’ll see) a ‘knowledge account’ of judgment, not of assertion.<sup>23</sup>

As proponents of the knowledge account of assertion deploy the lottery case, Susan is not entitled to assert that Andrew’s ticket has lost because she does not know that it has lost. (Imagine that though the winning ticket has been drawn, the result has not yet been announced.) There’s a key background consideration: all Susan knows about Andrew’s ticket is that, like any ticket, its chances of winning are one-in- $n$ , for some very large  $n$ . Let’s agree in advance both that Susan is not entitled to make this assertion, and that if Susan has no evidence other than the improbability that the ticket has won, she does not know that the ticket has lost. More specifically, let’s agree both (i) that Susan’s assertion in the standard lottery case is an Austinian abuse and (ii) that we should explain this fact about her illocutionary predicament by noting that her evidence does not suffice for knowledge that the ticket has lost. But the cases in section I help us see that (ii) admits of two readings: (a) her evidence does not suffice for knowledge in her context, or (b) her evidence does not suffice for knowledge in Andrew’s context. Both of these readings of that clause in (ii) are true as we’re imagining the case. But the knowledge account of assertion claims that it’s the truth of (a)—not of (b)—that explains why (i) is true.

Building on my argument in section I, I’ll challenge that claim by coining two variants on the standard lottery case. In the first variant, S’s evidence suffices for knowledge in A’s context but not in her own, and her assertion therefore is *not* an Austinian abuse. In the second variant, S’s evidence suffices for knowledge in her context but not in A’s, and her assertion therefore *is* an Austinian abuse. Each variant parallels a case considered in section I. The first variant parallels the case in which Sheila is entitled to tell Aaron that his muffin is nut-free. The second variant parallels the case in which Siegfried (as we concluded by imagining him) is not entitled to tell Ann that her spouse is cheating. Put together, the two variants show that it’s the truth of (b)—not of (a)—that explains why (i) is true, contrary to the knowledge account of assertion.

Begin by comparing the standard lottery case—that is, assertion of a lottery proposition with no evidence other than the ticket’s one-in- $n$  chance of winning—with a variant in which a bit more evidence is in play. To set terms for the comparison, imagine that there is a clear intuitive difference in the doxastic contexts of the speaker and her addressee. Say Susan and the new speaker Sally are each addressing an interlocutor, respectively Andre and Aiden, who has recently purchased a lottery ticket as a last-gasp measure in dealing with foreclosure on his house. Each addressee is wondering whether to pass up a neighbor’s offer of space on a moving van—leaving within the hour for a warehouse where the addressee’s furniture can be temporarily stored—in the irrational hope that his lottery ticket will be declared the winner on tonight’s evening news. Each speaker wants to tell each addressee that his ticket hasn’t won, and that he needs to get his furniture into that van as quickly as possible lest it be dumped in the rainy street the following morning. But Sally has some evidence that Susan lacks: she has overheard an apparent authority at the television station where she works apparently refer to the lottery winner as ‘she.’ This doesn’t, of course, rule out many alternatives to what she asserts: it’s possible she misheard, or that the authority misspoke, or that the apparent authority wasn’t really an authority, et cetera. Though the evidence does not suffice to close Sally’s deliberation on the matter, given her epistemic needs, it does perhaps count as sufficient evidence for Aiden in his very different—and rather dire—context. Though both Susan and Sally would be more broadly blameless in making the assertion, given the manifest irrationality of their addressees, Sally but not Susan would be blameless *qua* asserter. For only Sally has a satisfactory (we’re assuming) reply to her addressee’s inevitable question: “I know the odds are against it, but why not think mine is the winning ticket?” Though Susan could not be criticized for intervening to get Andre to load his furniture into that van, her failure to have a satisfying reply to that question shows that she could be criticized for the form her intervention took.

Before considering further this difference between Susan and Sally, let’s pause to note an interesting complexity in how practical exigencies make a doxastic standard rise or fall. The exigencies that define the doxastic contexts of each of our new addressees—Andre and Aiden—impose higher costs on his mistakenly *failing* to believe that his ticket has lost the lottery than on his mistakenly believing that it has lost. If either addressee mistakenly *fails* to have that belief (that is, keeps his mind open, though his ticket has in fact lost) and acts accordingly (by refusing his neighbor’s offer), then his precious furniture will be dumped in the rainy street and most of it ruined; but if the addressee mistakenly has that belief (that is, concludes that his ticket has lost, when it has in fact won) and acts accordingly (by accepting his neighbor’s offer), then he’ll merely have wasted the labor of moving these items to the warehouse. We may assume that he won’t get the lottery winnings in either case—in the latter case, say, because this is his only opportunity to form a belief about the matter. Without this assumption, it is less clear that the standard falls. But with the assumption it seems clear enough that having one’s practical exigencies ‘increase’ in this way causes the doxastic standard to fall.

One may have a general intuition that an ‘increase’ in one’s practical exigencies can only raise the doxastic standard for believing a given proposition, but our new case reveals that whether they raise or lower the standard depends on *how* the exigencies have ‘increased.’<sup>24</sup> Though the cases are not parallel in other respects, we might usefully compare how Andre’s or Aiden’s exigencies function with how Sheila’s function when she wonders whether Aaron’s muffin is nut-free: her exigencies, by contrast, impose higher costs on her mistakenly believing that her muffin is nut-free than on her mistakenly *failing* to believe that it is nut-free. If Sheila mistakenly has that belief, and acts accordingly by eating the bite that Aaron would offer her, then she’ll fall violently ill; but if she mistakenly *fails* to have that belief, and acts accordingly by avoiding the muffin, then she’ll merely have missed this opportunity for a tasty treat. The fact that Sheila’s exigencies are intuitively greater than Aaron’s makes her doxastic standard rise relative to his. But the fact that Andre’s or Aiden’s exigencies are intuitively greater than Susan’s or Sally’s makes his doxastic standard fall relative to hers. So, despite the intuitive difference in whether the speaker’s or the addressee’s context is more ‘urgent,’ these new cases parallel the case of Sheila and Aaron in this one important respect: in all these cases, the addressee’s doxastic standard is lower than the speaker’s.

But the fact that Andre’s and Aiden’s doxastic standards are lower than Susan’s and Sally’s does not on its own make it properly informative for Susan or Sally to tell Andre or Aiden that his ticket has lost. As we’ve seen, it seems that Susan would still commit an illocutionary abuse if she performed that speech act (however blameless her speech may be in every other respect, given the irrationality of her addressee). But the difference in the evidence available to Susan and Sally appears to entail this difference between them: unlike Susan, Sally does not fail to be properly informative when she simply tells her interlocutor—imagine she cannot say anything more—that his ticket has lost. We’ll consider presently what this difference between Sally and Susan says about Susan’s illocutionary predicament. For now, note the parallel between Sally’s predicament and Sheila’s predicament in the nut-allergy case. Because neither Sheila nor Sally is entitled to judge what she asserts—that is, to draw that deliberative conclusion ‘for herself’—neither would be warranted in believing what she asserts. We may therefore stipulate that neither actually believes what she asserts—not merely for non-epistemic reasons but for the epistemic reason that she would not be epistemically warranted in believing what she asserts.<sup>25</sup> But neither Sheila nor Sally appears to be committing an illocutionary abuse when she tells her interlocutor what she does not believe. We’ll consider presently why the same cannot be said of Susan, even when her interlocutor, Andre instead of Andrew, confronts the same low doxastic standard as that confronted by Sally’s interlocutor, Aiden.

Can we coin a lottery case to parallel the case from section I in which Siegfried would commit an abuse if he told Ann that her spouse has betrayed her? In a second variant on the standard lottery case, consider the predicament of Aiden’s spouse, Sabine. Imagine Sabine’s evidence as we imagined Sally’s in the first variant, but imagine that Sabine shares Aiden’s doxastic context—since she co-owns the

furniture that may soon be left in the street. On that evidential basis it appears that Sabine could informatively tell Aiden that his ticket has lost but not Susan's original interlocutor Andrew. Imagine that Andrew would rightfully scoff at Sabine's weak evidence if he learned of it: such weak evidence plainly does not suffice to settle the question for him, even for the nonce. Imagine it clear to all that, with only that evidence to add to the statistical evidence, he cannot rule out an alternative that exigencies of his context require that he rule out before closing doxastic deliberation: the alternative that, despite both the statistical evidence and Sabine's rather weak experiential evidence, his ticket has nonetheless won. Thus imagined, we find in Sabine's assertion a counterexample to the sufficiency of the knowledge account of assertion, just as we found counterexamples to its necessity in Sheila's and Sally's assertions. We may coherently imagine at once both that, unlike Susan, Sabine *does* know what she asserts when she tells Andrew that his ticket has lost and that her assertion is not properly informative. It follows that even when a speaker knows that *p* while asserting that *p*, that knowledge does not ensure that her assertion will count as properly informative.

It might seem that we now have an alternative diagnosis of why Susan commits an abuse when she tells her original interlocutor Andrew that his ticket has lost. She cannot tell him, we might say, because like Sabine she cannot let Andrew know that his ticket has lost. Our alternative to the knowledge account of assertion would simply reconstrue the epistemic constraint as arising not from the speaker's doxastic context but from that of her addressee. It would follow that there is nothing distinctive about lottery cases: they are merely like many other cases in which the speaker cannot properly assert because she cannot thereby inform her addressee. This is the substance of what I've thus far argued. But that 'simply' and 'merely' are too quick. I believe there *is* something distinctive about lottery cases. The alternative diagnosis of pure lottery cases—the lottery proposition asserted with only the statistical evidence—is fine as far as it goes. But we should pursue a deeper diagnosis of the speaker's epistemic predicament. Broadening our focus beyond the speaker's narrow epistemic needs to include those of her addressee has put us in the right spot, but we haven't yet hit paydirt.

Before we dig deeper, let's see why the shallower approach to lottery cases won't quite work on its own. It would be simply a mistake, I think, to hold that the sole linguistic error in a lottery assertion is the error on which we've thus far focused—failing to do justice to your addressee's epistemic needs. That is the core *illocutionary* error, because the illocutionary question here is interpersonal, a matter of securing without exploiting your audience's 'uptake' in the respect that Austin explored. That said, we should distinguish a non-illocutionary but nonetheless linguistic act, *merely* asserting, from the illocutionary act of telling. It is possible *merely* to assert that *p*—that is, to assert that *p* without telling anyone that *p* and thus without adopting the interlocutory stance that we've been emphasizing. Imagine, for example, your exclamation as you conclude a solitary deliberation, "Why of course: *p*!" When we imagine a case of *mere* assertion—assertion with no one told—we cannot imagine that the act is an Austinian abuse of the sort we're investigating, since it seems impossible to assert that *p* insincerely (or uninformatively) without telling someone that *p*.<sup>26</sup> Imagine you come out with a solitary 'My ticket has lost!' as such a mere

assertion. The assertion involves a linguistic error: you are not entitled to assert that. But the error cannot involve the second-personal considerations that we've emphasized, since you are not addressing anyone. And the error cannot be simply that you do not know that your ticket has lost—since, as we've seen, knowing that *p* is neither necessary (cf. Sheila) nor sufficient (cf. Siegfried) for being entitled to assert that *p*.<sup>27</sup> What error, then, is common to all lottery assertions, those with and without an addressee?

We need to place more emphasis on the datum that the speaker's assertion in a standard lottery case does not appear to be *merely* an error, as if Susan had merely failed to grasp—as we may imagine Sabine does when she addresses him—that Andrew's doxastic context is more demanding than her own. If we let our diagnosis of lottery cases rest with an analogy between Susan's predicament and Sabine's, we'd have a superficial explanation of each speaker's error but no explanation of our deeper sense that Susan's assertion is incoherent in a way that Sabine's is not. It is an important datum that Sabine does not appear confused in the way that Susan appears confused. Since Sabine can inform her spouse Aiden but not Andrew, her error as she addresses Andrew might well involve mere (perhaps nonculpable) ignorance of the differences between his context and Aiden's—no sign that she is confused about anything. We do not have a full explanation of Susan's incoherence until we know how to weigh a further feature of her predicament: that given her evidence she cannot informatively tell, not merely Andre or Andrew, but any conceivable addressee that any particular ticket has lost. Moreover, it is difficult to see how this fact about her predicament could have escaped her grasp. There is, in sum, something not merely erroneous about her speech act but more deeply incoherent. Since it does not depend on any particular interlocutor, this incoherence would infect a mere assertion as well.

My deeper claim, which I'll develop in section III, is that we can understand this further feature—the feature that distinguishes the incoherence that we find in a pure lottery case from the illocutionary error that we find in a case like Sabine's—by grasping an important connection between assertion and judgment. I don't mean merely the connection that we've already observed between assertion and a hypothetical judgment framed by 'if I were in his context.' I mean a connection between assertion and judging from one's own point of view—genuinely concluding or foreclosing deliberation in the way that would inform belief. The connection is simply this: because Susan cannot thus improve any conceivable addressee's epistemic position, she cannot without confusion judge that the ticket has lost. This inability to conceive how she could improve any possible addressee's epistemic position and her inability to form a coherent judgment appear to go hand in hand. The lottery case—that proposition asserted with only that evidence—thus reveals how the capacity for judgment rests on the capacity for assertion.

### III. The J-Account

For ease of reference, let's call the knowledge account of assertion *the A-account*, and let's call the alternative I'm about to develop *the J-account*, since we can (for now) regard it as a knowledge account of judgment. (Again, in section V



we'll see that the account I'm developing is a *conclusive-warrant* account of judgment, not strictly a knowledge account.)<sup>28</sup> In the previous section I explained how the A-account misdiagnoses the *error* in lottery assertions. In this section I'll explain how the A-account misdiagnoses the *incoherence* in lottery assertions. The latter misdiagnosis derives, I'll argue, from a conflation. The A-account treats what is actually a failure of judgment as a mere failure of knowledge. One who makes a lottery assertion does, of course, fail to know what she asserts, but the most revealing reason why she fails to know it is that she cannot without confusion judge what she asserts. We need the J-account to pursue this deeper explanation.<sup>29</sup>

My argument for the J-account will come in two stages. I'll first make a *prima facie* argument articulating the intuition that Susan cannot coherently judge that Andrew's ticket has lost. I'll then pursue a deeper argument to explain why that intuition is correct. As I'll emphasize, one could accept my *prima facie* argument while rejecting my deeper explanation and thereby accept the core of the J-account over the A-account. Where it makes sense to distinguish the two, I'll accordingly speak of the *core* and the *extended* J-accounts. Though the core account does not require it, we would miss something important if we failed to pursue the deeper explanatory argument. As we'll see, the knowledge account of judgment benefits from being supplemented by what we might inelegantly call 'the assertion account of the role of knowledge in judgment'—in part because the supplement rearticulates what we may regard as the core (if misarticulated) insight of the knowledge account of assertion.

Here's the *prima facie* argument. Even setting aside the question whom she could conceivably inform by asserting the judgment, it is plausible that Susan cannot without confusion judge that Andrew's ticket has lost. Imagine she judges that the ticket has lost by this reasoning:

- (1) It is very improbable that the ticket has won.
- (2) So the ticket has lost.

The inference is obviously invalid. So any such deliberative concluding would be obviously illicit. Now ask more broadly: with evidence only of this improbability, could Susan know the proposition that she judges? It seems she cannot. For her evidence fails to rule out the very possibility that her judgment explicitly denies: that the ticket has won. It is plausible that she should not judge that the ticket has lost *because* she does not know that it has lost.

We can make the proposal vivid by an analogy with the interpersonal case. Here is how Timothy Williamson marshals intuitions in defense of the knowledge account of assertion:

Intuitively, my grounds are quite inadequate for that outright unqualified assertion ['Your ticket did not win'], even though one can construct the example to make its probability on my evidence as high as one likes, short of 1, by increasing the number of tickets in the lottery. You will still be entitled to feel some resentment when you later

discover the merely probabilistic grounds for my assertion. I was representing myself to you as having a kind of authority to make the flat-out assertion which in reality I lacked. I was cheating.<sup>30</sup>

Williamson goes on to note: “There is a special jocular tone in which it is quite acceptable to say ‘[Come off it—] Your ticket didn’t win,’ but the tone signals that the speaker intends not to make a flat-out assertion. In the imagined example, I do not use that tone.” It is not clear that one can joke *with oneself* in any analogous sense, but let’s stipulate that the jocular tone is not present in Susan’s self-directed discourse. Her judgment, let’s assume, is flat-out.

We can easily imagine Susan analogously attempting to cheat *herself* by reasoning silently: ‘There are thousands of tickets! Andrew’s hasn’t won. Wait . . . how do I know that? Perhaps it has.’ Should she fail to catch the cheating till later, after it has led her to some action that gets her in trouble, she would naturally resent—at any rate, regret—this judgment. Since the consideration that you do not know that *p* always cuts against your doxastically concluding that *p*—“Wait, do I really know this?” you say to yourself as you find yourself on the verge of drawing the conclusion—it seems we should embrace a knowledge account of judgment. If you do not know that *p*, you should keep your mind open on the question whether *p*. It may be alright to hypothesize that *p*, but you should not bring your doxastic deliberation to a conclusion: *you should not flat-out judge*.<sup>31</sup> Though it is a mistake to view judgment as mere self-directed assertion,<sup>32</sup> judgment seems constrained by a norm of knowledge in the way that Williamson claims is characteristic of assertion.

Here then is the core of my alternative diagnosis of the incoherence on display in a pure lottery case. What seems incoherent in the lottery assertion, I propose, is that it is hard to see how the speaker could have judged what she asserts. Since she so transparently fails to know what she asserts, we must wonder how she could feel entitled to that doxastic-deliberative conclusion. True, the conclusion is very likely correct. But it is obvious to everyone that she does not know that it is correct. So what is she doing presenting herself as entitled to the conclusion? On my diagnosis, the explanatory failure lies deeper than Susan’s failure to know what she asserts. What explains the incoherence of her assertion is instead a failure of judgment.

This *prima facie* argument does not appeal to the second-personal dimension that I’ve been emphasizing. As far as the *prima facie* argument for the J-account goes, we could view it as merely ‘internalizing’ the A-account—treating judgment as ‘internalized’ assertion, or treating assertion as ‘externalized’ judgment, and then arguing that judging that *p* implies an implicit claim to know that *p*. True, judging is not an illocutionary act, so the rule for when you’re entitled to judge does not have a directly social upshot. But we’re now explaining not the error of doing something that you’re not entitled to do but the incoherence of trying to do something that you cannot without confusion regard yourself as doing. And it may seem that the reason why Susan cannot count as coherently judging that Andrew’s ticket has lost, on her evidence, is simply that she does not regard herself as satisfying a necessary

epistemic condition on judging—just as the A-account presents her as not regarding herself as satisfying a necessary epistemic condition on asserting.<sup>33</sup>

So what's wrong with merely 'internalizing' the A-account in this way? Why should we view the incoherence in Susan's assertion that Andrew's ticket has lost in the second-personal terms that I've been emphasizing? It is an important datum that the evidence in play in a lottery case distinctively fails to permit the speaker to help *any* addressee to know that the ticket has lost. Moving from the prima facie argument to the deeper explanation why it is sound, my proposal is not merely that this speaker cannot without confusion judge what she asserts because she cannot know what she asserts. I propose that we explain her inability without confusion to judge what she asserts by appeal to her inability to inform any possible interlocutor through the assertion. The second-personal dynamic that I highlighted in sections I and II is not manifest in the lottery assertion as such, so it is easy to overlook its bearing on the case. The key is to see how the second-personal dynamic, whereby the speaker informs her addressee, is not merely absent but *excluded* from the speaker's relation to *anyone* to whom she could possibly address that assertion. It is this exclusion, I propose, that explains the incoherence in the speaker's assertion, because it is the exclusion that explains why the speaker cannot without confusion judge what she asserts.<sup>34</sup>

The issue I'm engaging turns on why this is the correct order of explanation. As a preliminary question, let's ask why the impossibility of informing anyone that *p* explains your inability to judge that *p* rather than vice versa: why not think instead that the impermissibility of judging that *p* explains the impossibility of informing anyone that *p*? We've already seen how the latter puts the explanation back to front, since we've seen how your failure to judge that *p* does not prevent you from informing others that *p*. You can inform someone that *p*, despite not judging that *p*, when the context in which you judge places higher epistemic demands on you than the context in which you assert, where the context in which you assert is your addressee's context rather than the context that is narrowly your own. So if we're going to run such an explanation in either direction it must be from the impossibility of informing anyone to the impermissibility of judging. But that leads us to a deeper question: why should we be using this interlocutory framework to explain judgment in the first place?

Let's back up and see where we stand. The dialectic that we're pursuing in this paper starts from the A-account, with its claim that the illocutionary rule governing assertion requires that the speaker know what she asserts. Though proponents of the A-account do not gloss this requirement in second-personal terms, as the requirement that the speaker be able to help her addressee come to know what she asserts, that is the approach that we're pursuing (based on the argument in section I). Our approach contradicts the A-account on the specific question of illocutionary error: the A-account holds that this is the error of asserting what you do not know, whereas our account argues that failing to know is neither necessary nor sufficient for this species of error. We can put the point like this: the claim that you 'outwardly' stake when you assert that *p* is not merely an expression of the claim that you 'inwardly' stake when you judge that *p*. By 'staking' a claim

I mean undertaking an appropriate species of commitment. Again, as we've seen, you can properly stake either claim without its being proper for you, in context, to stake the other. It follows that we should not view assertion as 'externalized' judgment or judgment as 'internalized' assertion. Reflection on the examples of context-asymmetry in sections I and II suggests that even when the speaker asserts what she also judges, the two species of act have crucially different functions.

How might we characterize this difference? The possibilities of properly asserting what you could not (in your context) properly judge and of properly judging what you could not (in your context) properly assert suggest this formulation: when the speaker asserts what she also judges, she stakes a claim of epistemic authority by judging and a claim of illocutionary authority by bringing the claim of epistemic authority to bear, via her assertion, on her addressee's epistemic needs. She thus stakes two fundamentally different claims. Both claims are epistemic, but the claim staked in assertion is not to epistemic authority *simpliciter*. When S tells A that p, S stakes a claim to be offering A epistemic assistance on the question whether p from the perspective of A's doxastic context, bringing S's own evidence to bear on A's epistemic needs. The two possibilities that we've been emphasizing thus fall nicely in place: (i) S can properly claim the epistemic authority inherent in judging that p even while acknowledging that her evidence does not suffice to entitle her to tell a given addressee that p, and (ii) S can properly tell a given addressee that p even when she does not regard her evidence as entitling her to judge that p.

As we'll see in section IV, we must take care in formulating the second possible case. Though we may usefully say in case (ii) that S hypothetically judges that p—that S would judge that p if S were in A's context—when S regards herself as thus doing justice to A's epistemic needs, the reference to this hypothetical judgment does not do any work in helping us to understand why S regards herself as entitled to tell A that p. The claim S stakes when she makes the assertion is a claim to be able to help A, not the claim that she would make that judgment if she were in A's context. It is not wrong to gloss such a case in terms of a hypothetical judgment, but the gloss does not explain what's going on. In such a case, as in every proper assertion, S is aiming to do justice to her addressee's epistemic needs, not merely to show him what she would judge if she were in his context. (As we'll see when we consider an objection in section IV, this is a crucial feature of the extended J-account.) The reason she would, as she thinks, form that judgment if in A's context is that in aiming to do justice to A's epistemic needs she aims to do justice to the epistemic needs of anyone who shares those needs.<sup>35</sup>

Here then is how the extended J-account explains the core idea that judgment is guided by a norm of knowledge (that is, of sufficient warrant). Picking up a strand from earlier sections, we may formulate the extended J-account as follows. Your aim when you judge, we may say generally, is to constitute yourself as a potential informant.<sup>36</sup> You need not, we've just noted, aim to put yourself in position to let *all* your interlocutors know that p when you judge that p. But you cannot judge that p while feeling unable to let *any possible* interlocutor know that p. Again, 'possible interlocutor' means anyone with whom you can realistically imagine yourself

discussing the matter.<sup>37</sup> If you could not let any possible interlocutor know that *p*, then you could not give anyone your word that *p*. If you could not give anyone your word, then you lack any ‘word’ to give. If you lack any ‘word’ to give—if you do not feel epistemically capable of informing any possible interlocutor whether *p*—then you have not formed a judgment whether *p*.

By putting the point in terms of your ‘epistemic’ capacities, we abstract from your physical capacities, your psychological capacities, or your purely social capacities. You can of course judge that *p* while acknowledging that you cannot tell anyone that *p* because someone has gagged you or given you a drug that induces great shyness, or simply because no one will listen to you. And again, by ‘anyone’ here we mean any *possible* interlocutor. If you do not regard yourself as epistemically capable of informing any possible interlocutor on the question, then when someone asked you whether *p*, you would reply ‘I just can’t say,’ meaning not that you lack words to utter but that you lack a word—*your* word, your assurance—to give. More simply, you would mean that you have not made up your mind.<sup>38</sup>

Of course, you cannot close doxastic deliberation by imagining just anyone—say, someone in a context in which standards are very lax—able to rely on your judgment. You imagine yourself able to inform someone in a context with standards that match or exceed your own. Here is another way to put it. To close doxastic deliberation, you imagine being able to meet pertinent challenges to whether you are epistemically entitled to your conclusion. To constitute a challenge, the doxastic context of your imagined challenger must generate a standard that is appropriately high. And to be pertinent, the imagined challenge must be one that you are obliged to meet. To be a challenge that you are obliged to meet, the imagined challenger must be someone whom you have told, or at least whom you imagine yourself entitled to tell, the proposition that you judge. As we’ll see in the next section, this feature of the interlocutory framework requires elaboration in reply to a natural objection.

#### IV. Replies to Two Objections

We obviously lack space for a full defense of the extended J-account, given the other issues also under discussion in this paper, but let me fend off two natural objections before moving on to Moore’s paradox.<sup>39</sup> Briefly put, the objections are these. First, how could satisfaction of the second-personal epistemic condition that I’m proposing actually suffice to entitle you to judge? Second, is the condition really even necessary? Can’t you conclude doxastic deliberation without taking yourself to meet this strong epistemic condition?

Here’s an elaboration of the first objection. The extended J-account explains the point of your judging that *p* by construing your aim as implicitly second-personal: you aim to put yourself in position to inform a possible interlocutor that *p*. But then what should we say of possible or actual interlocutors in less demanding doxastic contexts—for example, of nut-allergic Sheila’s merely nut-disliking interlocutor Aaron? The question refocuses a worry that emerged at the end of the previous section. We noted there that such interlocutors are insufficiently

challenging. This implies the further result that the interlocutory emphasis on its own yields at most a necessary, not a sufficient, condition on judgment. You cannot without confusion conclude doxastic deliberation unless you feel able to inform a possible interlocutor on the matter, and you should not conclude unless you are so able, but being able to inform a possible interlocutor does not suffice to entitle you to the conclusion, nor does feeling able suffice to get your deliberation closed.<sup>40</sup> Satisfaction of what interlocutory condition would suffice? We can say that you must be in position to inform a possible interlocutor *who shares your doxastic context*. But how does that differ from saying that you must meet the epistemic standard that defines your context? The interlocutory dimension seems not to be doing any work.

In reply, let me emphasize a dimension of the interlocutory perspective that I have thus far discussed only in passing. As we noted in the previous section, when Sheila tells Aaron that the muffin is nut-free, she does not merely represent to him that she would judge that the muffin is nut-free if she were in his doxastic context. She assures Aaron that this is the conclusion that *he* is entitled to draw. Sheila is not, in other words, merely presenting herself to Aaron as believing that *p*, or as such that she would believe that *p* if she were in his context. She is inviting him to rely on her, to take her word for it, that the muffin is nut-free. She would issue the same invitation to someone who shared her doxastic context—if she felt she could thereby inform such an interlocutor.

To say that judging is feeling able to inform a possible interlocutor who shares your doxastic context does not merely repeat that judging is feeling that you've met that epistemic standard. It's hard to see how anything of epistemological interest follows simply from your feeling that your evidence meets the epistemic standard in your context. The problem isn't merely that your feeling might be wrong. The problem is that we need to understand how meeting this standard has genuinely epistemic implications—that is, how it lives up to its status as an *epistemic* standard. Of course, your feeling that you've met the standard is a feeling that you know, but the problem is that we need to understand why knowing is something anyone should care about in this way. In the extended J-account that we're pursuing, we seek an explanation of the J-account's core claim that there's a knowledge norm on judgment. And here is where it helps to bring assertion back into the story, appealing to the 'assertion account of the role of knowledge in judgment' that I mentioned at the beginning of the previous section. 'My evidence meets the epistemic standard for *p*!' is not yet an assertion that *p*, since staking a claim about one's evidence falls short of staking a claim about what that evidence shows. Asserting is not just presuming success in evidence-acquisition but a distinctively epistemic species of second-personal commitment. And it provides the explanation that we seek if we regard judgment as likewise commissive. The knowledge norm on judgment derives from what the A-account correctly regards as the basis of an epistemic norm on assertion. As we've seen, there is no knowledge norm on assertion. But the A-account is correct that there is nonetheless an epistemic norm on assertion, and we're following parallel reasoning toward identifying an epistemic norm on judgment. In each case, the reasoning emphasizes the commissive nature of the act.

What is this commitment? Our second-personal perspective on judgment explains it as the commitment that would inform assurances you might give others.<sup>41</sup>

It may help if I try to express directly (if crudely) the intuition on which the extended J-account trades. This isn't the argument (which, remember, has a dialectical structure) but the intuition on which the argument draws. To assure someone that *p* is to commit yourself to *p*—whether sincerely or insincerely—in a way that is familiar and central to what we do with the concept of knowledge. ('You know whether *p*, so you can tell him.' 'How could you tell him that *p* when you don't know that *p*?' To see the point of these interventions, substitute 'assure' for 'tell.') Without opportunities for such assurances, it would be hard to see how anything could count as concluding a deliberation rather than merely letting it lapse. Without others to judge *for*, it would be hard to see what point there could be to the concept of judgment. You could do just as well merely hypothesizing and thereby avoiding conclusions.

I'll conclude my treatment of the J-account by replying to a second, more general objection, targeting not my emphasis on interlocution but the whole idea that judging involves an implicit claim to know. Do we not sometimes intelligibly believe what we fail to credit ourselves with knowing? Let's take for granted that believing that *p* entails judging that *p*. Are not such cases counterexamples to the J-account? We can explain away such apparent counterexamples by emphasizing an aspect of the subject's interlocutory situation that we have not yet discussed.

Let's begin by reminding ourselves that if you do not think you *know* that there is a predator stalking you, for example, you can refrain from judging that there is a predator stalking you without rendering yourself motionless and doomed. You can form the hypothesis that there is a predator stalking you, and then make evasive maneuvers while continuing your deliberation. If that is your attitude toward the proposition, then you should not tell your companion that there is a predator stalking you but merely that there *may* be, or that you suspect there is. You have a choice between shifting modality or substituting a different proposition, most naturally one not directly about the hypothetical predator but about this deliberation that you have not yet closed.

We sometimes mark these distinctions—between hypothesis and judgment, between first- and second-order judgments—by asserting 'I believe that *p*, but I don't know that *p*.' I don't have in mind cases in which you assert 'I believe that *p*' rather than '*p*' merely in order to express hesitancy on the question whether *p*. Such cases are common, but they would not constitute a counterexample to the J-account. We can treat such cases along the lines of what Richard Moran calls the Presentational View, on which saying 'I believe that *p*' rather than merely '*p*' serves primarily to express not the speaker's belief that *p* but instead his or her hesitancy in asserting that *p*.<sup>42</sup> As many philosophers note, the most natural way to express your belief that *p* is to assert '*p*'—not 'I believe that *p*.' I'm interested in apparent counterexamples to the J-account in which the speaker really does refer to his or her state of belief. Do these less common cases show that you can judge that *p* without treating yourself as knowing that *p*?



Here is a journalistic example: “Richard Armitage, the Deputy Secretary of State in Bush’s first term, told me, ‘I think Iran has a secret nuclear-weapons program—I believe it, but I don’t know it.’”<sup>43</sup> In repeating ‘I believe’ after his initial ‘I think,’ Armitage’s full assertion appears not merely to express hesitancy but to testify to his belief state. So how can he consistently assert that he believes that Iran has a secret nuclear-weapons program while also appearing to assert that he lacks sufficient warrant for drawing that conclusion? In saying that he does not know that Iran has a secret nuclear-weapons program, Armitage is acknowledging that the epistemic standard for assessing such a claim, especially as made by a former Deputy Secretary of State, is very high. The disavowal of knowledge makes sense only because of this feature of the background. Compare: “Richard Armitage, looking out his office window at the torrents of rain pummeling the glass, told me, ‘I think it’s raining out—I believe it, but I don’t know it.’” This could make sense only if Armitage and his interlocutor are engaging in some epistemology, with a skeptical challenge looming in the background. But then the disavowal of knowledge would play the same role as before in acknowledging the unusually high epistemic standard. Since the standard is unusually high, the self-ascription of belief in each case has a hypothetical air: I am disposed to act as if *p*, and I would count myself as believing that *p* if the epistemic standard were not so unusually high. Armitage’s ‘I believe it’ amounts only to a hypothesis plus an indication that the standard is unusually high.

Though the J-account does not hold that you must judge that you know that *p* (or have sufficient warrant for *p*) in order to judge that *p*, it can thus accommodate cases in which you coherently self-ascribe belief—that is, say ‘I believe’ expressing not mere hesitancy but your doxastic state of mind—without self-ascribing knowledge. What you would be saying is that you *would* know, and therefore be entitled to judge, if the epistemic standard in your interlocutory context were not so unusually high. It is an acknowledgment not only that but how the epistemic standard governing judgment is sensitive to context: since you are telling *me*, and my doxastic context imposes a higher epistemic standard, you do not feel entitled simply to tell me what you would feel entitled to judge relative to the needs of other possible interlocutors.

Reframing the dialectic between the A-account and the extended J-account in light of our elaborations in this section, we can now make better sense of the intuition, pressed by proponents of the A-account, that there is a claim of epistemic authority specifically in the speech act of telling someone that *p*. When you judge that *p*, says the J-account, you presume yourself able to inform some possible interlocutors that *p*. When you tell someone that *p*, the presumption gets more specific and is informed by your specific manner of address. Does this manner of address yield a specifically linguistic presumption of authority? No, says the J-account, because we will need to appeal to such an interlocutory relation even in judgment. The proponent of the A-account is wrong, then, to assume that we can so easily separate the question of your epistemic authority in judging that *p* from the question of your epistemic authority in giving an interlocutor your word that *p*.

What's distinctive of the lottery case, on the J-account, is that it doesn't matter how we imagine your interlocutor: given that your evidence consists merely of the long odds of winning, which specifically and saliently fail to rule out the possibility that the ticket in question has won, there is no one to whom you could without confusion claim the authority at the core of an assertion that the ticket has lost. The obviousness of that fact about your predicament—that there's no one whom you could engage in this way—in turn explains your inability without confusion to claim the epistemic authority at the core of judgment. Even if your addressee isn't aware of his ticket's low probability of winning—your one piece of evidence, so your one chance to be helpful in this dimension—it manifests a confusion to bring this fact into view by telling him that his ticket has lost. (There are, of course, other ways to do so.) It is hard to see how you could be imagining your epistemic authority if you tried to form the judgment that his ticket has lost—that is, without significant confusion about what it takes to count as such an authority. And the point doesn't apply merely to this or that interlocutor but to an arbitrary interlocutor—that is, to possible interlocutors in general. In being unable without confusion to claim the epistemic authority inherent in telling an arbitrary interlocutor that his or her ticket has lost, you are unable to claim the epistemic authority inherent in judging that the ticket has lost. This inability to judge that the ticket has lost ensures that when you assert that the ticket has lost you are not merely in error, through an Austinian abuse, but more deeply incoherent.

### V. Moorean Assertions

I've argued that you *should* judge only what you can informatively assert, where what you can informatively assert depends on the epistemic standards implicit in the contexts of possible interlocutors. But *can* you judge even when you do not presume yourself entitled by those standards? Moore's paradox codifies a lingering oddity in such a case. I have elsewhere tried to explain how you can in a context of self-mistrust coherently judge what you at the same time explicitly regard yourself as failing to know or even to believe,<sup>44</sup> but for present purposes I'll rest with the less controversial assumption that such a Moorean judgment, and any assertion expressing it, must be incoherent. I'll develop a contrast between a Moorean assertion that expresses a Moorean judgment and different sort of Moorean assertion, one modeled on a variation of the Sheila case that we discussed in section I. Though both sorts of Moorean assertion count as Austinian abuses, I'll argue that the latter sort need not manifest any incoherence. We'll see how the parallel between the incoherences in lottery assertions and in self-mistrustful Moorean assertions reveals the nature of that incoherence. And we'll see how the possibility of a coherent Moorean assertion reveals especially clearly the nature of the illocutionary error committed by Moorean and lottery assertions alike.

Let me begin with the illocutionary side of Moorean 'absurdity.'<sup>45</sup> Why does a Moorean assertion—'P, but I don't believe that p'—violate the norm governing assertion and thereby constitute an illocutionary abuse? The problem, as with lottery assertions, lies in seeing how your assertion could count as informative. In

terms of the contrast I'm going to develop, the problem is that you typically cannot express a judgment that you do not trust by simply asserting the proposition. In light of my broader view, I want to avoid the assumption that judgment implies belief. So I'm formulating the J-account in terms of conclusive warrant, not in terms of the mental state of knowing, which many philosophers conceive as including belief. Without unusual background conditions, however, an assertion expresses not only judgment but belief: when you tell A that p, you typically thereby invite A to rely on your judgment in the way that you thereby show you rely on your judgment in believing that p. You cannot non-abusively—in the Austinian sense, if not in others—invite someone's trust when you have not even secured your own. Note, here again, that what appears to be a problem with asserting that p is more specifically a problem with the second-personal act of trying to tell people that p. You cannot tell someone that p without representing yourself as thereby informing him that p. And you cannot in this way—that is, illocutionarily—inform someone that p if without further explanation you go on to inform him that you do not believe that p.

Moorean cases allow us to see especially clearly the second-personal nature of this illocutionary error. Note that if you could cancel the implication that you are inviting your hearer to trust your assertion that p—if, in other words, you could make it clear that while you are asserting that p you are not telling him or otherwise attempting to inform him that p—then there would be no reason why this first conjunct of your Moorean assertion should by itself represent you as believing that p. And if it did not so represent you, then there would be nothing illocutionarily wrong with proceeding to the second conjunct and reporting that you lack this belief. Such a speech act would prove enormously difficult to pull off, however. Though I do not regard it as impossible given unusual background conditions, I'm setting aside that possibility here. The problem in any ordinary case is that if you are *addressing* this hearer and you assert that p, then you thereby count as telling him that p—which is inconsistent, in the second-personal way we've elucidated, with also telling him that you do not believe that p.

To make a Moorean assertion without illocutionary abuse you'll need to introduce some complexity into your mode of address. You might try this: tell A but not overhearer B that p, and then whisperingly tell B but not out-of-earshot A that you do not trust your judgment that p. There: you've intelligibly—to B—pulled off a non-abusive Moorean assertion. It's hard to see how you could count as making a non-abusive Moorean assertion *flat-out*, since without such redirection you would need to offer so much explanation and exculpation that you would typically compromise or undermine the assertive force of your utterance.<sup>46</sup> But with the redirection . . . voilà, a non-abusive—to one hearer, at least—flat-out Moorean assertion.<sup>47</sup>

Such redirection makes possible a non-abusive Moorean assertion even without self-doubt. Consider again the case of Sheila and Aaron, but imagine that Sheila goes on to assert as an aside to Ben, a friend well aware of her nut allergy and concerned that she not trigger it: “. . . still, I don't believe that the muffin is nut-free.” If Ben has also overheard the first conjunct of Sheila's assertion—“The muffin is nut-free . . .”—as addressed not to him but to Aaron, then he is in position to

make non-abusive sense of Sheila's total assertion. It may look at first glance as if Sheila mistrusts her own judgment, but we've been clear all along that she does not. She tells Aaron that the muffin is nut-free because she regards herself as able to inform him that it is nut-free, even though she does not herself, in her higher-stakes context, judge that it is nut-free. And Ben, we'll assume, understands what's going on. Drawing on his understanding of how she copes with her severe allergy, he fully expects that she doesn't believe what she tells Aaron and therefore isn't surprised by her aside confessing disbelief. From Ben's perspective, Sheila is meeting her illocutionary obligations at every turn: helping Aaron believe what he's warranted believing in his context, but also helping Ben believe what he's warranted believing in the context that we assume he shares with her.

One might suspect that Sheila's Moorean assertion could at best count as a permissible violation of the illocutionary norm from Ben's perspective, just as one might suspect that its first conjunct could at best count as a permissible violation from Aaron's perspective. We considered the latter suspicion in section I and found that there is no reason to credit it. Can we adapt that argument for the Moorean case? There are actually two questions here. First, can Ben reasonably avoid the thought that Sheila's assertion to Aaron at best amounts to a white lie or some other permissible violation of the illocutionary norm? We've already covered this question in our treatment of the original Sheila case, where we saw how she could not give an illocutionary justification for refraining from making the assertion. If there is no illocutionary justification for not asserting, it cannot be an illocutionary error to assert. This reasoning applies from Ben's perspective just as it would from any other. Second, must we regard Sheila as committing an illocutionary abuse in telling Ben that she does not believe that *p* when she knows he overheard her just tell Aaron that *p*? Well, we're imagining two things about Ben: that he has no stake in the question whether *p* and that he knows and sympathizes with Sheila's predicament vis-à-vis the allergy that leads her not to believe that *p*. So it would be pointlessly misleading for her not to acknowledge to Ben what he already suspects is the case. As we've seen, illocutionary consistency with her assertion to Aaron does not require it. And there is no illocutionary justification based in a contextual asymmetry for withholding the second conjunct of the Moorean assertion from Ben.

We're assuming that it must be an illocutionary abuse to tell someone that *p* when you mistrust your own judgment on the question whether *p*, but Sheila's case shows that such mistrust need not be present in Moorean assertion. Asserting 'p but I don't believe that p' can be non-abusive toward an overhearer simply because the speaker fails to judge that *p*. Though it may create a prima facie appearance of self-mistrust, the second conjunct—'but I don't believe that p'—need not actually manifest any self-mistrust. The J-account would explain why the speaker does not make this judgment. By the J-account, another way for such a speaker to say how things stand is—with appropriate redirection—'p, but I don't know that p.'

We thereby get a unified treatment of the 'belief' and the 'knowledge' versions of Moore's paradox.<sup>48</sup> Proponents of the A-account claim as a virtue of their account that it explains both of these versions of Moore's paradox.<sup>49</sup> But their explanation reveals that they have misidentified the explanandum, since they regard all Moorean

assertions as equally odd. This generates a further reason to reject the A-account, beyond the reasons given in earlier sections: it cannot explain the difference between an abusive Moorean assertion and a Moorean assertion that is not abusive through redirection. Assuming that my treatment of Sheila's assertion can generalize to similar cases in intuitively compelling ways, we get the important result that there are fairly ordinary situations in which a Moore-paradoxical assertion is not paradoxical at all but simply the most direct way accurately to express your state of mind. Our intuitions of absurdity in Moore's Paradox concern not what is said or even the saying of it but the strictly illocutionary question of *how that saying is addressed*.

There is another difference between a Moorean assertion that manifests self-mistrust and a Moorean assertion that manifests a presumption that the addressee's doxastic context is less demanding than the speaker's. When we view these assertions from an overhearing perspective like Ben's above, we see that the first manifests a rationally incoherent state of mind while the second does not. We can revealingly frame this contrast by focusing on the speaker's claims to epistemic authority. Sheila's assertion does not—from the third personal perspective of an overhearer—undermine her claim to epistemic authority. But a self-mistrustful Moorean assertion does undermine the speaker's claim to epistemic authority on the matter asserted—to addressee and overhearer alike. While undermining one's own authority in this way need not be an illocutionary abuse from an overhearing perspective, the overhearer must nonetheless regard such a self-mistrustful speaker as confused, since she is in a rationally incoherent state of mind. In this respect there is no difference between an overhearing perspective and the perspective of the speaker's addressee.

There is no inherent error—neither an illocutionary error nor a rational error—in failing to believe what you tell someone in the way that Sheila fails to believe what she tells Aaron. And it is not necessarily irrational to mistrust your own faculty of judgment in a way that prevents you from forming a judgment. But it is rationally incoherent to mistrust your own judgment in the active way that a Moorean assertion presents you to your addressee as doing—that is, to judge that *p* but, without abandoning the judgment, fail to believe that *p*.<sup>50</sup>

How might this diagnosis of the rational incoherence manifested by self-mistrustful Moorean assertions parallel our earlier diagnosis of the incoherence manifested by lottery assertions? When we focus on the speaker's claim of epistemic authority, we can see clearly one respect in which a self-mistrustful Moorean assertion resembles a lottery assertion. In a self-mistrustful Moorean assertion, as we've just noted, it is hard to see how the speaker could without confusion claim on behalf of her judgment the epistemic authority that her assertion represents it as having. In a lottery assertion, as we noted earlier, it is hard to see how the speaker could without confusion judge what she asserts. These are not precisely the same failures. But the extended J-account lets us see a common element. The speaker cannot coherently judge what she asserts when she makes a lottery assertion because there is no possible interlocutor to whom she can without confusion claim requisite epistemic authority. While the speaker of a self-mistrustful Moorean assertion does profess to judge what she at the same time asserts that she does not

believe, the second part of her act makes her seem confused in the first part. Does she really judge what she represents herself as judging? Like a lottery assertion, a self-mistrustful Moorean assertion makes the speaker look oddly insincere—even to a third party. The problem again lies in the difficulty in understanding how she could actually judge what she professes to judge. How could she judge that *p* when she so actively worries that she is not entitled to the judgment?

Lottery assertions and self-mistrustful Moorean assertions have this important feature in common: each makes the speaker appear confused about an aspect of the second-personal epistemic authority that I've argued lies at the core of both assertion and judgment. When you make a lottery assertion you appear confused about the aspect that requires that there be some possible addressee toward whom the authority could be exercised. When you make a self-mistrustful Moorean assertion you appear confused about the aspect that requires that the authority be exercised with appropriate self-consistency—in particular, that you presume only as much authority as you feel confident you're entitled to. In each case, the problem appears to be that you do not take appropriate responsibility for personal relations implied by the claim to authority that you stake when you assert. In a lottery assertion you appear to abjure responsibility for possible relations in which you would stand to addressees whose epistemic needs you could imagine yourself engaging. In a mistrustful Moorean assertion you appear to abjure responsibility for the self-relations that such engagement would require of you.

The akratic predicament that you represent yourself as suffering in a Moorean assertion is thus not merely irrational but more broadly irresponsible. That is the pragmatic contradiction at the core of all Moorean assertions, a contradiction in the speaker's claim of second-personal epistemic authority. It is an Austinian abuse to claim a second-personal species of authority while explicitly directing your addressee's attention to something that undermines that authority. The abuse may be explicable in a way that absolves you from the confusion I just described, as a Moore-assertive Sheila would appear third-personally. Or it may be more puzzlingly self-undermining in a way that transforms irrationality in your self-relations into confusion about how you relate to others.<sup>51</sup>

## Notes

<sup>1</sup> In this paper I'll use "assert what you know (believe)" to mean *assert the same thing as you know (believe)*—that is, *assert a proposition that you also know (believe)*—and never to mean *assert that you know (believe)*. The knowledge account of assertion does not, obviously, say that you should go around asserting *that* you know stuff (or that this is the only thing you should assert!). Avoiding the lexical ambiguity would require repeating cumbersome formulations where I presume the intended meaning is perfectly clear. (Likewise for "know what you assert"; this will mean *express knowledge when you assert* and never *know which proposition you are asserting*.)

For the knowledge account of assertion, see (for starters) Peter Unger, *Ignorance: A Case for Scepticism* (Oxford: Oxford University Press, 1975), Chapter VI; Michael Slote, "Assertion and Belief," in Jonathan Dancy (ed.), *Papers on Language and Logic* (Keele: Keele University Library, 1979); and especially Timothy Williamson, *Knowledge and Its Limits* (Oxford: Oxford University Press, 2000), Chapter 11; and Keith DeRose, "Assertion, Knowledge, and Context," *Philosophical Review* 111 (2002), 167–203. I cite some more recent work in note 6 below.



<sup>2</sup> In section V, we'll also consider the 'knowledge' version: 'p but I don't know that p.'

<sup>3</sup> For this form of illocutionary 'unhappiness,' see J. L. Austin, *How To Do Things With Words*, second edition (Cambridge: Harvard University Press, 1975), 13–17. Austin himself used the term 'abuse' to mark a failure by either interlocutor, but especially by the speaker, to live up to the norms inherent in the nature of the illocutionary act in play. We assume that the act has 'come off'—that is, has been performed without a 'misfire'—and now the question is what further requirements this performance constitutively places on speaker and addressee. Austin called one category of abuse "Insincerities," but his editors note that in other versions of the lectures Austin called this category "Dissimulations" (18). (Note that my own argument will raise a deeper question about Austin's terminology. On page 40 Austin claims that "one element in lying" is not believing what you assert and that this lack of belief brings lies into the category of Insincerities. I'll question whether not believing what you tell someone must amount to insincerity. Austin may have missed an interesting complexity in this category of abuse.)

<sup>4</sup> Another reason not to rest with the 'knowledge' formulation is that many philosophers think that knowing that p involves believing that p, in which case believing that p would figure in the standard for judging that p and, since I assume that you cannot believe that p without judging that p, judging that p would absurdly itself figure in the standard for judging that p. The view I'll defend is not open to this objection. It is equally important not to interpret the view as involving the thesis that judging is *claiming* to know (or to have conclusive warrant) since that suggests that judging is *judging* that you know (or have conclusive warrant), and an apparently vicious regress ensues. (On the second point, see note 40 below.)

<sup>5</sup> As Bernard Williams observes in *Truth and Truthfulness* (Princeton: Princeton University Press, 2002), one thing that may lead a speaker to reflect on whether he really does believe what he has just spontaneously asserted, beyond mere curiosity whether it's true, is "that he is engaged in trustful conversation with another who relies on him, and the question is whether he can give that person to believe the proposition. In doing that, he may well, in such a case, give himself to believe it as well. It is the presence and needs of others that help us to construct even our factual beliefs" (193–4). I return to this observation in note 38 below.

<sup>6</sup> Consider the alternatives to the knowledge account of assertion developed by Igor Douven ("Assertion, Knowledge, and Rational Credibility," *Philosophical Review* 115 (2006), 449–485) and Jennifer Lackey ("Norms of Assertion," *Noûs* 41 (2007), 594–626). These accounts emphasize the difference between your knowing that p and its being "reasonable" for you to believe that p (Lackey) or p's being "rationally credible" for you (Douven)—arguing that the latter, not knowledge, is all that's required for you to avoid criticism in asserting that p. Douven considers the possibility that the norm requiring rational credibility does not require actual belief (460–1), but his treatment assumes that the only sort of case in which you might be entitled to assert what you do not believe is a case in which you will soon believe it, or in which you are otherwise on the verge of believing it, because the proposition in question is rationally credible for you. Lackey argues from cases in which the speaker is stipulated not to believe what she asserts—where the key assumption is that she would nonetheless be reasonable in believing it. The kind of case we'll discuss dispenses with both assumptions: the speech act does not warrant criticism despite our stipulating that the speaker (i) does not believe (or stand on the verge of believing) what she asserts and (ii) would not (either by her own lights or objectively) be reasonable in believing it. (For a bit more discussion of Lackey's argument, see notes 8 and 21 below. I agree with the broad thrust of her criticisms (section 4) of Matthew Weiner's appeal to the notion of 'secondary propriety' in his attempt to avoid altogether the idea that there are epistemic constraints on assertion ("Must We Know What We Say?" *Philosophical Review* 114 (2005), 227–251).)

<sup>7</sup> The first assumption, thus interpreted, may require some elaboration. It is of course possible to be untruthful without being insincere—when, for example, your untruthfulness derives from an error that you have not detected. Here you would fail to manifest the virtue of accuracy, a way of being untruthful that need not involve insincerity. And it is possible to be uninformative without being insincere—when, for example, you have an erroneous view of what your conversational context presupposes. But we tend to assume that sincerity is at least a necessary condition for asserting truthfully and for asserting informatively (again, in the aretaic senses). An insincere speaker is not, we assume, asserting truthfully (in the aretaic sense). And an insincere speaker is not, we assume, asserting informatively (in the aretaic sense): we tend to assume that an insincere speaker is intentionally failing to provide the information



that her addressee needs from her. Even if it is in rare cases possible for an insincere speaker to assert truthfully or to assert informatively, we tend to assume that sincerity is a necessary condition for being both truthful and informative.

<sup>8</sup>We need the parenthetical qualifications because we need to set aside cases in which the speaker's failure to believe does not reflect her assessment of available evidence. I'm inclined to think that Lackey's cases of "selfless assertion" ("Norms of Assertion," section 2) fall into this category. More explicitly, I'm inclined to think that Lackey's cases face a dilemma: either the speaker fails to believe what she asserts for non-epistemic reasons (she 'cannot bring herself' to believe something for which she takes herself to have sufficient warrant) or her assertion is a violation of the norm for assertion (because it is not properly informative in the respect we're about to discuss). (For a bit more on Lackey's cases, see note 21 below.)

<sup>9</sup>Another part of that ideal is accuracy, which informs a norm requiring you to take the care appropriate to ensuring that your beliefs are true. This is not an ideal or norm specifically of assertion.

<sup>10</sup>I won't keep repeating this qualification. See note 8 above.

<sup>11</sup>We could expand this definition to include cases in which the 'conclusion' is non-inferential and therefore does not actually emerge from any doxastic deliberation: though you don't in such cases conclude in any inferential sense that *p*, you do conclude in the sense that you treat the question whether *p* as settled and are thereby disposed not to deliberate whether *p*. The definition could also be expanded to include the state of doxastically having concluded that *p*. Since nothing in my argument will depend on my not making these expansions explicit, making them explicit would pointlessly complicate discussion.

<sup>12</sup>We lack space for more than a brief treatment of the role of self-trust in the intrapersonal dynamic informing judgment and belief. For fuller discussion, see my "Judging as Inviting Self-Trust," in preparation.

<sup>13</sup>In a fuller treatment we could allow for degrees of credence by specifying degrees of doxastic-deliberative closure or commitment: when you treat your evidence as ruling out a range of relevant alternatives to your knowing, you count as doxastically committed to that degree. (For a discussion of how this notion measures up to the traditional notion of subjective probability, see my "Reflection, Disagreement, and Context," *American Philosophical Quarterly*, 49:2 (April 2012).)

<sup>14</sup>As we'll see in section III, my approach to judgment builds on Gilbert Harman's approach to 'full acceptance' in *Change in View* (Cambridge, Mass.: MIT Press, 1986). Though I don't think judgment is the same as belief or full acceptance (since I think it is possible to judge that *p* without believing or fully accepting that *p*), what I say about judgment is compatible with what Harman says about belief and full acceptance here:

Belief in or full acceptance of *P* involves two things. First, one allows oneself to use *P* as part of one's starting point in further theoretical and practical thinking. Second, one takes the issue to be closed in the sense that, when one fully accepts *P*, one is no longer investigating whether *P* is true. Granted, one may continue investigating in order to get evidence that will stand up in court or for some other reason, but one is no longer investigating in order to find out whether *P* is true. In fully accepting *P*, one takes oneself to know that *P* is true. (*Change in View*, 47)

My departure from Harman lies in my thesis (articulated in sections III and IV as the 'extended J-account') that your *concluding* deliberation involves your being *committed* to the conclusion in a way that goes beyond your merely ceasing to deliberate or investigate. I'll argue that the commitment that defines judgment amounts to a kind of assertoric commitment, since it involves treating yourself as committed to possible interlocutors.

<sup>15</sup>I discuss cases like these from another angle in "Assurance and Warrant," *Philosophers' Imprint*, forthcoming. I use them there to engage the epistemology of testimony rather than the natures of assertion and judgment.

<sup>16</sup>This use of terminology from the contextualism debate will naturally raise questions that we lack space to discuss. But here are two brief elaborations. First, I'll speak of A's 'doxastic,' rather than of his 'epistemic' context, but I'll speak of his 'epistemic needs' because 'doxastic needs' would be misleading, suggesting the controversial thesis that A may weigh practical considerations in doxastic deliberation. Second, I'll use 'doxastic standard' and 'epistemic standard' interchangeably, and I'll remain neutral on how we should understand the nature of the shift across contexts, sometimes speaking as if it is

the standard that shifts and sometimes as if it is the range of relevant alternatives. (For a discussion of this last issue, see Jonathan Schaffer, "What Shifts? Thresholds, Standards, or Alternatives?," in Gerhard Preyer and Georg Peter (eds.), *Contextualism in Philosophy* (Oxford: Oxford University Press, 2005), 115–130. Like Schaffer (see also his "Contrastive Knowledge," in Tamar Szabo Gendler and John Hawthorne (eds.), *Oxford Studies in Epistemology*, Volume 1 (Oxford: Oxford University Press, 2005), 235–271) and like Edward Craig (whose *Knowledge and the State of Nature* (Oxford: Oxford University Press, 1990) Schaffer cites), the present account treats deliberative questions as paramount (Schaffer calls them questions of inquiry). But it is nonetheless compatible with alternative conceptions of the shift.)

<sup>17</sup> In a slightly different case we might imagine Siegfried spelling out the assurance: 'Look, this isn't a question of life or death. It's not as if you're wondering whether you have sufficient evidence to believe that your spouse has an incurable illness. There, the standard really is extremely high. While the standard is high in the present case—higher than it would be if we were talking about someone else's marriage—you have a superstitious view of what's at stake. If he's having an affair, as it's clear enough he is, the two of you will deal with it.' A friend making this sort of intervention typically won't say all this, of course. And the point is that Siegfried needn't. The force of the elaboration is already contained in the assurance that he gives in the initial act of telling.

<sup>18</sup> Some may question this claim because they view the presence of the gossiped-about as lowering the standard—since an erroneous belief may now at least be challenged by the person it concerns. That strikes me as an odd view; I find myself inclined to believe much of the gossip that I read about distant celebrities, but much less of the gossip I hear about acquaintances with whom I interact and about whom I feel a corresponding obligation to keep an open mind. Still, nothing turns on the question. Readers who dispute my claim may simply alter the example to yield the low-stakes-to-high-stakes speaker-addressee dynamic that I'm assuming.

<sup>19</sup> If you find it too difficult to imagine that Sheila asserts the same proposition that she fails to believe here—because of the context-sensitivity of 'free' ('not enough to taste' versus 'not even a trace')—switch to an example without such a term. Imagine that the issue for S and A is whether the home of a friend of A's child contains a cat or a dog. The friend has issued party invitations, and A needs to know because he has a mild allergy to cat dander. In light of A's need S can assure him, based on what her reliable child has told her, that the animal is a dog. But S herself has a morbid fear of cats: confrontation with a cat would send her into a panic that would affect her emotions for days. So she reasonably does not regard her (very good) evidence as settling the question for her.

<sup>20</sup> Bernard Williams and Richard Moran have cast doubt on whether your sincerity in asserting should be understood simply in terms of your directly manifesting your belief. Williams focuses on two sets of cases: cases in which your interlocutor does not deserve the truth from you and cases in which your interlocutor deserves more than the (bare minimally literal) truth (*Truth and Truthfulness*, Chapter 5). Moran emphasizes an aspect of the second-personal dimension I'm developing, arguing for a distinction between believing the speaker and believing merely what the speaker believes. But his counterexamples to the thesis that sincere assertion requires manifesting a belief with that content are cases in which it requires misrepresenting what you judge true (or believe) because your actual judgment is repressed or otherwise inaccessible to you ("Problems of Sincerity," *Proceedings of the Aristotelian Society*, 105 (2005), esp. 340–1). We may accept these important corrections to over-simple depictions of sincerity while noting that they do not address the issue I'm raising. The issue I'm raising addresses not complexities in the relation between sincerity and what you judge true, but a respect in which sincerity need not manifest an actual judgment that you've made at all. (Moran's project here is continuous with his assurance-theoretic approach to testimony in "Getting Told and Being Believed," *Philosophers' Imprint* (2005). I develop an alternative version of that approach in "Telling as Inviting to Trust," *Philosophy and Phenomenological Research* 70:3 (May 2005), and "Assurance and Warrant," *op. cit.*)

<sup>21</sup> Note the difference between this sort of case and the sort of case that has been cited to undermine the thesis that S cannot inform A that p unless S knows that p herself. (See Jennifer Lackey, "Testimonial Knowledge and Transmission," *Philosophical Quarterly* 49 (1999), 471–90; and Peter Graham, "Conveying Information," *Synthese* 123 (2000), 365–92.) Say S is a conscientious science teacher but also a devout creationist, and instructs her students in evolutionary theory despite disbelieving the theory herself. Such a testifier can arguably allow her audience to know what she does not regard herself as

knowing and does not regard her audience as thereby coming to know. This would perhaps show two things (though I don't believe it does), neither of which pertains to the present point: (i) that a plainly insincere testifier and (ii) that an unknowledgeable testifier can nonetheless let her audience know what she tells them. In our cases, by contrast, S does regard her audience as able to come to know on her say-so and is therefore arguably not insincere (though I take no firm stand on how that term should be applied). In "Norms of Assertion" (*op. cit.*), Lackey modifies the case in an attempt to make it a counterexample to the knowledge account of assertion by stipulating that the creationist science teacher believes that she is giving her audience some evidence for what she asserts, even though she believes neither what she asserts nor (apparently) that she can thereby let her audience know what she asserts (599). The stipulation is supposed to ensure that there's no illocutionary abuse of the sort we've been considering, but we may wonder if the teacher really is being properly informative, in our sense, given that she (apparently) does not believe that her speech could contribute to her audience's coming to know what she tells them. In any case, this creationist science teacher is quite unlike our Sheila in one key respect: though she doesn't believe, the science teacher is *entitled* to believe what she asserts; but Sheila (as we're imagining her) is not entitled to believe what she asserts. My case is therefore a very different counterexample from Lackey's (even assuming that hers is a counterexample).

<sup>22</sup> The term 'second-personal' echoes Stephen Darwall's rhetoric in *The Second-Person Standpoint* (Cambridge: Harvard University Press, 2006), which I discuss critically in "Assurance and Warrant." I agree with Darwall in tracing the basis of second-personal reasons to illocutionary norms. But I disagree with him about the nature of the relevant illocutions. As I explain at length in that paper, I oppose the Kantian tenor of his project and in particular the distinction between care and respect that he uses to frame it.

<sup>23</sup> Again, I'll defend a 'conclusive warrant' account of judgment, not a 'knowledge' account. (See again note 4, as well as the argument to come in section V.) But it will in the meantime be illuminating to treat the two 'knowledge accounts'—of assertion and of judgment—in parallel. The reasons for preferring a 'conclusive warrant' to a 'knowledge' account of judgment do not apply to the knowledge account of assertion.

<sup>24</sup> If one remains doubtful, I could of course specify different details to illustrate the difference that I'm emphasizing between Sally's circumstances and Susan's (e.g. one in which A's doxastic context imposes a low standard, say because the payoff for the winning ticket is trivial, or because A doesn't need the money).

<sup>25</sup> Again (see notes 6, 8 and 21 above), here is where the cases that I'm emphasizing differ crucially from Lackey's cases of 'selfless assertion' ("Norms of Assertion," section 2).

<sup>26</sup> Can you, for example, make an insincere assertion in a diary that you emphatically intend no one, not even your own future self, to read? Yes, of course. But then it seems you must at least be imagining or fantasizing a reaction to your assertion, and this perspective would count as that of your addressee. Without such a mode of address, it's hard to see what point there could possibly be to the insincere assertion. Perhaps, in a rather esoteric species of self-deception, you can without such fantasy 'tell yourself' something insincerely. But 'telling yourself' that p is not the same as judging that p (cf. note 32 below). Are you trying to 'brainwash' yourself into believing that p by saying it over and over? That is not at all the same as asserting that p.

<sup>27</sup> For a parallel observation about Moorean mere assertions, see note 50 below.

<sup>28</sup> And see again notes 4 and 23.

<sup>29</sup> With the act of mere assertion now in play, we can diagnose how proponents of the A-account overlook the J-account. When you tell A that p, you assert that p but do not necessarily judge that p. Because you do not necessarily judge that p, your act of telling is not as such constrained by a knowledge norm. When you merely assert that p, by contrast, your act *is* constrained by a knowledge norm, but only because *mere* assertion—assertion with no one told—necessarily expresses a judgment. We can thereby see how an account that focuses on mere assertion instead of telling may fail to grasp the connection between assertion and norms of judgment.

<sup>30</sup> *Knowledge and Its Limits*, 246.

<sup>31</sup> If you should judge that p only if you know that p, then if you do not know that p you should adopt some sub-doxastic attitude towards the proposition, such as hypothesis, supposition, or 'mere' acceptance. Do adjust your dispositions to reflect your subjective probabilities. But do not commit

yourself to the proposition in the way characteristic of judgment. (What sort of commitment is that? The extended J-account gives one answer.)

<sup>32</sup> 'I kept telling myself that p,' does not mean that the speaker judged that p. It means that she tried (unsuccessfully) to get herself to judge that p. (For more on this distinction, see my "Judging as Inviting Self-Trust.")

<sup>33</sup> One could view this as a version of Gilbert Harman's view of judgment (or 'full acceptance'). For the contrast between my approach and Harman's, see note 14 above.

<sup>34</sup> We can use a further variant of our standard case to illustrate this deeper argument. Imagine a new speaker, Sophie, who speaks from a low-standard context (like Sabine's) but who (unlike Sabine) does not have anyone who shares her context actually available as interlocutor. It seems that Sophie can nonetheless coherently judge that a given ticket drawn from that lottery has lost—simply because she can coherently imagine informing such an addressee that the ticket has lost. But Susan cannot coherently even imagine such an addressee. This difference in their merely possible interlocutors seems to explain the difference in what Susan and Sophie can judge.

<sup>35</sup> Remember: as we're using the terms, reference to a given 'doxastic context' is equivalent to reference to the 'epistemic needs' of those who are in that context. See again note 16 above.

<sup>36</sup> For an argument that the point of the concept of knowledge is to pick out potential informants, see Edward Craig's *Knowledge and the State of Nature*. But Craig does not consider the issue of context-sensitivity that I'm emphasizing. My approach also echoes John MacFarlane's approach to the semantics of knowledge-ascriptions. Indeed, my second-personal emphasis yields a basis for MacFarlane's distinctive view, semantic relativism about knowledge-ascriptions—or, as we might put it from our different angle, ascriptions of warrant for concluding deliberation with a judgment. (For an articulation and defense of such relativism, see John MacFarlane, "The Assessment Sensitivity of Knowledge Attributions," in Tamar Szabo Gendler and John Hawthorne (eds.), *Oxford Studies in Epistemology*.) The relativism emerges as follows (though this argument does not resemble any of MacFarlane's arguments). Working from an example like that of Siegfried and Ann in section I, imagine that S refrains from telling A that p because the standard is too high for S to be able simply to *inform* her interlocutor that p. So S is treating the context of a possible interlocutor—that is, of someone distinctively depending on S's status as knowing that p—as the context to which the truth value of the knowledge claim is sensitive. Why think of a possible interlocutor as someone distinctively dependent on the speaker's status as knowing? A quick answer is that only one who knows can inform. But we've seen why that is too quick: S can inform A that p, by telling A that p, without even judging (much less knowing) that p—as long as S would judge (and know) that p if she were in A's doxastic context. So we need a longer answer that draws on the J-account: when S judges that p she aims to meet the epistemic standard determined by the context of a possible interlocutor—that is, of someone wondering whether S can let him know. That confirms not sensitive invariantism, which treats the standard as sensitive to the subject's context, or contextualism, which treats it as sensitive to an ascriber's context, but relativism, which treats it as sensitive to the context of one concerned to assess the knowledge-ascription as true or false.

<sup>37</sup> The emphasis on discussion complicates matters. As we'll see presently, it will not help you close doxastic deliberation to imagine yourself entitled to tell someone whose doxastic context is undemanding because he has no interest or stake in whether you are right. Such a person is not an *interlocutor* in the fullest sense. Nor will it help you to imagine yourself entitled to tell someone whose doxastic context is less demanding than your own. We'll consider the latter issue fully in section IV.

<sup>38</sup> Note that we've returned to Bernard Williams's observation (see note 5 above) that a speaker might 'give herself' to believe that p by asking whether she can 'give another' to believe that p. But I'm suggesting that we go beyond Williams and hypothesize that the epistemic needs of others provide the point of our having a concept of judgment, of actually concluding doxastic deliberation and thereby *committing* ourselves to the proposition rather than merely suspending deliberation without such commitment. Without the second-personal dynamic that I'm emphasizing, we may wonder what point there would be not only to assertion but to the notion of a doxastic concluding that assertion purports to express.

<sup>39</sup> For a broader defense of the J-account, see my "Judging as Inviting Self-Trust." I don't claim that either reply is the only way to elaborate the accounts of assertion and judgment given thus far—as if the theses defended in sections II and III impose the entire elaboration on their own. I claim only that

each elaboration is a natural extension of those earlier discussions and conveys distinctive attractions to the J-account, as well as inevitable further controversies.

<sup>40</sup> Talk of a ‘feeling’ here follows David Owens’s talk of the ‘impression’ of sufficient warrant that enables you to close deliberation (*Reason Without Freedom* (London: Routledge, 2000), 35 and Chapter 3). Owens convincingly argues that such a feeling or impression cannot take the form of a judgment, but he also argues that it must be ‘non-reflective,’ a conclusion that I challenge in “Judging as Inviting Self-Trust.”

<sup>41</sup> Again, though this formulates the point in terms that suggest the ‘assurance view’ of testimony (see note 20 above), the point does not itself depend on the controversial features—how could a mere assurance generate epistemic warrant?—of that approach. One need not accept the assurance view of testimony to accept the present point, or indeed any claim in this paper. This paper takes no position on the nature of testimonial warrant.

<sup>42</sup> Moran, “Self-Knowledge: Discovery, Resolution, and Undoing,” in *European Journal of Philosophy* 5:2 (1997), 144; and *Authority and Estrangement: An Essay on Self-Knowledge* (Princeton: Princeton University Press, 2001), 71.

<sup>43</sup> Seymour Hersh, “The Iran Plans,” *The New Yorker*, 17 April 2006.

<sup>44</sup> For this argument, see my “Judging as Inviting Self-Trust.”

<sup>45</sup> This is Moore’s term, which is ambiguous between the *error*-reading and the *incoherence*-reading (‘A Reply to My Critics,’ in P. Schilpp (ed.), *The Philosophy of G. E. Moore* (La Salle: Open Court), 543).

<sup>46</sup> Could you pull it off by getting your single interlocutor A to think you were addressing one or both of the two conjuncts only to yourself as inward musing? For a discussion of such indirection from a different but related angle, see my “Telling as Inviting to Trust,” section II. Such indirection is rhetorically more complex than the redirection described in text, since it requires that the speaker herself stand in for the second addressee.

<sup>47</sup> Note well that I’m not claiming merely that it can be non-abusive to assert each conjunct in succession, provided that you switch audiences at the ‘and.’ I’m claiming that it can be non-abusive to assert the entire conjunction to a single audience—provided that that audience is not the intended addressee, but a mere overhearer, for one of conjuncts. It is the distinction between addressee and overhearer that enables us to make sense of this possibility, a distinction not in play in the traditional literature on Moore’s Paradox.

<sup>48</sup> For the knowledge version, see G. E. Moore, *Commonplace Book: 1919–1953* (London: Allen & Unwin, 1962), 277.

<sup>49</sup> See, for example, DeRose, “Assertion, Knowledge, and Context,” 180–1.

<sup>50</sup> The incoherence would arise even in a Moorean *mere* assertion (in the sense distinguished at the end of section II)—an assertion with no one told. Judging that p without believing that p would be irrational in a way that mirrors the irrationality of judging, all things considered, that one ought to  $\varphi$  at  $t$  without choosing or intending to  $\varphi$  at  $t$ . Some philosophers regard such practical akrasia as possibly rationally coherent—say, when Huck Finn mistrusts his manifestly untrustworthy judgment that he ought to betray Jim—and on similar grounds might regard mistrust in an untrustworthy doxastic judgment as possibly rationally coherent. (See, for example, Nomy Arpaly, “On Acting Rationally against One’s Best Judgment,” *Ethics* 110 (2000), 488–513, and Karen Jones, “Emotion, Weakness of Will, and the Normative Conception of Agency,” in A. Hatzimoysis (ed.), *Philosophy and the Emotions* (Cambridge: Cambridge University Press, 2003), 181–200. But neither author discusses the doxastic case.) In “Rational Requirements and ‘Rational’ Akrasia,” *Philosophical Studies*, forthcoming, I argue that we can explain the intuition that there is something rationally praiseworthy in mistrusting a manifestly untrustworthy judgment, while also insisting that the akratic state itself is rationally incoherent, via the independently defensible thesis that one can reason ‘upstream’ from an akratic failure to act on a judgment to abandoning the judgment, just as one might reason ‘downstream’ to overcome a bout of akratic temptation. But nothing in the present paper turns on whether akrasia is necessarily rationally incoherent. It would preserve the force of the present argument if I claimed merely that such akrasia is *typically* irrational.

<sup>51</sup> I wrote this paper for a workshop on Assertion and Sincerity organized by David Owens at the University of Sheffield in December 2009. Thanks to workshop participants (especially to Paul Faulkner, Michael Ridge, and Roy Sorensen) for helpful feedback. Thanks also to a very helpful referee for *Noûs*.