

9

Linguistic Intuitions in Context: A Defense of Non-Skeptical Pure Invariantism

John Turri[†]

He either knows for certain, or he heard from the right people. I will ask him. And he will tell me.

—Batman¹

9.1. Introduction

Epistemic invariantism is the view that the truth-conditions of knowledge ascriptions don't vary across contexts. Epistemic purism is the view that purely practical factors can't directly affect the strength of your epistemic position. The combination of purism and invariantism, pure invariantism, is the received view in contemporary epistemology. It has lately been criticized by contextualists, who deny invariantism, and by impurists, who deny purism. A central charge against pure invariantism is that it poorly accommodates linguistic intuitions about certain cases. In this chapter, I develop a new response to this charge. I propose that pure invariantists can explain the relevant linguistic intuitions on the grounds that they track the propriety of indirect speech acts, in particular indirect requests and denials. In the process we learn

[†] For helpful conversations and feedback on this paper, I thank Anthony Booth, Ian MacDonald, Rachel McKinnon, and Darrell Rowbottom. As always, special thanks go to Angelo Turri. This research was supported by the Social Sciences and Humanities Research Council of Canada, the National Endowment for the Humanities, the British Academy, the Association of Commonwealth Universities, and an Ontario Early Researcher Award.

¹ Quote-o-rama, "Batman: As the Crow Flies," available at <<http://www.otd.com/~paul/Quote/batman.html>> [last accessed October 21, 2013].

an important methodological lesson about how to effectively marshal linguistic intuitions in favor of—or in opposition to—a theory.

First I will explain pure invariantism and its rivals, contextualism and impurism (section 9.2). Then I will present a central argument against pure invariantism (section 9.3). Next I will review some previous responses to the argument (section 9.4). Lastly I will develop a new response (sections 9.5 and 9.6) and conclude with a methodological lesson (section 9.7).

9.2. The Received View Explained

One of the liveliest philosophical debates over the past two decades concerns the semantics of knowledge ascriptions. Competing views here divide into two camps: contextualist and invariantist.

Contextualists claim that the truth-conditions of knowledge ascriptions, such as ‘I know that *P*’ or ‘He doesn’t know that *P*’, vary with features of the speaker’s context.² The hallmark of contextualism—“the result that contextualists insist on” (DeRose 2009, p. 61)—is that a speaker in one context could truly say ‘*S* knows that *P*’, while simultaneously a speaker in a different context could refer to the same person at the same time with respect to the same proposition, and truly say ‘*S* doesn’t know that *P*’. It is possible that some such disagreements are merely verbal, contextualists say, because the speaker’s context determines how strong an epistemic position *S* must be in with respect to *P* in order for the speaker to speak truthfully by saying ‘*S* knows that *P*’. Consequently, *S* might meet the laxer requirements selected by the one speaker’s context but fail to meet the stricter requirements selected by the other speaker’s context, and this is why they could both speak truthfully.

Invariantists deny the hallmark of contextualism. The truth-conditions of knowledge ascriptions, according to invariantism, are insensitive to features of the speaker’s context. Every context selects the same requirements for being in a strong enough epistemic position for knowledge. No matter how different our respective contexts are, if you say ‘*S* knows that *P*’, and I say ‘*S* doesn’t know that *P*’, and we’re simultaneously referring to the same person by ‘*S*’ and the same proposition by ‘*P*’, then our disagreement is *not* merely verbal. At least one of us speaks falsely.

A second debate has emerged over the past decade, which cuts across the contextualism/invariantism debate. It concerns the nature of knowledge itself, in particular how it relates to purely practical matters. Suppose that Naomi and David both have a true belief that *P*, that they both base their belief on equally good and compelling evidence, that they are equally reliable on the question at hand, that they are equally alert to counterevidence, and so on. In short, on all relevant truth-related matters, Naomi and David are equal. In virtue of this, let’s say that they’re in an *equally strong epistemic position* relative to *P*. Given that they both have a true belief and are in an

² A quick note on quotations: I use single quotes to mention or name expressions; I use double quotes as scare quotes and for direct quotation.

equally strong epistemic position, could it nevertheless turn out that *only one of them knows P*? Could it turn out that, say, Naomi knows *P*, but David doesn't? Could a purely non-epistemic (i.e. a "practical") feature of David's situation prevent him from knowing *P*? For example, if David has more at stake than Naomi does on whether *P* is true, and holding all else equal, could that prevent David from knowing *P*? Purists answer 'no'. Impurists answer 'yes'.³ The hallmark of impurism is that "what makes a true belief into knowledge is not entirely an epistemic matter" (Stanley 2005, p. 2).

Leading contextualists endorse purism.⁴ By contrast, a debate rages between purists and impurists in the invariantist camp. Impure invariantists claim that although the truth-conditions of knowledge attributions are not sensitive to the speaker's context, they nevertheless are sensitive to practical features of the context of the subject under evaluation (i.e. the *S* in '*S* knows that *P*'). Pure invariantists deny *both* that the truth-conditions of knowledge attributions vary with the speaker's context *and* that they vary with changes in the practical features of the subject's context.⁵

A third debate, as ancient as the previous two are recent, cuts across all the positions discussed so far: the debate over skepticism. Do we know most, or at least many, of the things we ordinarily take ourselves to know? Or to put the matter in a way contextualists might find more congenial, do we usually, or at least often, speak truthfully when we say people 'know' things? Skeptics answer 'no', non-skeptics answer 'yes'. These are not perfectly precise views, since 'many', 'most', 'usually', and 'often' are vague terms. Skepticism comes in degrees.

We could end up with eight different positions, depending on how we settle the three debates.⁶

- Non-skeptical pure invariantism (Williamson 2005; Turri 2010a)
- Skeptical pure invariantism (Unger 1975)
- Non-skeptical impure invariantism (Hawthorne 2004; Stanley 2005)
- Skeptical impure invariantism
- Non-skeptical pure contextualism (DeRose 2009; Cohen 2005)
- Skeptical pure contextualism

³ 'Purism' is Jeremy Fantl and Matt McGrath's (2009) term. It is also called "intellectualism" (Stanley 2005; DeRose 2009). Impurism is also called "practicalism" (Grimm 2011) and "pragmatic encroachment" (Jon Kvanvig's coinage).

⁴ Keith DeRose (2009, p. 188 n.4, p. 189) remarks, "One of the intuitive attractions of contextualism is that it allows one to uphold intellectualism [i.e. purism] while delivering certain desired results about key test cases," and that his "allegiance to intellectualism" has motivated his sustained defense of contextualism. See also Stewart Cohen (2005). However, Fantl and McGrath (2009, pp. 35–6) point out that there might be hints of impurism in David Lewis's (1996) defense of contextualism. Fantl and McGrath even express some sympathy for impure contextualism (see also McGrath 2010, §6).

⁵ Timothy Williamson (2005) calls impure invariantists "sensitive invariantists" and pure invariantists "insensitive invariantists." DeRose (2009) calls pure invariantists "classical invariantists"; Jason Stanley (2005) and John MacFarlane (2005) call them "strict invariantists"; Baron Reed (2010) calls them "stable invariantists."

⁶ Here I set aside consideration of *assessment relativism* about knowledge attributions. See MacFarlane (2005).

- Non-skeptical impure contextualism (Greco 2010)
- Skeptical impure contextualism

Some of these positions are unoccupied in the current literature and will be ignored here. (It's hard to even imagine any motivation for, say, skeptical impure contextualism or skeptical impure invariantism.) Some important participants are sometimes difficult to classify. For example, Jeremy Fantl and Matt McGrath are typically classed with John Hawthorne and Jason Stanley as non-skeptical impure invariantists, but Fantl and McGrath's primary allegiance is to the less specific position of non-skeptical impurism, and they express openness to non-skeptical impure contextualism (see Fantl and McGrath 2009, p. 53; McGrath 2010, §6).

Impurism and contextualism are both "shifty" positions. They entail that either the semantic content or truth-value of knowledge ascriptions can shift in ways that traditional epistemology simply overlooks. By contrast, pure invariantism is a fully "stable" position. It entails that the semantic content and truth-value of knowledge ascriptions neither do nor can shift in the ways that impurists and contextualists claim that they can and often do.

Pure invariantism is orthodoxy. And since philosophers have traditionally been (mostly) non-skeptical, non-skeptical pure invariantism has the distinction of being *the received view* in contemporary epistemology. Impurists and contextualists are philosophical rebels. Is their rebellion justified?⁷

9.3. The Received View Contested

Why rebel against the received view? Some argue for skepticism or express sympathy for it (e.g. Unger 1975; Reed 2007; BonJour 2010). But there have always been skeptics and they have always been a distinct minority. The leading contemporary challenge to orthodoxy is motivated by linguistic intuitions about verbal behavior. Many philosophers argue that these intuitions provide strong evidence against pure invariantism. They argue that either impurism or contextualism best explains the intuitive linguistic facts.⁸

Simple thought experiments elicit the relevant intuitions. In particular they are elicited by descriptions of entirely realistic pairs of cases which seemingly reflect how

⁷ An overzealous and irreverent traditionalist might characterize the debate as "the embattled *pure hearted* fending off the combined strength of the *imps* and *cons*." A less colorful but more temperate and dignified characterization of the debate is offered in the main text.

⁸ Impurists and contextualists have an in-house debate over which view *better* explains the data, once pure invariantism has been eliminated. I won't concern myself with that debate here. The important point for present purposes is that impurists and contextualists agree that pure invariantism can't do the job. (Assessment relativists like MacFarlane (2005) argue that their view is even better than impurism and contextualism.) I should note that Stanley (2005), whom I'm lumping in with the generic opponents of orthodoxy, has a more nuanced and slightly ambivalent take on the role intuitions play in the dialectic. "These intuitions are not intended simply to be data for an epistemological theory, as the grammaticality of various sentences may be taken to be data for a syntactic theory. Rather, the role of my appeal to our intuitions about these particular cases is to make vivid our commitment to the conceptual connection between knowledge and practical reasoning" (Stanley 2005, pp. 97–8; compare p. 12). Earl Conee (2013, p. 76) also expresses some ambivalences about whether the data are best described as "intuitions" or "intuitive responses" based on something else.

competent speakers use ‘know’. One case, call it ‘LOW’, features a protagonist (‘Low Pro’) who sincerely utters ‘I know *P*’ in an ordinary “low-stakes” setting. The key intuition here is that Low Pro’s verbal behavior is natural and appropriate. The other case, call it ‘HIGH’, features a protagonist (‘High Pro’) who sincerely utters ‘I don’t know *P*’ in a “high-stakes setting.” The key intuition here is that High Pro’s verbal behavior is natural and appropriate. Importantly, the cases are set up so that Low Pro and High Pro are in equally strong epistemic positions with respect to *P*: they have the same evidence, are equally reliable, are equally alert to counterevidence; are just as confident that *P* is true. And they are otherwise similarly situated, with *one exception*: in HIGH the stakes are considerably higher because much more rides on whether *P* is true. It is perhaps a harmless oversimplification to say that the only difference between Low Pro and High Pro is that High Pro has more to worry about.⁹

Consider a concrete pair of such cases.¹⁰

LOW FLIGHT: Stewart is in the Atlanta airport, waiting to board his flight. A fellow traveler seated nearby looks up from his laptop, stretches, turns to Stewart and says, “I’ve been traveling all day and it’ll be a relief to get home to Detroit. A layover would be annoying. Say, do you happen to know whether this is a direct flight to Detroit?” With his itinerary in hand, Stewart answers, “Yes, I do—it’s direct to Detroit.”

HIGH FLIGHT: Stewart is in the Atlanta airport, waiting to board his flight. Suddenly a man dressed in a uniform and carrying a small hardshell cooler rushes down the concourse, stops in front of Stewart’s gate and breathlessly says to Stewart, “I’m an organ courier transporting a kidney to a patient in Detroit. I need a direct flight to Detroit, or the kidney will spoil. Do you know whether this is a direct flight to Detroit?” With his itinerary in hand, Stewart answers, “Sorry, I don’t know [whether it is]. You should check with an airline official.”

In each case, based on the itinerary Stewart believes throughout that the flight is direct to Detroit, and his belief is true.

The cases reflect how we ordinarily speak.¹¹ Speakers are disposed to be more generous in their ‘knowledge’ attributions in LOW cases, more sparing in their ‘knowledge’ attributions in HIGH cases, and likely to deny ‘knowledge’ in HIGH cases. This is

⁹ Jonathan Schaffer (2006) has a very different take on which data best support contextualism, focusing instead on the “contrast” and “question sensitivity” of knowledge ascriptions.

¹⁰ The pair is inspired by Cohen’s (1999, p. 58) widely discussed airport case, but whereas Cohen’s case involves *both* elevated stakes *and* the explicit mention of a specific error possibility, my cases don’t involve the mention of specific error possibilities. This will be important below when we consider Patrick Rysiew and Jessica Brown’s defense of pure invariantism.

¹¹ Or so it seems to theorists reflecting on the matter from the armchair. As with other aspects of actual patterns and tendencies in ordinary usage of ‘know’, the matter is ripe for empirical investigation. As much recent work in experimental philosophy and psychology has suggested, sometimes we’re surprised by what we find when we carefully look and see (e.g. see Beebe and Buckwalter 2010; Beebe and Jensen 2012; Beebe 2013; Weinberg et al. 2001; Swain et al. 2008; Feltz and Zarpentine 2010; Starmans and Friedman 2012; Schaffer and Knobe 2012; Sripada and Stanley 2012; Pinillos 2012; Myers-Schulz and Schwitzgebel 2013; Murray et al. 2013; Turri 2012, 2013, Under review; and Turri and Friedman Forthcoming).

all “utterly natural” and “integral to the ordinary use of ‘know’” (Williamson 2005, p. 217; compare Stine 1976, p. 274).

But why should this undermine pure invariantism? How do we go from the propriety and naturalness of such behavior to the denial of pure invariantism? According to Timothy Williamson:

Presumably, the endorsement rests on a methodological principle of charity, by which, very roughly, we should prefer to interpret speakers as speaking... truly rather than falsely (*ceteris paribus*). Shifting standards seems to give us more flexibility to assign to ‘know’ a charitable reference. (Williamson 2005, p. 220; compare Fantl and McGrath 2009, ch. 2, and Davis 2007, p. 430)¹²

Keith DeRose’s most recent work (2009, ch. 2) confirms this by clarifying “the methodology that takes us from the data to a contextualist conclusion,” or more generally, from the data to the denial of pure invariantism. The data are the “intuitive” facts that the cases feature sincere, natural and appropriate speech, along with the observation that the speech isn’t based on mistaken beliefs about the situation. DeRose calls this “the best possible type of evidence” we could have against invariantism in epistemology, and exactly the same type of evidence that leads us to reject, say, invariantism about indexicals or gradable adjectives. DeRose’s key methodological claim is that there is a “general presumption” that when competent speakers “are not basing their claims on some false beliefs they have about underlying matters of fact, how they naturally and appropriately describe a situation, especially by means of common words, will be a true description.” Consequently, DeRose concludes, it is “a bad strike against” a theory “if it rules [that such descriptions are] false, as it seems invariantism will have to rule with respect to one or the other of” High Pro and Low Pro (2009, pp. 50–1, p. 67).

In short, if we combine (a) the fact that sincere, natural, appropriate, non-misinformed descriptive uses of ‘know’ in ordinary language are guided by noticeably different standards in different contexts, and (b) the charitable methodological principle that such descriptive uses are true, then we have evidence that “militate[s] strongly” against pure invariantism.¹³ Pure invariantism is too inflexible to charitably explain all the data. It predicts that when knowledge ascriptions in different contexts superficially contradict one another, at least one of them is false. And it predicts that knowledge doesn’t come and go depending on how much is at stake. But a charitable explanation of our linguistic behavior leads us to reject at least one of those predictions. Our linguistic behavior is shifty. Pure invariantism can’t shift gears.

¹² Note that Williamson is characterizing the line of thought, not endorsing it. He defends pure invariantism.

¹³ DeRose also reports intuiting directly that the utterances in question are true (e.g. 2009, p. 49 n.2). But he doesn’t simply rest with this. He acknowledges that the intuition that a speaker’s utterance is “appropriate” is stronger than the intuition that it’s true (2009, p. 50). DeRose then proceeds to argue that appropriateness is powerful evidence for the truth, as I describe in the main text. This is a good approach because it begins with less controversial data—intuitions about the *generic propriety* of speech—upon which all parties to the debate are more likely to agree.

To help organize the discussion, I'll understand the argument against pure invariantism as follows, where 'proper' abbreviates 'sincere, natural, appropriate and not based on misinformation.'

(*Anti-PI*)

1. Low Pro's and High Pro's utterances are proper. (Premise)
2. If their utterances are proper, then both of their utterances are true. (Premise: Charity Principle)¹⁴
3. So both of their utterances are true. (From 1 and 2)
4. If both of their utterances are true, then pure invariantism is false. (Premise)
5. So pure invariantism is false. (From 3 and 4)

9.4. The Received View Defended

How should pure invariantists respond? The argument is valid, so they must reject at least one premise. This section reviews several previous responses. The next two sections develop a new response.

Here is one strategy for rejecting line 4.¹⁵ Knowing *P* requires confidently enough believing *P*.¹⁶ But high stakes induce people to be epistemically cautious, perhaps excessively so. We need more evidence to maintain the same level of confidence. Recall that High Pro and Low Pro have the same evidence. The high stakes erode High Pro's confidence to the point where he no longer knows *P*. So High Pro speaks truthfully when he denies that he knows. But this is because he fails to satisfy the invariant requirements of knowledge, not because the meaning of 'knows' or the requirements of knowledge are shift.

In response, critics deny that High Pro is best understood as losing confidence in the relevant sense. At least, the case doesn't have to be interpreted that way. There is another way of understanding it that suffices for the critic's purpose. DeRose (2009, pp. 190–3) argues that this objection to line 4 presupposes an "unstable" conception of confidence, whereas a "stable" conception is "more natural" and "correct." On the stable conception, high-stakes situations don't typically diminish our confidence. Rather, the same level of confidence is present in both LOW and HIGH. This level of confidence disposes us to act one way in LOW but more cautiously in HIGH.¹⁷

¹⁴ Alternatively, line 2 might have said, "The best explanation for 1 is that both of their utterances are true." Pure invariantist arguments against line 2 have indeed focused on finding a satisfying alternative explanation of the utterances' propriety.

¹⁵ Inspired by Kent Bach (2010), though I'm not sure he would endorse my way of putting it. The proposal is modeled on his, but he has a different sort of case (a "third-person" case) in mind when he makes his proposal. Compare also Hawthorne's "belief removal model" (2004, p. 169).

¹⁶ It has recently been disputed whether belief is genuinely required for knowledge. For example, Myers-Shulz and Schwitzgebel (2013) and Murray et al. (2013) present empirical evidence that, on the ordinary concept of knowledge, knowledge doesn't require belief. For an opposing view, see Rose and Schaffer (2013), and Buckwalter et al. (2013).

¹⁷ See also Brown (2005, pp. 147–8).

Another response is to reject line 2 on the grounds that Low Pro's utterance is false because he fails to meet the demanding standards of knowledge. Low Pro's utterance nevertheless seems proper because it's close enough to the truth for practical purposes (Unger 1975). One way of sharpening this proposal is to invoke the phenomenon of loose talk (compare Conee 2005 and Davis 2007). It would be laborious to always be perfectly precise. It's convenient to speak loosely when present purposes don't require precision. Some contexts tolerate very liberal approximations of the truth whereas others require great precision. Our variable strictness in ascribing knowledge follows the same pattern, one might argue, which explains the shiftiness in our use of 'knows' while maintaining that the truth-conditions of knowledge ascriptions are invariant.

The loose-talk strategy faces two challenges. First, it seems too hospitable to skepticism. Indeed, Unger deployed it as part of his overall case for a radical form of skepticism, to help explain why we *falsely say* that we know many things. As Conee (2005, p. 52) puts it, on this view "only the most conspicuous facts of current perception, the clearest memories, triple-checked calculations, and the like will" enable knowledge. Other things being equal, I prefer a less skeptical defense of pure invariantism.¹⁸ Second, knowledge-talk doesn't fit the profile of loose talk (MacFarlane 2005, p. 784). You say, "I'm going to the store. Do we have any coffee left?" and I respond, "No, it's all gone." If you challenge me, "Actually, there are a few grounds left in the jar," a natural comeback on my part would be, "[Of course] I meant it was *just about* all gone [and so you should buy some more]." By contrast, in typical cases where someone is challenged, "Actually, you don't know that," they don't say, "the point is that I *just about* know it," or, "[Of course] I meant that I *nearly* knew it."

Patrick Rysiew (2001, 2007) rejects line 2 for different reasons. His proposal is premised on the uncontroversial observation that an utterance communicates more than its literal content. We presuppose that our conversational partners are cooperative and, consequently, that they strive to make their speech relevant. This presumption enables us to communicate information beyond what we literally say, a phenomenon which I'll call *suggestion* or *conveyance*. Knowledge attributions are no exception: they can also suggest or convey information. On Rysiew's view, High Pro speaks properly but falsely when he says that he doesn't know. High Pro says that he doesn't know in order to avoid suggesting certain other false things. And in the context it seems more important to avoid the false suggestion than to speak the literal truth. This is why we intuit that his utterance, although false, is proper.

The details of Rysiew's influential proposal are important so let's consider it more carefully. Knowing *P* requires being in a good enough epistemic position relative to *P*. For convenience Rysiew understands being in a good enough epistemic position as being able to rule out the relevant alternatives, on some moderately strong but

¹⁸ Compare Stanley's (2005, p. 84) remarks on a contextualist version of the loose-talk strategy: "This is not a very satisfying way of 'rescuing' ordinary knowledge attributions. Indeed, one may wonder whether it has any advantages over skepticism at all."

non-skeptical and invariant understanding of ‘relevant alternatives’.¹⁹ We’re often able to rule out all the *relevant* alternatives but we’re almost never able to rule out *all* alternatives. Let ‘*C*’ name one of these irrelevant alternatives that, on a specific occasion, you can’t rule out. Despite being irrelevant, *C* might still be *conversationally salient* (‘salient’ for short). And in a context where *C* is salient, saying ‘I know *P*’ suggests that you can rule out *C*. It suggests this because *C* is salient, you’re presumed to be cooperative, and so it’s expected that you wouldn’t say ‘I know *P*’ unless you (wanted to communicate that you) can rule out *C*. So in order to avoid falsely suggesting that you can rule out *C*, you say ‘I don’t know *P*’. And since conversational salience is highly context-sensitive, Rysiew argues, this pragmatic account can explain the shiftiness of our knowledge attributions.

Rysiew’s proposal faces a challenge. It’s not generally true that when asserting *R* suggests *S*, asserting not-*R* similarly suggests not-*S*. For example, if a colleague asks me my opinion of a certain restaurant and I respond, “The best I can say about that place is that the staff is polite,” this suggests that their food isn’t good. But if I instead say, “The best I can say about that place isn’t that the staff is polite,” this does *not* suggest that their food *is* good. To take another example, if my partner asks whether we’ll be on time for the reception, and I say, “It’s possible,” this suggests that I am not confident that we will be on time. But if I instead say, “It’s not possible,” this does not suggest that I am confident that we will be on time—far from it!

In light of these points, reconsider Rysiew’s proposal. Suppose that saying ‘I know *P*’ suggests that you can rule out all the salient alternatives, even the epistemically irrelevant ones you can’t rule out, such as *C*. Now when asked whether you know *P*, one way to avoid suggesting that you can rule out *C* is to *not say anything at all*. But that would be rude. A politer way is to hedge with ‘Maybe I know, but *C* sure is hard to rule out’ or ‘I’m tempted to say “I know,” but then again, there is the possibility of *C*’. Those hedges have the advantage of *being true*. Why opt for the false ‘I don’t know’ when you could avoid suggesting the falsehood by expressing a readily available truth? This defies our expectations of competent, cooperative speakers. The upshot is that although Rysiew’s view predicts that High Pro won’t say ‘I know *P*’, it doesn’t explain why he goes so far as to say ‘I don’t know *P*’.²⁰

But does it really defy our expectations? As Brown (2006, p. 415ff) notes, it’s not uncommon for speakers to say something literally false in order to convey a relevant truth. Such speech often seems proper, as in the following example. It’s lunchtime. You’re hungry and would like some company over a meal. “Have you had lunch?” you ask me. “No, I haven’t had lunch,” I respond. Arguably what I say is literally false because I have eaten lunch *at some time* in the past; but it conveys a relevant truth, namely, that I haven’t eaten lunch *today*. I could easily have truthfully said ‘I haven’t

¹⁹ Rysiew’s approach isn’t wedded to relevant alternatives theory. This will be important below.

²⁰ Compare Fantl and McGrath (2009, p. 41). Why achieve your conversational purpose “by lying” they ask, when it could be achieved “just as well with the truth?” See also DeRose (2009, pp. 111–24).

eaten lunch today' and had the same effect, yet it still seems appropriate for me to prefer the shorter false utterance instead. Presumably it's appropriate because you aren't liable to be misled by my literal words into thinking that I had somehow, amazingly, lived over thirty years without eating lunch *even once*; rather, you can be counted on to infer that I intended to communicate that I hadn't eaten lunch today. On similar grounds, Rysiew and Brown might argue, it's appropriate for High Pro to prefer 'I don't know' to true, but less convenient formulations such as 'Maybe I know, but *C* sure is hard to rule out'.

Brown is clearly right that there is precedent for this sort of explanation. But one worry is that the precedent differs importantly from the High cases that it's invoked to help explain. For example, in HIGH FLIGHT the organ courier can't be counted on to avoid being misled by Stewart's literal words. The presumption is that the courier will believe Stewart when Stewart says 'I don't know'. And the courier has no special reason to think that Stewart is merely trying to avoid falsely suggesting that he can rule out a salient but irrelevant alternative. So whereas I know that you won't be misled when I say 'I haven't eaten lunch', Stewart doesn't have any reason to think that the organ courier won't be misled when he says 'I don't know'.

Rysiew and Brown's view faces another challenge: some high cases lack an irrelevant but salient alternative, which their view requires in order to explain the relevant behavior. The organ courier mentions no error possibility to Stewart. No one mentions that Stewart's itinerary might contain a misprint, or that the captain might mistakenly land the plane in Pittsburgh, or anything else. LOW FLIGHT and HIGH FLIGHT seem to feature exactly the same set of alternatives: *it's a direct flight to Detroit* versus *it's not a direct flight to Detroit*. It's totally implausible to suggest that *it's not a direct flight to Detroit* is relevant in LOW FLIGHT but irrelevant in HIGH FLIGHT. And it's equally implausible to suggest that *it's not a direct flight to Detroit* is salient in HIGH FLIGHT but not in LOW FLIGHT. Thus Rysiew and Brown's proposal can't explain why Stewart says 'I don't know'.

Rysiew and Brown might respond that heightened stakes naturally prompt us to start worrying about additional alternatives, even if they're not mentioned. For example, they might say it's natural that in HIGH FLIGHT Stewart's thoughts will turn to the possibility that his itinerary contains a misprint, or the possibility that the pilot will make an unnecessary, unauthorized stop. The problem with this response is that it doesn't seem *essential* to the case that Stewart begins worrying about such possibilities. Even if we stipulate that he's cool under pressure and doesn't begin to worry about those possibilities, it seems neither unnatural nor inappropriate for him to deny that he knows.

At this point it's worth recalling the generic version of Rysiew's proposal, stated in terms of *strength of epistemic position* rather than *relevant alternatives*. It's uncontroversial that knowledge requires a true belief plus a strong enough epistemic position. How strong? We're not giving an analysis of knowledge, so it's harmless to answer 'strong enough to know'. But *strong enough to know* doesn't entail *strong enough for*

everything. Some purposes might require a position stronger than what knowledge strictly requires. If that's the case, then we can explain High Pro's behavior even without a conversationally salient but epistemically irrelevant alternative, as follows. In both LOW FLIGHT and HIGH FLIGHT Stewart's epistemic position relative to *it's a direct flight to Detroit* is strong enough to know and strong enough to satisfy a stranger's idle curiosity. But in neither case is it strong enough for directing personnel in life-and-death medical matters. So Stewart knows in both cases and Stewart speaks falsely in HIGH FLIGHT when he says 'I don't know'. Nevertheless, he also speaks appropriately because he communicates that his epistemic position isn't strong enough for the courier's conversationally salient purposes.

Understood this way, Rysiew and Brown's view contrasts interestingly with the loose-talk proposal discussed earlier. Just as it was argued earlier that Low Pro speaks falsely but appropriately because present purposes recommend loose usage, here it is argued that High Pro speaks falsely but appropriately because present purposes recommend, as it were, guarded usage.

Overall, although Rysiew and Brown's defense of pure invariantism is impressive, it faces difficult challenges. I'll raise two further concerns about their view before presenting my own proposal in the next section.

First, it's noteworthy that competent speakers in High cases respond similarly whether you ask them 'P?' or 'Do you know whether P?' (see Turri 2010b, 2011). Consider this exceedingly minor revision of HIGH FLIGHT.

Stewart is in the Atlanta airport, waiting to board his flight. Suddenly a man dressed in a uniform and carrying a small, hardshell cooler comes rushing down the concourse, stops in front of Stewart's gate, and breathlessly says to Stewart, "I'm an organ courier transporting a kidney to a patient in Detroit. I need a direct flight to Detroit, or the kidney will spoil. Is this a direct flight to Detroit?" With his itinerary in hand, Stewart answers, "Sorry, I don't know. You should check with an airline official."

The only difference here is that the courier asks Stewart 'Is this a direct flight to Detroit?' instead of 'Do you know whether this is a direct flight to Detroit?' The very same answer, 'I don't know', serves equally well and seems to have the same effect on the conversation in both versions of the case. It might be too much to ask for an identical explanation of both versions, but we should expect the explanation to be similar for both. It's not clear that Rysiew and Brown's view can offer this.

Second, Rysiew says that it's "essential" to his proposal that "our untutored intuitions about the truth conditions of various sentences are generally insensitive to the semantic/pragmatic distinction" (2007, p. 648; cf. p. 660 n.31). That is, our intuitive estimation of speech as true or false doesn't distinguish the truth-value of what we literally say, on the one hand, from the truth-value of what we suggest, on the other. Instead we "tend to identify" the truth-value of an utterance with the truth of "the most salient proposition a speaker actually communicates" in saying it (Rysiew 2001, p. 487). Moreover, Rysiew extends this point to explain why some mistakenly *believe* that they don't know certain

things, as follows (2001, pp. 502–3). Saying ‘I know that I’m not a brain in a vat’ would falsely suggest that I can rule out the irrelevant possibility that I am a brain in a vat.²¹ As a result, I not only refrain from saying ‘I know I’m not a brain in a vat’, but I also mistakenly “come to believe” that I don’t know that I’m not a brain in a vat.

Suppose that attributing such confusion to us is essential to Rysiew’s view.²² Then one might question whether the anti-skeptical preference that Rysiew and I share is properly motivated. G. E. Moore (1959, p. 193ff) pitted the obviousness of commonsense knowledge attributions against the skeptic’s wherewithal. And many non-skeptical epistemologists follow him in that regard. But if it’s correct that we begin doing epistemology with long habits of mistaking false knowledge ascriptions for true ones, and true knowledge ascriptions for false ones, then that poisons the well of Moorean data, thereby weakening Moore’s hand—perhaps even to the point where he no longer knows that he has one.

The objections and concerns I’ve raised might not debilitate Rysiew and Brown’s view. I don’t claim that they are insurmountable. But they are enough to motivate me to look for an alternative response to Anti-PI. Of course, alternative responses needn’t compete with one another. They could be complementary.

9.5. Indirect Speech Acts

The responses to line 2 covered in the last section are widely labeled, in Keith DeRose’s memorable phrase, “warranted assertability maneuvers” or “WAMs” for short (DeRose 1999, p. 196ff; 2009, p. 83ff). Brown helpfully encapsulates the essence of a WAM:

At the core of a WAM is the idea that the intuitions about contextualist cases [e.g. High/Low pairs] can be explained by appeal to the truth-value of the propositions pragmatically conveyed by knowledge attributions, rather than the literal truth-value of those attributions. (Brown 2005, p. 150)

A WAM is one way to develop a pragmatic account of such cases. But it isn’t the only way. In the remainder of this section, I’ll lay the groundwork for a pragmatic account that doesn’t fit the WAM model, based on speech-act theory.²³

Orders and requests are ways to direct people. More specifically, they are speech acts we perform in order to direct people. Sometimes we direct them overtly by saying things like ‘pass the salt’ or ‘give me the information.’ We could even use an explicit performative, as in ‘I request that you pass the salt’ and ‘I hereby command you to give me the information.’ But explicit performatives are awkward and overt direction is impolite, so we usually make requests indirectly. Often we do this by asking questions, as with ‘can you pass the salt?’, ‘could you give me the information?’, ‘would you mind not stepping on my foot?’

²¹ This possibility is irrelevant, on Rysiew’s view, because nobody accepts it (2001, p. 499).

²² I’m not convinced that Rysiew is right when he says that this is essential to his view, but set that aside.

²³ The remainder of this section is heavily indebted to John Searle’s (1979, ch. 2) discussion of indirect speech acts.

and ‘do you want to join us for dinner?’ Another way to make an indirect request is by stating that we have certain preferences, as with ‘I would like (you to pass me) the salt,’ ‘I hope you’re able to give me the information,’ and ‘I need you to get off my foot.’

That these are all ways of indirectly requesting is supported by the fact that we can felicitously add ‘please’ to what we say, as in ‘please, can you pass the salt?’ and ‘would you please give me the information?’. It is also supported by the fact that we naturally respond to the questions as if they were requests. We *hear them as* requests. If at the dinner table you ask, “Can you pass the salt?” I pass you the salt without further ado, just as I would if you directly requested it in the imperative mood. Similarly, if you ask, “Can you pass the salt?” and I respond, “No,” you don’t suddenly start worrying that my arms are paralyzed or that I’m lying to you. You understand me to be denying your request, not answering the question itself.

Making a request is your primary purpose in asking, “Can you pass the salt?” That you’re also asking a question is incidental. It’s mutual knowledge that we both already know that I can pass the salt, so you’re neither seeking information nor trying to lead me down a path of self-discovery. Moreover, if I responded directly to your literal question by saying, “Yes, I am indeed able to do that,” or, “Why do you want to know?” it would be interpreted as either humorous (if said while passing the salt) or uncooperative (if said despite not passing the salt).

Call the performance of an indirect speech act *indirection*. We can distinguish different types of indirection. *Conventional indirection* is accomplished by using idioms, which usage has established as indirectional devices. Examples of conventional indirection are ‘how about passing me the salt?’ as a way of requesting the salt, or ‘I’ll be keeping an eye on you’ as a way of warning or putting someone on notice.

Conversational indirection is accomplished by exploiting features specific to the conversational context, along with general communicative principles and background knowledge. Here is an example:

MAN: Let’s go to the movies tonight.

WOMAN: I have a lot to prepare for a major court case scheduled early tomorrow morning.

The man makes a direct proposal. Normally the woman’s response would count as a denial of the proposal. But that’s not because ‘I have a lot to prepare for a major court case scheduled early tomorrow morning’ is conventionally associated with denying proposals. Rather, it’s because, in the context, it’s clear that the best way to make sense of her assertion is that she wants to communicate that she will not be going to the movies tonight with him. To accept the proposal, all she had to say was ‘sure,’ but instead she chose to assert that she had a time-consuming task to complete, which would usually prevent her from having enough time to go to the movies. She wouldn’t have said that unless she was politely declining the proposal.

Not all cases of indirection fall neatly into either conventional or conversational. Many seem to fall somewhere in between. Questions involving ‘can,’ ‘would,’ and

‘could’ are unlike idioms, in that they retain a literal compositional meaning, they admit of direct responses to their literal content, and their literal translation into other languages can preserve their indirectional potential.²⁴ Yet they are also unlike ‘I have to prepare for a major court case’, in that their default status is to be *heard as* requests. Stage-setting is required to hear ‘can you pass the salt?’ primarily as a question about your abilities rather than as a request to pass the salt. By contrast, stage-setting is required to hear ‘I have to prepare for a major court case’ primarily as a denial rather than an assertion.

Closely related to the point about how readily we hear certain formulations as indirect requests, conventional and conversational indirection also differ in whether asking for clarification is felicitous. It would be positively odd for a competent speaker to ask for clarification upon being asked, “How about passing me the salt?” Sincerely asking in turn, “Just to clarify, are you asking me to pass you the salt?” would come across as completely obtuse. By contrast, in cases of conversational indirection, asking for clarification is typically felicitous. In the example above, it would not be out of order for the man to follow up with, “So is that a ‘no’?” or, “We’re not going, then, right?” On this dimension, typical indirect requests featuring ‘can’, ‘would’, and the like behave more like conventional than conversational indirection. Normally if you ask, “Can you change the channel?” and I sincerely ask in turn, “Are you asking me to change the channel?” you’re likely to think me annoyingly dimwitted.

(If we want a label for the in-between cases of indirection involving ‘can’, ‘would’, and the like, let’s call them *convensational indirection*.)

We can deny indirect requests either directly or indirectly. Responding with a simple ‘no’ is blunt and potentially impolite, but still in order. Answering ‘no’ is heard not as a commentary on the speaker’s abilities—i.e. not as answering the literal question—but as a denial of the request. It’s natural to add ‘sorry’, which makes it more polite, though still direct. And if a request is answered with a ‘no’, earnestly replying in turn, “That’s just not true—you’re an able-bodied adult fully capable of passing the salt,” or, “That’s a lie and you know it!” will come across as either incoherent (if the original request was made directly in the imperative), or coherent but obtuse (if the original request was made indirectly in the interrogative).

A very common way of indirectly denying an indirect request is to echo the verb of the literal original question and add a negation. If you ask, “Can you give me the information?” then my response, “I can’t,” is felicitous. If you ask, “Do you want to help me with this?” then my response, “I don’t,” is similarly felicitous. It is also felicitous to echo the entire original question, and add ‘no’, ‘sorry’, and other respectful niceties, as in ‘I am sorry, but no, unfortunately I can’t give the information to you’, and ‘It hurts me to say so, but no, I don’t want to help you with this.’ Taken literally, and setting aside the niceties, these indirect denials are literal assertions about the speaker’s inability to do something

²⁴ Compare to paradigm cases of idioms, such as ‘they tied the knot’, ‘keep an eye out’, and ‘how about the weather lately?’

or about the believer's mental state. But they are heard primarily as denials. Earnestly replying in turn, "That's just not true. You're able to give me the information—it's right there in your hand, and all you have to do is hand it to me," would come across as coherent but obtuse and could be met with, "You're right, I *am* able to. I'm just not going to."

9.6. WAM! BAM! POW!

It's time to relate our discussion of indirection back to the argument against pure invariantism, Anti-PI. The ultimate goal is to defend pure invariantism without courting skepticism. The strategy for achieving this goal is to deny line 2, on the grounds that High Pro speaks properly but falsely when he says 'I don't know'. The proposed tactic for implementing this strategy is to appeal to High Pro's primary illocutionary intention in uttering those words, and thereby drive a wedge between our intuition that he speaks properly, on the one hand, and the further claim that he speaks truthfully, on the other. High Pro's primary illocution is warranted, which is why his speech is proper.

Because I'm appealing to *primary illocutionary warrant*—or what we might call *primary oratorical warrant*—and because I want to emphasize how my proposal both resembles and differs from WAMs, I hereby name my proposal a *POW*.

POWs and WAMs are alike because they both appeal to a basic form of assessment that comes naturally to us when we're interpreting speech, but is clearly distinct from the assessment of the truth-value of the literal propositional content of the speaker's direct speech act, which, it so happens, is an assertion in all the key cases in the literature (whence the title '*warranted assertability maneuver*'). Call any device that fits this mold a *basic assessment maneuver*, or a *BAM*. POWs and WAMs are species of the BAM genus. They distinguish themselves by appealing to different forms of basic assessment. A WAM explains intuitions by appealing to the truth-value of propositions pragmatically conveyed by knowledge ascriptions, rather than the literal truth-value of the knowledge ascriptions themselves. By contrast, a POW explains the intuitions by appealing to the aptness of the indirect speech acts performed by speakers making knowledge ascriptions, rather than the literal truth-value of those ascriptions. Indirect speech acts aren't limited to speech acts with truth-evaluable, assertive propositional content. WAMs generate their plausibility from the fact that pragmatically conveyed propositions attract our attention and figure centrally in our evaluation of speech. POWs generate their plausibility from the fact that primary illocutionary intentions attract our attention and loom large in our evaluation of speech.

The POW I propose focuses on High Pro's primary illocutionary intention to *deny his interlocutor's request*.

Normally the point of asking someone 'Do you know whether *P*?' is *not* to better understand their epistemic standing regarding *P*. Rather, the questioner's primary illocutionary intention is to *get told whether P*. Asking 'Do you know whether *P*?' is a way of indirectly asking '*P*?' (The same applies to other know-wh questions featuring when/how/why/where/what, as well as formulations substituting 'if' for 'whether'.) As

evidence of this, notice that ‘Do you know whether *P*?’ is *heard as* asking whether *P*, we respond to it precisely that way, and we expect people to respond to it that way. If you ask, “Do you know whether the game is tonight?,” a response of ‘I do know’ could be interpreted as either playful or humorous, but is otherwise interpreted as uncooperative. And again, the same is true for other know-wh questions. If you ask, “Do you know when the flight boards?,” a simple ‘Yes I do’ is uncooperative. Stage-setting is required for us to hear ‘Do you know whether *P*?’ primarily as posing the question it literally asks. Its default status is to be heard as the question ‘*P*?’ (The same applies to ‘Can/could/would you tell me whether *P*?’, ‘Might you know whether *P*?’, ‘If you don’t mind, I’d appreciate it if you could tell me whether *P*’, etc.)

To ask someone ‘*P*?’ is to request information. Usually we don’t seek information just for the sake of it. Often we seek it because we’re going to rely on it in planning or evaluating courses of action.

Now let’s return to HIGH FLIGHT. When the courier asks Stewart, “Do you know whether this is a direct flight to Detroit?,” he’s requesting Stewart to provide information that will be relied on in making a serious decision. Making this request is the courier’s primary illocutionary intention. It should be obvious that this is what the courier is doing. To confirm this observation, notice that the courier’s prompt exactly fits the model of indirect requests. And it would have been natural for the courier to frame his final sentence in any of these ways:

- Can you tell me whether *P*?
- Could you tell me whether *P*?
- Would you tell me whether *P*?
- Might you know whether *P*?
- If you don’t mind, I’d appreciate it if you could tell me whether *P*.
- It would be very helpful if you could tell me whether *P*.

Adding ‘please’ would also be perfectly natural and appropriate.

It’s highly unusual for ordinary passengers such as Stewart to be asked to be relied on for information relating explicitly to such decisions. Thus it’s perfectly understandable and appropriate for Stewart to not want to grant the courier’s request. Moreover, it’s not obviously in Stewart’s self-interest to grant the request and undertake responsibility, and he’s not morally required to do so either, because there are people nearby whose job it is to provide that information. So Stewart is warranted in denying the request. And that is exactly what Stewart does, through indirection. By saying ‘I don’t know’, Stewart *indirectly denies* the courier’s indirect request, in the typical way: he echoes the courier’s question and adds a negation. The courier says ‘Do you know . . .?’ and Stewart replies with ‘I don’t know.’²⁵ By contrast, had Stewart replied with ‘Yes, I do—it’s a direct

²⁵ Stewart might also be *encouraging* the courier to find another source of information on the matter. It’s natural, though not required, to interpret Stewart that way.

flight', he would have thereby accepted the courier's request and encouraged the courier to rely on him, thereby undertaking responsibility.

As evidence that Stewart is indirectly denying a request, notice how Stewart's response exactly fits the profile of indirect denial. I already noted the echoic formulation. It would be perfectly natural for Stewart to express himself in this context by saying any of the following:

- No.
- No, sorry.
- No, sorry, I don't.
- No, I'm sorry, I wish I could help, but unfortunately I don't know.
- I'm sorry, I'd like to help, but I'm not the person to ask about that.

It would also be odd for the courier to respond directly to the literal content of Stewart's assertion. If Stewart says, "No, sorry, I don't know," it would not be in order for the courier to respond, "Do you have any evidence for that assertion?" or, "I doubt that that's true." The natural response is more like, "Okay, thanks."

When evaluating Stewart's speech, our immediate response is not to think that he said something false. This is because our evaluation tracks, in the first instance, the fact that his primary illocutionary intention is warranted (i.e. it tracks primary oratorical warrant). Stewart's point is to deny the request, and denials are neither true nor false. That Stewart denies the request by making an assertion is *incidental* and *typically ignored*, just as it is incidental and typically ignored when we indirectly deny a request by saying 'I can't'.

The proposal thus far has granted that Stewart denies the request by literally asserting that he doesn't know. This strikes me as the most plausible account of the situation. But it's worth noting that there is a more radical proposal in the neighborhood, namely, that 'I don't know' functions idiomatically as a way of denying a request for information, which complements the fact that 'Do you know' functions idiomatically as a way of requesting information. On this alternative POW, the *only* speech act Stewart performs is a direct denial, by uttering an idiomatic expression. And whereas I have proposed that Stewart incidentally makes a false assertion that we ignore, the more radical POW denies that there even is an assertion to be ignored.

This brings me to *the essential point* in response to line 2 of the argument. Even if High Pro speaks literally falsely under such circumstances, we should expect this to be ignored and we should expect his speech to both be and seem proper.

The present proposal seems perfectly fitted to explain the asymmetry in verbal behavior in Low versus High cases. As stakes rise, people tend to become more unwilling to be relied upon and undertake responsibility for information crucial to decision-making. And this tendency is not unreasonable, even when holding constant their confidence and the strength of their epistemic position.

This POW avoids problems faced by Rysiew and Brown's WAM. First, it doesn't court methodological danger from the skeptic's corner. It isn't part of my proposal that

we mistakenly judge that a false knowledge ascription²⁶ is true because it conveys true information. Rather, we simply ignore the false ascription because it is incidental and unimportant. Second, and relatedly, this POW doesn't leave us wondering why Stewart chose to say something false. Stewart's utterance follows the typical format of indirect denials. Third, it helps explain why people in HIGH cases respond similarly to both 'Do you know whether *P*?' and '*P*?' They respond similarly because 'Do you know whether *P*?' is simply heard as an indirect request to answer '*P*?' Moreover, this POW does all this without interpreting High Pro's speech as being based on misinformation, and without claiming that High Pro's confidence diminishes.

9.7. Conclusion

J. L. Austin (1956/7, p. 11 n.5) once wrote that when we're investigating why we use certain words in certain situations, and why certain linguistic behavior is appropriate, we should "forget, for once and for a while, that other curious question 'Is it true?'" It's not clear how consistently Austin thinks we ought to avoid asking 'Is it true?', but avoiding it altogether is certainly unwise. After all, truth often matters, even if it's almost never the only thing that matters. If our discussion here is any indication, we can glean some guidance on when it would be wise to look beyond the question of truth, at least for a while, namely, when the speakers themselves aren't primarily concerned with performing a truth-evaluable speech act. When their primary illocutionary intention is, for example, to give an order or deny a request, the explanation of our reaction to their speech shouldn't be expected to track the truth of what they literally say. In such a case, it's a distinct possibility that the truth-value of their literal speech will be irrelevant not only to their own take on the situation, but also to our intuitive assessment of their behavior as well.

References

- Austin, J. L. (1956/7). "A Plea for Excuses," *Proceedings of the Aristotelian Society* 57, pp. 1–30.
- Bach, K. (2010). "Knowledge In and Out of Context," in J. K. Campbell, M. O'Rourke, and H. S. Silverstein (eds), *Knowledge and Skepticism*. Cambridge, MA: MIT Press, pp. 137–63.
- Beebe, J. (2013). "A Knobe Effect for Belief Ascriptions," *Review of Philosophical Psychology* 4, pp. 235–58.
- , and Buckwalter, W. (2010). "The Epistemic Side-Effect Effect," *Mind and Language* 25, pp. 474–98.

²⁶ I say knowledge "ascription" even though it's literally a knowledge denial (i.e. an assertion that knowledge is not present), because I want to avoid confusion between this sort of denial, on the one hand, and denying a request, on the other.

- ., and Jensen, M. (2012). “Surprising Connections Between Knowledge and Action: The Robustness of the Epistemic Side-Effect Effect,” *Philosophical Psychology* 25, pp. 689–715.
- BonJour, L. (2010). “The Myth of Knowledge,” *Philosophical Perspectives* 24, pp. 57–83.
- Brown, J. (2005). “Comparing Contextualism and Invariantism on the Correctness of Contextualist Intuitions,” *Grazer Philosophische Studien* 65, pp. 136–64.
- . (2006). “Contextualism and Warranted Assertibility Manoeuvres,” *Philosophical Studies* 130, pp. 407–35.
- Buckwalter, W., Rose, D., and Turri, J. (2013). “Belief through Thick and Thin,” *Noûs*. DOI: 10.1111/nous.12048.
- Cohen, S. (1999). “Contextualism, Skepticism, and the Structure of Reasons,” *Philosophical Perspectives* 13, pp. 57–89.
- . (2005). “Knowledge, Speaker and Subject,” *Philosophical Quarterly* 55, pp. 199–212.
- Conee, E. (2013). “Contextualism Contested,” in M. Steup, J. Turri, and E. Sosa (eds), *Contemporary Debates in Epistemology*, 2nd ed. Malden, MA: Wiley-Blackwell, pp. 60–9.
- Davis, W. (2007). “Knowledge Claims and Context: Loose Use,” *Philosophical Studies* 132, pp. 395–438.
- DeRose, K. (1999). “Contextualism: An Explanation and Defense,” in J. Greco and E. Sosa (eds), *The Blackwell Guide to Epistemology*. Malden, MA: Blackwell.
- . (2009). *The Case for Contextualism: Knowledge, Skepticism, and Context, Volume 1*. Oxford: Oxford University Press.
- Fantl, J., and McGrath, M. (2009). *Knowledge in an Uncertain World*. Oxford: Oxford University Press.
- Feltz, A., and Zarpentine, C. (2010). “Do You Know More When It Matters Less?” *Philosophical Psychology* 23, pp. 683–706.
- Greco, J. (2010). *Achieving Knowledge: A Virtue-Theoretic Account of Epistemic Normativity*. Cambridge: Cambridge University Press.
- Grimm, S. (2011). “On Intellectualism in Epistemology,” *Mind* 120, pp. 705–33.
- Hawthorne, J. (2004). *Knowledge and Lotteries*. Oxford: Oxford University Press.
- Lewis, D. (1996). “Elusive Knowledge,” *Australasian Journal of Philosophy* 74, pp. 549–67.
- MacFarlane, J. (2005). “The Assessment Sensitivity of Knowledge Attributions,” in T. Gendler and J. Hawthorne (eds), *Oxford Studies in Epistemology*. Oxford: Oxford University Press, pp. 190–205.
- McGrath, M. (2010). “Contextualism and Intellectualism,” *Philosophical Perspectives* 24, pp. 383–405.
- Moore, G. E. (1959). *Philosophical Papers*. New York: Collier Books.
- Murray, D., Sytsma, J., and Livengood, J. (2013). “God Knows (But Does God Believe?),” *Philosophical Studies* 166, pp. 83–107.
- Myers-Schulz, B., and Schwitzgebel, E. (2013). “Knowing That P Without Believing That P,” *Noûs* 47, pp. 371–84.
- Pinillos, Á. (2012). “Knowledge, Experiments and Practical Interests,” in J. Brown and M. Gerken (eds), *New Essays on Knowledge Ascriptions*. Oxford: Oxford University Press.
- Reed, B. (2007). “The Long Road to Skepticism,” *Journal of Philosophy* 104, pp. 236–62.
- . (2010). “A Defense of Stable Invariantism,” *Noûs* 44, pp. 224–44.
- Rose, D., and Schaffer, J. (2013). “Knowledge Entails Dispositional Belief,” *Philosophical Studies* 166, S19–S50.

- Rysiew, P. (2001). "The Context-Sensitivity of Knowledge Attributions," *Noûs* 35, 4, pp. 477–514.
- . (2007). "Speaking of Knowledge," *Noûs* 41, pp. 627–62.
- Schaffer, J. (2006). "The Irrelevance of the Subject," *Philosophical Studies* 127, pp. 87–107.
- ., and Knobe, J. (2012). "Contrastive Knowledge Surveyed," *Noûs* 46, pp. 675–708
- Searle, J. (1979). *Expression and Meaning*. Cambridge: Cambridge University Press.
- Sripada, C., and Stanley, J. (2012). "Empirical Tests of Interest-Relative Invariantism," *Episteme* 9, pp. 3–26.
- Stanley, J. (2005). *Knowledge and Practical Interests*. Oxford: Oxford University Press.
- Starmans, C., and Friedman, O. (2012). "The Folk Conception of Knowledge," *Cognition* 124, pp. 272–83.
- Stine, G. (1976). "Skepticism, Relevant Alternatives, and Deductive Closure," *Philosophical Studies* 29, pp. 249–61. Reprinted (1st ed. 2000; 2nd ed. 2008), in E. Sosa and K. Jaegwon (eds), *Epistemology: An Anthology*. Malden, MA: Blackwell.
- Swain, S., Alexander, J., and Weinberg, J. (2008). "The Instability of Philosophical Intuitions: Running Hot and Cold on Truetemp," *Philosophy and Phenomenological Research* 76, pp. 138–55.
- Turri, J. (2010a). "Epistemic Invariantism and Speech Act Contextualism," *Philosophical Review* 119, pp. 77–95.
- . (2010b). "Prompting Challenges," *Analysis* 70, pp. 456–62.
- . (2011). "The Express Knowledge Account of Assertion," *Australasian Journal of Philosophy* 89, pp. 37–45.
- . (2012). "Is Knowledge Justified True Belief?," *Synthese* 184, pp. 247–59.
- . (2013). "A Conspicuous Art: Putting Gettier to the Test," *Philosophers' Imprint* 13, pp. 1–16. <<http://hdl.handle.net/2027/spo.3521354.0013.010>> (last accessed November 20, 2013).
- . (Under review). "The Problem of ESEE Knowledge."
- ., and Friedman, O. (Forthcoming). "The Folk Epistemology of Lotteries," in J. Beebe (ed.), *Advances in Experimental Epistemology*. Continuum.
- Unger, P. (1975). *Ignorance: A Case for Skepticism*. Oxford: Oxford University Press.
- Weinberg, J., Nichols, S., and Stich, S. (2001). "Normativity and Epistemic Intuitions," *Philosophical Topics* 29, pp. 429–60.
- Williamson, T. (2005). "Contextualism, Subject-Sensitive Invariantism and Knowledge of Knowledge," *Philosophical Quarterly* 55, pp. 212–35.