

The Problem of False Denials: Invariantism and Error

Geoff Pynn

Abstract

Invariantists often attempt to offer pragmatic explanations of the linguistic data that motivates contextualism. I consider some such attempts and conclude that they cannot escape the claim that speakers systemically err in their judgments about who knows what. Like it or not, moderate invariantists must impute widespread error to ordinary speakers.

1 Contextualist cases

The primary evidence for contextualism is linguistic. Contextualists present pairs of cases where our intuitions about the propriety of a particular knowledge ascription and the corresponding knowledge denial shift. In one case—the ‘easy context’ case—a speaker seems appropriately to ascribe knowledge of a proposition to a subject whose epistemic position with respect to that proposition is only moderately good. In the other case—the ‘hard context’ case—a speaker seems appropri-

ately to deny knowledge of the same proposition to the same subject, despite the fact that her epistemic position is just as it was in the first case.¹

The contextualist explanation for this phenomenon is that the truth-conditions of knowledge-ascribing sentences can shift with features of the context in which they are uttered. While the sentences used in each context appear on the surface to contradict each other, they do not. Thus the knowledge ascription in the easy context and the knowledge denial in the hard context are both true (or, at least, can be). This is why they are appropriate.

Invariantists deny that the truth-conditions of knowledge-ascribing sentences can shift in this way. If a knowledge-ascribing sentence is true when uttered in a one context, then the corresponding knowledge-denying sentence uttered in another context is false. However, most invariantists do not wish to deny that our intuitions about the propriety of each sentence shift with context as the contextualist maintains. To blunt the force of this argument for contextualism, then, invariantists must offer some alternative explanation for these intuitions.²

2 Invariantism and seemingly appropriate falsehoods

Invariantism comes in several flavors. *Skeptical* invariantists think that knowledge demands that subjects stand in a *very* strong epistemic position, and thus that it is seldom if ever true to say that we know things. *Moderate* invariantists think knowledge is less demanding, and hence that it is often true to say that we

know things. Orthogonal to this distinction is the distinction between *sensitive* and *insensitive* invariantism. Insensitive invariantists think that the requirements for knowledge are more-or-less the same for every subject in every situation. Sensitive invariantists think that the requirements for knowledge can shift with features of a subject's situation, and hence that a subject in a single epistemic position with respect to p may know p in one situation but not in another.³ Most sensitive invariantists should also be counted as moderates, since they hold that many situations determine relatively relaxed requirements for knowledge, and hence that many of our ordinary knowledge ascriptions are true.

All invariantists must acknowledge that we find it appropriate for speakers in some range of contexts to make assertions which, according to invariantism, are false. Moderate invariantists must say this about our hard context knowledge denials and skeptical invariantists about our easy context knowledge ascriptions. Thus moderate invariantists face:

The Problem of False Denials Why do we find it appropriate for speakers to deny knowledge to subjects in hard contexts, even though such denials are typically false?

Skeptical invariantists face an analogous problem:

The Problem of False Ascriptions Why do we find it appropriate for speakers to ascribe knowledge to subjects in easy contexts, even though such ascriptions are typically false?

Because most invariantists are moderates, I focus here on the problem of false

denials. Appropriately modified, most of what I say here could be said about the problem of false ascriptions, too.

3 Two explanatory strategies

One strategy for addressing the problem is *pragmatic*. It says that hard context knowledge denials are appropriate because they pragmatically convey something true, even if what those denials literally say is false. Another strategy is *error theoretic*. It says that in hard contexts speakers erroneously judge that subjects lack knowledge; this error leads them to regard such denials as appropriate. A compelling pragmatic explanation must say what truth is conveyed by a false hard context knowledge denial and how such denials manage to convey this truth. A compelling error theoretic explanation must say what it is about hard contexts that leads us to mistakenly judge that speakers don't know things.

The boundary between the strategies is blurred. Pragmatic explanations of the envisaged sort might posit some error, for example widespread confusion between what an assertion literally says and what it pragmatically conveys. And an error theoretic explanation might have the consequence that, since speakers reliably make the posited error, false denials reliably convey certain truths. But we can highlight the distinction with three differences.

First, the error-theoretic approach traces our intuition that hard context denials are appropriate to a mistaken judgment that they are true. The pragmatic strategy traces our intuition that they are appropriate to the fact that they are in-

formative.

Second, the error-theoretic approach explains why false denials *seem* appropriate, but might stop short of saying that they *are* appropriate; after all, they result from a mistaken judgment. On the assumption that assertions based on mistaken judgments are typically inappropriate, an error-theoretic explanation might well rule that false denials really *aren't* appropriate, even if they seem to be. The pragmatic strategy, by contrast, says that the informativeness of false denials in fact makes them appropriate, and not just seemingly so.

Third, the error-theoretic approach pairs well with a revisionary agenda. Once we learn that our false hard context knowledge denials result from systematic error, we should stop making them—or so one might plausibly argue. The pragmatic approach promises an account of our false denials that not only explains why we make them, but justifies our practice of doing so.

Thus the pragmatic approach is more charitable towards our ordinary judgments and respectful of our ordinary ways of talking. These are *prima facie* reasons to favor it. But can pragmatics do the job? In what follows I shall argue that the pragmatic approach faces serious challenges, and that the most straightforward way to meet those challenges implies that speakers are prone to make a particular sort of error—an error which would itself imply that their false denials are the result of mistaken judgments about what subjects know. Invariantists are stuck with an error theory.⁴

4 Some pragmatic explanations

Let's distinguish between the epistemic standard a subject must meet in order to know something—the standard_K—and the epistemic standard that is salient in a conversation, which I'll call the standard_C. For invariantists, the standard_K is fixed across contexts. The standard_C, however, can vary from context to context.⁵ For moderate invariantists, the standard_K is not particularly demanding; i.e., knowing *p* always requires being in a merely good epistemic position with respect to *p*, and never an extraordinarily strong one. But the standard_C can become quite demanding. Whatever factors the contextualist maintains drive up the standard_K, the invariantist can maintain will drive up the standard_C.

For example, the invariant standard_K might be loose enough that someone can always know that her car is parked in the garage without having conclusive evidence that her car has not just been stolen from the garage. In certain contexts, though, we might become focused on the possibility that the car *has* been stolen from the garage. In these contexts, the standard_C might become so high that one meets it with respect to the proposition that one's car is parked in the garage only if one *does* have conclusive evidence that it has not just been stolen. Thus one could know that one's car is parked in the garage without meeting the epistemic standard_C with respect to that proposition.

With this distinction in mind, we can begin to examine some pragmatic strategies for solving the problem of false denials.

4.1 Relevance

Patrick Rysiew and Jessica Brown each argues that in contexts where (a) the standard_C is significantly inflated and (b) the conversation is focused on whether *S* meets the standard_C with respect to *p*, asserting that *S* knows that *p* pragmatically conveys:

(1) *S* meets the epistemic standard_C with respect to *p*.⁶

Note that in such a context, asserting that *S* knows that *p* would not *entail* (1). Given that the standard_C is inflated, one can meet the standard_K and hence know *p* even if (1) is false. So why does ascribing knowledge to *S* nonetheless convey (1)?

Rysiew and Brown claim that when the conversation is focused on whether (1) is true, an assertion that *S* knows that *p* would be irrelevant. Ascribing knowledge of *p* to *S* in such a context would thus violate Grice's Maxim of Relation—"be relevant!"⁷ Her listeners, assuming that she is heeding the Maxim of Relation and hence intends for her assertion to be relevant, infers that she intends to communicate (1). Thus, her knowledge ascription conveys (1).⁸

Rysiew and Brown each goes on to suggest that it follows from this that by *denying* that the subject knows *p*, the speaker thereby conveys the denial of (1); i.e., (2):

(2) *S* does not meet the epistemic standard_C with respect to *p*.⁹

The Rysiew-Brown account of false denials thus proceeds in two steps: first, the presumption of relevance leads a hard context knowledge ascription to convey (1), and second, this implies that a hard context knowledge denial conveys (2). If in the envisaged context (2) is true—as it is liable to be, given that the standard_C is unusually inflated—then a false denial is informative, and hence appropriate.

But each step of the explanation is flawed. First, when a conversation is focused on whether S meets an inflated standard $_C$, the literal content of a knowledge ascription is relevant to the conversation. Consider a parallel case. Suppose that we are discussing whether Brian is a billionaire. I say, “Brian’s net worth is at least \$500 million.” Though my assertion does not settle the question at issue, it is highly relevant to that question. Similarly, if knowledge requires meeting the standard $_K$, and I assert that S knows p , then I have asserted something highly relevant to whether S meets the inflated standard $_C$, even if my assertion does not settle that question.

Moreover, a speaker who makes a relevant assertion that doesn’t settle the question at issue thereby indicates that she does *not* intend to communicate anything that *does* settle that question. A listener who reasoned that my assertion that Brian’s net worth is at least \$500 million would be irrelevant if I did not intend to communicate that Brian was a billionaire would be badly misreading me. Similarly, if the question is whether S meets an inflated standard $_C$ with respect to p , and I assert that S knows that p , I’ve asserted something relevant to but incapable of settling the question. I am thereby indicating that I am do *not* intend to communicate something that *would* settle the question. My listener would not be licensed in concluding that I also meant to communicate (1).

Second, even if the presumption of relevance *did* lead knowledge ascriptions to convey (1), that would be no reason to think that knowledge denials convey (2). Consider Grice’s famous example of a relevance implicature: A says, “Mrs. X is an old bag.”¹⁰ In response, B says, “The weather has been quite delightful this sum-

mer, hasn't it?" B's response is relevant only if he intends to communicate that A's remark should not be discussed; his response thus conveys that A's remark should not be discussed. But if B had instead *denied* that the weather had been quite delightful, he would not thereby have conveyed that A's remark *should* be discussed. More generally, even if a knowledge ascription in a hard context pragmatically conveys (1)—as a relevance implicature, or by some other means—this would not imply that a knowledge denial conveys (2). The fact that asserting p conveys but does not entail q does not imply that denying p conveys the denial of q .¹¹

In a subsequent paper, Rysiew offers a more direct relevance-based explanation for why hard context knowledge denials convey (2)—one that does not detour through the (alleged) fact that knowledge ascriptions convey (1). There he says that when a speaker denies knowledge in a context focused on whether (1) is true, and his listeners have no reason to think he regards p as false or S as not in a good epistemic position relative to p , then the assumption that he intends to be relevant should lead them to infer that he means to communicate (2).¹²

But this explanation is equally unsatisfying. As Keith DeRose points out, even when the *issue* of whether (1) is true is more relevant than the *issue* of whether S knows p , a speaker who denies that S knows p thereby says something highly relevant to the conversation.¹³ Provided the listener takes the speaker to think that S truly believes p , a knowledge denial will indicate that the speaker thinks that S does not meet the standard _{K} . If S does not even meet the standard _{K} with respect to p , then surely S does not meet the standard _{C} with respect to p . So a knowledge denial, taken literally, would settle the relevant issue. Thus it's unclear

why a listener would be led to initiate a search for a more relevant proposition than the one literally asserted.¹⁴

However, this more direct strategy suggests another pragmatic explanation.

4.2 Exaggeration

If S does not satisfy the standard $_K$, then S does not satisfy an even more demanding standard. Suppose we are in a context where it is taken for granted that, if S does not know p , this is because S does not meet the standard $_K$ with respect to p , and also where the standard $_C$ exceeds the standard $_K$. Given this background, a knowledge denial will be taken to *entail* (2). In such a context, a knowledge denial will communicate (2).

Speakers who deliberately assert a falsehood that is stronger than p in order to communicate p are exaggerating. On the present proposal, when demanding standards become salient, speakers exaggerate by saying that subjects do not know things. Those exaggerations accurately communicate that the subjects in question do not meet the salient standards with respect to the propositions they are (falsely) said not to know.

One advantage this proposal has over the previous one is that it makes it clear how denying knowledge conveys (2). Still, hard context knowledge denials don't *feel* like exaggerations. Exaggeration is a deliberately non-literal use of language, and most sophisticated speakers are adept at discerning such uses. Hard context knowledge denials have none of the dramatic character of ordinary exaggeration; e.g., "It took me a million years to make it through Book 1 of the *Logical Investi-*

gations!"¹⁵ Here are two ways this difference is manifest:

First, we are typically willing to retract exaggerations under mild pressure. "No, of course he's not about to die of starvation. (He's just very hungry.)" But we are not similarly willing to retract hard context knowledge denials. "Yes, of course he knows that his car's parked in the lot. (He just doesn't have any evidence that it hasn't been stolen.)"

Second, we can employ devices like 'in fact' and 'literally' to indicate that we are not exaggerating.¹⁶ "He is, in fact / literally, about to die of starvation. Call 911." If hard context knowledge denials are exaggerations, we should hear a significant difference between "He doesn't know that his car's in the lot" on the one hand and "He doesn't, in fact, know that his car's in the lot" and "He literally doesn't know that his car's in the lot" on the other. But we do not; the latter sound no more literal than the first. So it is implausible to regard hard context denials as exaggerations.¹⁷

4.3 A false presupposition to the rescue

Let's suppose that the pragmatic theorist is right that hard context ascriptions communicate (1), and that hard context denials communicate (2). Neither the relevance nor the exaggeration accounts can explain why this is so. How else might we do so?

Imputing to speakers and listeners in hard contexts a false presupposition will do the trick. Proposal: when the standard_C is unusually inflated, speakers and listeners have a tendency to mistake it for the standard_K . Thus in hard contexts they falsely believe (3):

(3) S knows p only if S meets the standard _{C} with respect to p .

If (3) is not only believed, but *presupposed*—i.e., if all conversational participants believe (3) and attribute this belief to each other—then we can explain why knowledge ascriptions convey (1), and knowledge denials convey (2).

First, if in a hard context a speaker asserts that S knows that p , then her interlocutors, presupposing (3), will be in a position to conclude from what the speaker said that she believes that S meets the standard _{C} with respect to p . Thus, knowledge ascriptions in hard contexts communicate (1).

Second, if in a hard context a speaker denies that S knows p , and her interlocutors presuppose (3), they will be in a position to conclude that, provided the speaker does not regard S as failing to meet some other requirement for knowledge—e.g., that S fails to believe p , or that his belief is somehow “gettierized”—she believes that S does not meet the standard _{C} with respect to p . Thus her knowledge denial will communicate (2).

If we treat a hard context as one with an inflated standard _{C} and a false presupposition of (3), then we have a clean explanation for the pragmatic effects of knowledge ascriptions and denials in hard contexts, and a resolution to the problem of false denials. But the explanation comes at a cost.

5 An error theory

Obviously, if participants in hard context conversations falsely presuppose (3), then they have made an error. If, to adequately explain the pragmatic effects of

knowledge denials in those contexts, a pragmatic theorist must impute this systematic error to speakers, then she is no longer offering a purely pragmatic theory. Her account of the informativeness of false denials depends upon the tendency of speakers to make a certain mistake.

However, if (3) is presupposed in hard contexts, a far simpler account of false denials is available. In a hard context, the speaker correctly discerns that the subject fails to meet the standard_C with respect to p , and then, via her false presupposition that (3), mistakenly judges that the subject does not know p . This is what inclines her to make her false denial. Her listeners, also presupposing (3), go along. No circuitous route through purported pragmatic effects is needed to make sense of the propriety of her denial; the false presupposition of (3) is enough.

Why is it that when an inflated epistemic standard becomes salient, speakers tend to go wrong about what is required for knowledge? Obviously, this is a crucial question, and without an answer, the suggested error theory is incomplete. Here is a very brief sketch of an explanation. The notion of an epistemic standard is a technical one. Ordinary speakers have only a fuzzy grasp on this aspect of knowledge, and are not good at keeping track of multiple epistemic standards in a single context. When a highly demanding epistemic standard becomes salient in a conversation, speakers do not distinguish it from the epistemic standard for knowledge. This failure to distinguish the inflated salient standard from the standard required for knowledge leads them astray, and they mistakenly regard subjects who do not meet the salient standard as not being in a position to know.

This brief sketch is hardly sufficient, but I hope it can be developed into a more

precise and plausible story. And all moderate invariantists should hope so as well. For as I have argued, the problem of false denials can't be solved by pragmatics alone. Invariantists are stuck with some form of error theory—if not an impure pragmatic account, then a straight-up error theory like the one I've just described, or of some other sort.¹⁸

Notes

¹The best-known contextualist cases are the Bank Cases of DeRose 1992 and Cohen 1999's Airport Case. See also Stanley 2005's versions of the Bank Cases, and DeRose 2009, 4-5 for his new Thelma and Louise Cases.

²There are two ways to characterize the intuitions mustered by contextualists. The first is as I described them above: as intuitions that certain assertions are natural and appropriate. Contextualists usually go further than this, though. Keith DeRose, for example, describes contextualist cases as follows:

A 'low-standards' case [...] in which a speaker seems quite appropriately *and truthfully* to ascribe knowledge to a subject will be paired with a 'high-standards' case [...] in which another speaker in a quite different and more demanding context seems with equal propriety *and truth* to say that the same subject (or a similarly positioned subject) does not know. (DeRose 2009, 47, emphasis added.)

An intuition of the propriety of an assertion is not *ipso facto* an intuition that the content of what is asserted is true. Hans and Julie are married, but not to each other. Brian is talking with Tracie, who thinks that Hans is Julie's husband, and Brian corrects her:

(*) Hans and Julie aren't married.

Brian's assertion seems perfectly appropriate. Does it also seem that (*) is literally true? Not to my ears, though others may have a different reaction. I raise the example just to point out that

intuitions of propriety and intuitions of truth needn't sway together. Here I will not be concerned with the claim that our contextualist cases yield intuitions about the truth of the relevant knowledge ascriptions and denials. Whether or not they seem true, they certainly seem appropriate, and it is this intuition that I shall be concerned with.

Informed readers will know that in the last several years another view has emerged to rival both contextualist and invariantist accounts of these intuitions: relativism. See MacFarlane 2005 for a primer. Here I'm concerned with one cost of invariantism, and so far as I can tell this cost is not born by relativism.

³See Hawthorne 2004 and Stanley 2005 for versions of sensitive invariantism. Sensitive invariantists and contextualists tend to agree about the truth-values of our knowledge ascriptions and denials when the subject is the speaker or some other participant in the conversation; that is because the features of the context that, for the contextualist, fix the semantic content of "*S* knows that *p*" are, for the sensitive invariantist, the features of the subject's situation that determine how good her epistemic position must be for her to know *p*. But the verdicts of the sensitive invariantist and the contextualist diverge when it comes to ascriptions and denials concerning subjects in distant situations. Thus I'll restrict my discussion to third-person knowledge ascriptions and denials.

⁴Much of the discussion of what I am calling the "pragmatic strategy" has concerned what ? labels a 'Warranted Assertibility Maneuver', or WAM. I refrain from using DeRose's term for two reasons. First, DeRose claims that no WAM can be successfully employed to explain away the appearance of truth. If that is so, then if the apparent propriety of false denials leads us to mistakenly regard them as true, no WAM can successfully solve the problem of false denials. Second, it is clear from his later work (esp. ?) that DeRose intends 'warranted assertibility' to mean something like '*epistemically* warranted assertibility'. But a successful pragmatic account needn't show how false denials are *epistemically* appropriate; I take it that our reactions to contextualist cases are not so precise as that.

⁵For the sensitive invariantist, the standard_{*K*} is a function of the subject's situation, and hence is not the same for every subject in every situation. But once the relevant features of a subject's situation are fixed, then the standard_{*K*} is also fixed, and can come apart from the standard_{*C*} in a

particular conversation.

⁶Rysiew 2001, Rysiew 2007, and Brown 2006. Brown, unlike Rysiew, holds that it is not just the fact that the conversation is focused on whether *S* meets the high standard_{*C*} with respect to *p*, but also the practical importance to the participants of being wrong about *p* that triggers the pragmatic effect of ascribing knowledge to *S*. Also, it should be noted that Rysiew couches his view in terms of a relevant alternatives theory of knowledge. He distinguishes between those not-*p* alternatives that must be ruled out in order to know *p*, and those not-*p* alternatives that are salient in the conversational context; in some contexts the salient alternatives may not be among those that *S* must rule out in order to know *p*. On his view, then, a knowledge ascription conveys not (1) but (1R):

(1R) *S* is in a position to rule out all of the salient alternatives to *p*.

As Rysiew explains, though, the relevant alternatives framework is not essential to his general account, which can be described just as adequately in terms of standards, and not alternatives.

⁷See Grice 1989, 27. It's not at all clear what "relevance" amounts to here, as Grice himself was happy to admit: "Though this maxim is terse, its formulation conceals a number of problems that exercise me a good deal: questions about what different kinds and focuses of relevance there may be, how these shift in the context of a talk exchange, how to allow for the fact that subjects of conversation are legitimately changed, and so on. I find the treatment of such questions exceedingly difficult, and I hope to revert to them in later work."

⁸Rysiew 2001, 491f., Brown 2006, 426. Rysiew 2007, 638-639 offers a slightly different and more sophisticated version of this claim, but the criticism I shall offer below applies to it, as well.

⁹I say that they "suggest" this because each, immediately after arguing that relevance concerns explain why hard context knowledge ascriptions convey (1), simply states, without further explanation, that hard context knowledge denials convey (2). In a moment, we'll look at an apparently distinct account of (2)'s conveyance from Rysiew.

¹⁰Grice 1989, 35.

¹¹Rysiew 2007, 660, fn. 31 responds to a similar criticism raised by John MacFarlane (Mac-

Farlane 2005, 208). There Rysiew says that it is essential to his view that speakers are “generally insensitive to the semantic/pragmatic distinction,” and thus presumably that they are prone to regard the negation of a sentence that pragmatically conveys some proposition as also implying the negation of the pragmatically conveyed proposition. But if this is so, then it is entirely unclear how a knowledge ascription could convey (1) via the presumption of relevance in the first place. If a listener is insensitive to the distinction between the semantic properties and pragmatic effects of a knowledge ascription, what reason would she have for regarding her interlocutor’s assertion as irrelevant unless accompanied by the intention to communicate (1)?

¹²Rysiew 2005, 65.

¹³DeRose 2009, 122f.

¹⁴Consider again a parallel case: a speaker who asserted, “Brian’s net worth is no more than \$500 million” in response to the question of whether Brian was a billionaire literally says something highly relevant to that question, even if the issue of whether Brian’s net worth is \$500 million is less relevant to the conversation than the issue of whether it exceeds \$1 billion.

¹⁵Of course we can also exaggerate secretly, as when we claim to have read all of Husserl’s *Logical Investigations* when all we’ve really read is part of Book 1. Secret exaggeration does not wear its non-literality on its sleeve. But this sort of exaggeration—which is a kind of lying—is clearly not what is going on with hard context knowledge denials.

¹⁶This point is due to Hawthorne 2004, 120.

¹⁷There may be cases of “exaggeration” that do not wear their non-literality on their sleeve; e.g., some of the examples of what Bach 1994 calls “implicature”. For example, someone who says “I haven’t eaten breakfast” is typically interpreted to mean “I haven’t eaten breakfast today.” But one might think (as Bach appears to) that her utterance is literally true only if she has *never* eaten breakfast. If so, then she has said something literally false in order to communicate something weaker but true. Still, it would not be obvious to most speakers that she has spoken non-literally.

However, in this case (as in all of the cases Bach gives of such implicatures), to the extent that it is plausible to regard the utterance as literally false, it is also both natural to retract it (“Well, of course I *have* eaten breakfast, just not today”), and to employ devices signalling that it is intended

literally (“I haven’t eaten breakfast—literally, never!”). So it passes the two tests for exaggeration that (1) fails. Moreover, in all of Bach’s examples, the literally false utterance is syntactically very closely related to what is communicated; typically, the communicated proposition can be expressed with the addition of just a few words (e.g., “You’re not going to die” → “You’re not going to die *from that cut*”; “France is hexagonal” → “France is *roughly* hexagonal”). There is no such syntactic overlap between (1) and (2).

¹⁸For other, very different error-theoretic accounts of hard context false denials, see ? and ?.

References

- Bach, K. (1994). Conversational implicature. *Mind and Language*, 9:124–162.
- Brown, J. (2006). Contextualism and warranted assertibility manoeuvres. *Philosophical Studies*, 130:407–435.
- Cohen, S. (1999). Contextualism, skepticism, and the structure of reasons. *Philosophical Perspectives*, 13:57–89.
- DeRose, K. (1992). Contextualism and knowledge attributions. *Philosophy and Phenomenological Research*, 52:913–929.
- DeRose, K. (2009). *The Case for Contextualism*. Oxford University Press, Oxford.
- Grice, P. (1989). *Studies in the Way of Words*. Harvard University Press, Cambridge, MA.
- Hawthorne, J. (2004). *Knowledge and Lotteries*. Oxford University Press.
- MacFarlane, J. (2005). The assessment sensitivity of knowledge attributions. *Oxford Studies in Epistemology*, 1:197–234.
- Rysiew, P. (2001). The context-sensitivity of knowledge attributions. *Noûs*, 35:477–514.
- Rysiew, P. (2005). Contesting contextualism. *Grazer Philosophische Studien*, 69:51–69.
- Rysiew, P. (2007). Speaking of knowing. *Noûs*, 41(4):627–662.

Stanley, J. (2005). *Knowledge and Practical Interests*. Oxford University Press.