

Minimalism and the Limits of Warranted Assertability Maneuvers^{*}

BLAKE ROEBER

University of Notre Dame

Abstract: Contextualists and pragmatists agree that knowledge-denying sentences are contextually variable, in the sense that a knowledge-denying sentence might semantically express a false proposition in one context and a true proposition in another context, without any change in the properties traditionally viewed as necessary for knowledge. Minimalists deny both pragmatism and contextualism, and maintain that knowledge-denying sentences are not contextually variable. To defend their view from cases like DeRose and Stanley’s high stakes bank case, minimalists like Patrick Rysiew, Jessica Brown, and Wayne Davis forward “warranted assertability maneuvers.” The basic idea is that some knowledge-denying sentence seems contextually variable because we mistake what a speaker pragmatically conveys by uttering that sentence for what she literally says by uttering that sentence. In this paper, I raise problems for the warranted assertability maneuvers of Rysiew, Brown, and Davis, and then present a warranted assertability maneuver that should succeed if any warranted assertability maneuver will succeed. I then show how my warranted assertability maneuver fails, and how the problem with my warranted assertability maneuver generalizes to pragmatic responses in general. The upshot of my argument is that, in order to defend their view from cases like DeRose and Stanley’s high stakes bank case, minimalists must prioritize the epistemological question whether the subjects in those cases know over linguistic questions about the pragmatics of various knowledge-denying sentences.

Suppose that you go from a situation where it doesn’t matter much whether p is true to a situation where it matters a lot whether p is true. Does this change in your practical interests change the strength of epistemic position required for you to know that p , or the strength of epistemic position required for you to semantically express a true proposition by uttering the sentence ‘I know that p ’?

The pragmatic encroachment debate pits pragmatists against contextualists, and purist invariantists against both of them. By ‘pragmatism,’ I mean the view that knowledge depends at least in part on our practical interests, in the sense that a mere difference in practical interests can entail a difference in knowledge.¹ If pragmatism is true, there will be pairs of cases that differ only insofar as some stipulated difference in practical interests requires that they differ, and where, as a result of just this difference in practical

^{*} Penultimate draft. Please cite final draft in *Episteme*.

¹ See, for example, Fantl and McGrath (2002), Hawthorne and Stanley (2008), Ganson (2008), Fantl and McGrath (2009), Weatherson (2012), Ross and Schroeder (2014), and the portions of Hawthorne (2004) and Stanley (2005) that are explicitly about knowledge and practical interests.

interests, a difference in knowledge follows. Contextualism, in contrast, is a theory about the word 'knows.' According to contextualism, the word 'knows' can semantically express different properties or relations in different contexts of utterance, and a sentence containing the word 'knows' can therefore semantically express different propositions in different contexts of utterance.² Contextualism entails that, if you and I occupy different contexts when we utter some sentence of the form '*S* knows that *p*,' you might semantically express a true proposition while I semantically express a false proposition, even if we utter this sentence at exactly the same time.³

What motivates contextualism and pragmatism, respectively? Contextualism evolved as both a solution to the "skeptical problem" (DeRose 1995) and an explanation of our intuitions about the knowledge ascriptions and denials in cases like the following.

Low Stakes: Rachel and Keith are driving home on Friday afternoon. They plan to stop at the bank to deposit their paychecks, but, as they approach it, they notice that the lines inside are very long. Although they generally like to deposit their paychecks as soon as possible, it's not especially important in this case that they deposit their checks right away, so Keith suggests that they drive straight home and deposit their checks tomorrow morning. Rachel says, "Maybe the bank won't be open tomorrow. Lots of banks are closed on Saturdays." Keith replies, "No, I know it will be open. I stop in at the bank every Saturday. It's open until noon."

High Stakes: Rachel and Keith are approaching their bank on Friday afternoon, as in Low Stakes, and they notice the long lines. Keith again suggests that they go straight home and return to deposit their paychecks tomorrow morning, explaining that he stops at the bank every Saturday morning and that it's open until noon. But in this case, Rachel and Keith have just written a very large and important check. If their paychecks are not deposited before Monday morning, the important check they wrote will bounce, leaving them in a terrible situation. The bank isn't open on Sunday. Rachel reminds Keith of these facts. She then says, "What if our bank discontinued its Saturday hours since your last visit? Do

² See, for example, Cohen (1988), DeRose (1995), and Lewis (1996).

³ The label 'pragmatism' comes from Fantl and McGrath (2009). Though perhaps misleading, it does allow us to easily distinguish between the pragmatist thesis that knowledge depends at least in part on our practical interests and the *invariantist* thesis that the word 'knows' *does not* express different properties or relations in different contexts of utterance. (The more popular labels 'subject-sensitive invariantism' and 'interest-relative invariantism' both lack this desirable feature.)

you know that it will be open tomorrow?” Remaining just as confident as he was before that the bank will be open, Keith replies, “Well, no, I don’t know that the bank will be open. We had better stop and deposit the checks now.”⁴

While contextualism evolved as an explanation of our intuitions about the knowledge attributions and denials in cases like these, pragmatism evolved as a solution to the lottery paradox and as a *competing* explanation of our intuitions about cases like Low and High Stakes.⁵ So, while pragmatism and contextualism are logically consistent—pragmatists might be right about *knowledge* while contextualists are right about the word ‘knows’—pragmatism and contextualism are standardly viewed as rivals.⁶

Neither contextualism nor pragmatism is universally accepted among epistemologists, however. Where ‘purism’ names the denial of pragmatism, and ‘invariantism’ names the denial of contextualism, many epistemologists are purist invariantists. Purist invariantism entails (a) that Keith knows that the bank will be open in both Low Stakes *and* High Stakes if he knows that the bank will be open in either case, and (b) that Keith would semantically express a true proposition by uttering the sentence ‘I know that the bank will be open’ in both cases if he would semantically express a true proposition by uttering this sentence in either case. But (a) and (b) are consistent with any level of skepticism, and many purist invariantists aren’t sufficiently skeptical to deny that Keith knows that the bank will be open in Low Stakes.⁷ Where ‘moderate purist invariantism’ names purist invariantism in conjunction with the proposition that Keith knows that the bank will be open in Low Stakes, many epistemologists are moderate purist invariantists. Since moderate purist invariantists think we can get by in our epistemological theorizing without the contextualist thesis that ‘knows’ can express different properties in different contexts of utterance and the pragmatist thesis that knowledge depends at least in part on our practical interests, I will call moderate purist invariantism, simply, ‘minimalism.’

This paper concerns a specific attempt to defend minimalism from the contextualist or pragmatist conclusions that some epistemologists have wanted to draw from Low Stakes and High Stakes, and other similar cases. In Low Stakes, Keith utters the sentence ‘I know that the bank will be open,’ and then, in High Stakes, he utters the sentence ‘I don’t know that the bank will be open.’ Call these sentences ‘K’ and ‘¬K,’ respectively.

⁴ These cases are slight modifications of the cases in DeRose (1992), p. 913.

⁵ See Hawthorne (2004, pp. 2-3) for the lottery paradox. Stanley (2005) contains the paradigm pragmatist treatment of High Stakes.

⁶ This is not an accident. Contextualism and pragmatism clearly *are* rivals in at least the sense that contextualists and pragmatists can’t both be right about all of the purported motivations for accepting their respective views.

⁷ As we will see below, Wayne Davis is an exception.

Minimalists think that Keith says something false by uttering $\neg K$ in High Stakes. Pragmatists and contextualists think that he says something *true* by uttering $\neg K$ in High Stakes. Everyone agrees that Keith *appears* to say something true by uttering $\neg K$ in High Stakes, at least *prima facie*.⁸ The disagreement concerns the correct explanation of this appearance. As Jeff King and Jason Stanley note (2007, p. 134), a tempting minimalist explanation takes the form of the following “warranted assertability maneuver” (or ‘WAM’ for short):

Keith appears to say something true by uttering $\neg K$ in High Stakes because (a) he *pragmatically conveys* something true by uttering $\neg K$ in High Stakes and (b) we mistake the proposition that he pragmatically conveys by uttering $\neg K$ in High Stakes for what he *says* by uttering $\neg K$ in High Stakes.⁹

In this paper, I argue that WAMs are insufficient for defending minimalism from cases like High Stakes. Thus, I argue, defending minimalism from cases like High Stakes will require adopting a different strategy. In §§1-4, I discuss the WAMs defended by Patrick Rysiew, Jessica Brown, and Wayne Davis, and show why minimalists might be dissatisfied with each of them. In §5, I present my own WAM and suggest that it will succeed if any WAM will succeed. In §6, however, I argue that even my improved WAM must fail. The upshot of my argument will be that, in order to defend minimalism from High Stakes and similar cases, minimalists must prioritize the properly epistemological question whether Keith knows that the bank will be open over linguistic questions about what Keith pragmatically conveys by uttering the sentence ‘I don’t know that the bank will be open.’

1. The Rysiew/Brown WAM

According to Patrick Rysiew (2001), ‘knows’ is not a context sensitive word. It semantically expresses the same properties and relations in every context of utterance. As

⁸ Like others, I stipulatively define ‘says’ so that the proposition one semantically expresses by uttering a sentence in a context is identical to what one says by uttering that sentence in that context.

⁹ Note that, according to the variety of WAM responsible for the moniker ‘warranted assertability maneuver,’ we think Keith says something true by uttering $\neg K$ in High Stakes because we mistake the *warrant* he has for asserting $\neg K$ for the *truth* of the proposition that he asserts by uttering $\neg K$. Since it is implausible that we are mistaking the warrant that Keith has for asserting $\neg K$ for the truth of the proposition that he asserts that he asserts by uttering $\neg K$, I will ignore this variety of WAM. From here forward, I will assume that, if we are mistaking anything at all, we are mistaking a true proposition pragmatically conveyed by some utterance for what is *said* by that utterance.

a result, by uttering $\neg K$ in any context, Keith would semantically express the proposition *that he doesn't know that the bank will be open* (p. 487). Rysiew endorses a relevant alternatives semantics for the word 'knows.' On his view, an utterance of '*S* knows that *p*' expresses a truth only if *S* can rule out all of the relevant $\neg p$ alternatives, where a relevant $\neg p$ alternative is a scenario where *p* is false that is likely to obtain (pp. 487-8). In High Stakes, since Keith can rule out all the likely scenarios where the bank will not be open, and since Keith satisfies all of the other necessary conditions for knowing, Rysiew thinks that Keith says something false by uttering $\neg K$. But an irrelevant alternative may become salient in a context, and Rysiew thinks that an instance of 'I know that *p*' uttered in such a context would implicate that the speaker can rule out this irrelevant alternative (p. 490). When Rachel mentions the possibility that the bank has recently discontinued its Saturday hours, she makes this possibility salient. Given the evidence currently in Keith's possession, Keith cannot rule this possibility out. Since it is *unlikely* that the bank has recently discontinued its Saturday hours, however, this alternative is irrelevant, so it doesn't count against Keith's knowing. Thus, says Rysiew, even though Keith knows that the bank will be open, he cannot felicitously utter *K*. For if he did, he would implicate that he *can* rule out the possibility that the bank has recently discontinued its Saturday hours (p. 490). Since Keith knows that he cannot rule this possibility out, he utters $\neg K$ instead, and thereby implicates that he cannot rule out the possibility that the bank has recently discontinued its Saturday hours. Because the proposition that Keith cannot rule out the possibility that the bank has recently discontinued its Saturday hours is both true and easily mistaken for what Keith says by uttering $\neg K$, says Rysiew, it is tempting to think that Keith says something *true* by uttering $\neg K$ (p. 492).

This is the gist of Rysiew's WAM. According to Jessica Brown, however, Rysiew's WAM fails to account for Keith's *practical interests* (2006, pp. 421-3). Suppose that you take yourself to know that *p* in a situation where it doesn't matter whether *p*, and suppose that I mention some distant but un-eliminable scenario in which $\neg p$. According to Brown (p. 423), it would be perfectly appropriate for you to dismiss the scenario as unlikely and continue to insist that you know that *p*. Similarly, thinks Brown, in Low Stakes (where it doesn't matter when Keith and Rachel go to the bank), Keith could have felicitously uttered *K* in an exchange like the following.

... Rachel raises the possibility that their bank has recently discontinued its Saturday hours and then asks Keith whether he knows that the bank will be open tomorrow. Keith responds, "Yeah, I *do* know that the bank will be open

tomorrow. It's so unlikely that it has discontinued its Saturday hours since my last visit that we needn't worry about that possibility."

Even though Keith points out that their bank has almost certainly not changed its Saturday hours, the possibility that it *has* is still salient. Since Keith can't rule this possibility out, Rysiew's view predicts that Keith's utterance of K isn't felicitous. According to Brown, however, it *is* felicitous. Brown thinks Rysiew's account is at least on the right track, though, so she attempts to develop his account to capture the impact of practical interests on our intuitions about Keith's utterance of $\neg K$ in High Stakes. To do this, Brown starts with Grice's Maxim of Relation (1989, p. 32) and his well-known example of the motorist who tells the pedestrian that she is running out of gas. In Grice's example, the pedestrian utters the sentence 'there is a garage nearby' and thereby implicates that there is an *open* garage nearby. The scenario that Brown imagines differs slightly from Grice's. In Brown's scenario, there is only one garage nearby and the pedestrian knows that this garage is *closed*, so she utters the sentence 'there is no garage nearby.' Brown thinks that, just as the pedestrian in Grice's petrol example implicates that there is an open garage nearby by uttering just 'there is a garage nearby,' the pedestrian in her example implicates that there *is not* an open garage nearby by uttering just, 'there is no garage nearby.' According to Brown, "[w]hile this utterance is literally false, it pragmatically conveys the true claim that there is no open garage nearby" (p. 425). Brown thinks that an utterance of K or $\neg K$ in High Stakes would, respectively, be analogous to an utterance of 'there is a garage nearby' or 'there is no garage nearby' in her petrol scenario (p. 426). In High Stakes, Keith would have implicated a false proposition by uttering K. By uttering $\neg K$ instead, he thereby implicated a *true* proposition. Which proposition, exactly? In High Stakes, the salience of Keith's practical situation makes it clear that Keith needs to be in a very strong epistemic position with respect to the proposition that the bank will be open on Saturday. According to Brown, had Keith uttered K in High Stakes, this is exactly what he would have implicated. More specifically, says Brown, he would have implicated that his belief that the bank will be open matches the facts out to the nearest world in which the bank has recently discontinued its Saturday hours. Let '*m*' name the proposition that Keith's belief that the bank will be open matches the facts out to the nearest world in which the bank has recently discontinued its Saturday hours. Since Keith uttered $\neg K$ instead of K, says Brown, he implicated $\neg m$ instead of *m*. And since $\neg m$ is *true*, he implicated a *true* proposition by uttering $\neg K$ in High Stakes. Finally, says Brown, when we read High Stakes and consider Keith's utterance of $\neg K$, his utterance seems true because it

con conversationally implicates that $\neg m$, and we mistake $\neg m$ for what Keith says by uttering $\neg K$ (p. 424).

2. A Worry about the Generality of the Rysiew/Brown WAM

Brown's development of Rysiew's WAM does seem like an improvement, since it captures the impact of Keith's practical interests on our intuitions about High Stakes. It is clear that we might improve on both WAMs, however. Consider the following revision of High Stakes, which focuses *solely* on Keith's practical interests. Unlike the original version of High Stakes, this version does not mention the possibility that the bank has discontinued its Saturday hours since Keith's last visit.

Practical High Stakes: Rachel and Keith are driving home on Friday afternoon. They plan to stop at the bank to deposit their paychecks, but as they approach it they notice that the lines inside are very long. Keith suggests that they go straight home and return to deposit their paychecks on Saturday morning, explaining that he stops at the bank every Saturday morning and that it's open until noon. But in this case, Rachel and Keith have just written a very large and important check. If their paychecks are not deposited before Monday morning, the important check they wrote will bounce, leaving them in a terrible situation. The bank isn't open on Sunday, and Rachel reminds Keith of this fact. She then says, "It will be completely disastrous for us if we go home now and the bank isn't open tomorrow. Do you know that it will be open tomorrow?" Remaining just as confident as he was before that the bank will be open tomorrow, Keith replies, "Well, no, I don't know that the bank will be open tomorrow. We had better stop and deposit the checks now."

Just as Keith appears to say something true by uttering $\neg K$ in the original version of High Stakes, he appears to say something true by uttering $\neg K$ in this version of High Stakes. But in this version of High Stakes, nobody mentions any specific reason why the bank might not be open on Saturday, including the possibility that the bank has recently discontinued its Saturday hours. In fact we can stipulate that, in Practical High Stakes, neither Keith nor Rachel considers the possibility that the bank has recently discontinued its Saturday hours, and we can also stipulate that, in Practical High Stakes, whatever specific reasons they *are* considering, they aren't considering any of the *same* reasons. Giving these stipulations, we cannot explain why Keith seems to say something true by

uttering $\neg K$ in Practical High Stakes in terms of anything as specific as the possibility that the bank has recently discontinued its Saturday hours. Since Rysiew's and Brown's WAMs are just this specific, they don't explain why Keith seems to say something true by uttering $\neg K$ in Practical High Stakes. To explain why Keith seems to say something true by uttering $\neg K$ in Practical High Stakes, we need a more general WAM than the WAMs forwarded by Rysiew and Brown.

3. A WAM with General Application

While it is clear enough how Rysiew and Brown could revise their WAMs to meet this need, Wayne Davis's WAM is sufficiently general as it stands, with literally no need for revision. According to Davis (2007, p. 414), knowledge is "completely and non-defectively justified true belief," where a completely justified belief must be "either self-evident or based on evidence sufficient to establish its truth." For Davis, establishing the truth of a proposition is no easy task. Keith's evidence that the bank will be open doesn't establish the truth of the proposition that the bank will be open, for example, since "there are too many very real possibilities that [he] cannot rule out: that the bank changed its hours since [he] last checked, that an unexpected thunderstorm would lead to a massive power outage, and so on" (p. 414). Since the proposition that the bank will be open is not self-evident, and since Davis thinks that Keith's evidence doesn't establish the truth of this proposition, Davis agrees with pragmatists and contextualists that Keith says something true by uttering $\neg K$ in High Stakes, and he *disagrees* with them about Low Stakes. According to Davis, Keith says something *false* by uttering K in Low Stakes. To handle our intuition that Keith says something *true* by uttering K in Low Stakes, Davis then develops a WAM according to which, by uttering K in Low Stakes, Keith implicates that he is close enough to knowing for present purposes.

On Davis's view, Keith's utterance of $\neg K$ in High Stakes and his utterance of K in Low Stakes are respectively analogous to Kim's utterances of the sentences 'the coffee is all gone' and 'no, there may be enough for you' in the case below (p. 406).

Coffee: Kim takes time out from studying to make a cup of coffee. When the scoop comes up empty, she yells to Jerome, "the coffee is all gone!" Eric hears Kim yell this. After telling her that he needs a few coffee grounds for his science project, he then asks whether the coffee is really all gone. Kim responds without embarrassment, "no, there may be enough for you."

On Davis's view, Kim said something true when she uttered the sentence 'no, there may be enough for you,' but something false when she uttered the sentence 'the coffee is all gone.' She *implicated* that the coffee is close enough to being all gone for present purposes, however, and this *is* true. The reason she appeared to say something true by uttering 'the coffee is all gone,' on Davis's view, is that she implicated something true by uttering this sentence (p. 407). According to Davis, we should say roughly the same thing about Keith's utterance of K in Low Stakes. Since Keith does not know that the bank will be open in Low Stakes, Keith says something false by uttering K in Low Stakes. But Keith *is* close enough to knowing for present purposes, and this is what he implicates by uttering K in Low Stakes. The reason Keith appears to say something true by uttering K in Low Stakes, says Davis, is that he implicates something true and we mistake the true proposition that he implicates for what he *says* by uttering K in Low Stakes (p. 408).

Now, as should be obvious, *minimalists* cannot rely on this WAM for defending their view from Low and High Stakes, since minimalism maintains that Keith says something *true* by uttering K in Low Stakes. According to Davis, however, minimalists can simply revise his account of knowledge to develop a similar WAM for High Stakes. With a more permissive account of knowledge, says Davis, minimalists can maintain that Keith says something false by uttering $\neg K$ in High Stakes (since he *does* know that the bank will be open), but implicates the true proposition that, for present purposes, he is too close to *not* knowing that the bank will be open. According to Davis, minimalists can then explain why Keith seems to say something true by uttering $\neg K$ in High Stakes by insisting that we mistake what Keith implicates by uttering $\neg K$ in High Stakes for what he says by uttering $\neg K$ in High Stakes (p. 414).

4. A Possible Dissatisfaction with Davis's WAM

Davis's WAM seems preferable to Rysiew's WAM and Brown's WAM in at least one respect. In both High Stakes and Practical High Stakes, Rachel and Keith are deciding whether they should wait in line to deposit their checks Friday evening, or go straight home and deposit them Saturday morning. In both cases, deciding whether they should wait in line or go straight home is the sole purpose of their conversation. According to the WAM that Davis outlines, in High Stakes, Keith implicates that, for present purposes, he is too close to *not* knowing that the bank will be open. If Keith implicates this in High Stakes, he surely also implicates this in Practical High Stakes, since High Stakes and Practical High Stakes are identical with respect to the purpose of Rachel and Keith's conversation. So Davis's proposed WAM seems to work just as well for Practical

High Stakes as High Stakes. I find Davis’s WAM slightly dissatisfying for a different reason, however—a reason that I think counts against Brown’s WAM, but not Rysiew’s WAM.

As we noted above, when minimalists say that Keith implicates something true in High Stakes, and when they say that we mistake what Keith implicates in High Stakes for what he says in High Stakes, they are talking about *conversational* implicatures. According to a view often associated with Grice (1989), however, conversational implicatures are always explicitly cancellable. Call this ‘the strong Gricean view.’¹⁰ I find the strong Gricean view implausible, but some philosophers do not, and, if the strong Gricean view is correct, then Brown’s WAM and Davis’s WAM both rely on false claims.¹¹ Again, consider the conversation that Grice gives us in his well-known petrol example.

Motorist: I’m running out of gas. Is there a garage nearby?

Pedestrian: Yeah, there is a garage nearby—right around the corner.

In this conversation, the pedestrian utters ‘there is a garage nearby’ and thereby implicates that there is an open garage nearby. Had the garage been closed, however, the pedestrian could have explicitly cancelled this implicature by adding ‘but it is closed.’ That is, had the pedestrian uttered ‘yeah, there is a garage nearby—but it is closed,’ she would have successfully conveyed to the motorist that there is a garage nearby without conveying to the motorist that there is an open garage nearby.

Now, as we saw above, Brown takes it for granted that, because the pedestrian implicates that there is an open garage nearby by uttering just ‘there is a garage nearby,’ the pedestrian would have implicated that there is *not* an open garage nearby had she

¹⁰ I call this the *strong* Gricean view because it is arguable that Grice only endorses a weaker claim. Consider the difference between SC and WC below.

(SC) A putative conversational implicature is a genuine conversational implicature only if it is *both* explicitly *and* contextually cancellable.

(WC) A putative conversational implicature is a genuine conversational implicature only if it is *either* explicitly *or* contextually cancellable.

As an anonymous referee has helpfully brought to my attention, there are defensible interpretations of Grice on which Grice only endorses WC. This is important because only the stronger principle (SC) says that explicit cancellability is a necessary condition on genuine conversational implicature.

¹¹ See Davis (1998) and Bach (2006) for cases where someone plausibly conversationally implicates something that is an obvious entailment of what is said.

uttered just ‘there is no garage nearby.’ If the strong Gricean view is correct, however, then Brown is mistaken, since the proposition that there is no open garage nearby is not explicit cancellable. Consider the following answers that the pedestrian might have given to the motorist.

A1: No, there’s no garage nearby, but there is an open garage nearby.

A2: No, there’s no garage nearby, but I am not saying that there is no open garage nearby.

When the pedestrian utters ‘there is no garage nearby,’ she thereby communicates to the motorist that there is no open garage nearby. But if the strong Gricean view is correct, she does not conversationally implicate that there is no open garage nearby, since, downstream from an utterance of ‘there is no garage nearby,’ the proposition that there is no open garage nearby is not explicitly cancellable. As A1 and A2 show us, once the pedestrian has uttered ‘there is no garage nearby,’ it is not (in Grice’s words) admissible for the pedestrian to add that there is an open garage nearby, or that she is not saying that there is no open garage nearby, or anything like that. So, if the strong Gricean view is correct, then Brown is mistaken when she says that the pedestrian’s utterance of ‘there is no garage nearby’ implicates the true proposition that there is no open garage nearby.

If she is mistaken about this, however, then Davis is also mistaken when he says that, by uttering ‘the coffee is all gone’ in his coffee case, Kim implicates that the coffee is close enough to being all gone for the purpose of making coffee. As Davis himself notes, just as the pedestrian cannot utter the sentence ‘there is no garage nearby’ and then admissibly add that there is an open garage nearby, Kim cannot utter ‘the coffee is all gone’ and then admissibly add that there is enough left to make coffee (p. 411). So, both Brown’s WAM and Davis’s WAM rely on claims that turn out false if the strong Gricean view is correct.

As I said above, I doubt that the strong Gricean view is correct. Instead of rehearsing the arguments against the strong Gricean view, however, I want to provide a WAM that succeeds *whether or not* the strong Gricean view is correct. In the next section, I present a WAM that is both sufficiently general and immune to whatever worries might arise from the strong Gricean view. This WAM should satisfy both philosophers who accept the strong Gricean view and philosophers who reject it.

5. A Sufficiently General, Explicitly Cancellable WAM

Return once again to Grice's petrol case. In that case, the pedestrian utters 'there is a garage nearby' and thereby implicates that there is an open garage nearby. What explains this implicature? According to Grice's Cooperative Principle, the pedestrian must make her conversational contribution such as is required, at the stage at which it occurs, by the accepted purpose or direction of the conversation in which she is engaged (1989, p. 26). According to Grice's Maxim of Relation, the pedestrian must make her conversational contribution relevant (p. 27). The pedestrian does not seem to be flouting the Maxim of Relation when she tells the motorist that there is a garage nearby, and we have no reason to think that she is being uncooperative by telling the motorist just that there is a garage nearby. Yet the pedestrian must regard her own utterance as irrelevant if she does not think that the nearby garage is *open*. After all, the motorist obviously needs to buy gas, and the pedestrian (we may presume) knows that a closed garage is useless for this purpose. So, preserving the assumption that the pedestrian is observing the Cooperative Principle requires assuming that the pedestrian believes that the nearby garage is open. *This* (according to Grice) is why the pedestrian implicates that there is an open garage nearby by uttering just 'there is a garage nearby' (p. 32).

If this is why the pedestrian implicates that there is an open garage nearby by uttering 'there is a garage nearby,' however, then it seems that the pedestrian also implicates *that the motorist can get gas somewhere nearby* by uttering 'there is a garage nearby.' Just as the pedestrian must regard her utterance as irrelevant if she does not think that the nearby garage is open, she must regard her utterance as irrelevant if she does not think that the motorist can get gas somewhere nearby. After all, the pedestrian knows that the motorist is running out of gas. Preserving the assumption that the pedestrian is observing the Cooperative Principle requires assuming *both* that the pedestrian believes that the nearby garage is open and that the pedestrian believes that the motorist can get gas somewhere nearby. So, by uttering 'there is a garage nearby,' the pedestrian implicates both that there is an open garage nearby and that the motorist can get gas somewhere nearby.

Now consider the following conversation where the pedestrian gives a *negative* answer to the motorist's question.

Motorist: I'm running out of gas. Is there a garage nearby?

Pedestrian: No, there's no garage nearby. The nearest one's in Sleetmute.

As we saw above, if the strong Gricean view is correct and all conversational implicatures are explicitly cancellable, then the pedestrian does not conversationally implicate that there is no open garage nearby when she utters ‘there is no garage nearby.’ Plausibly, however, *even if* the strong Gricean view is correct, the pedestrian *does* implicate that the motorist cannot get gas anywhere nearby when she utters ‘there is no garage nearby.’ First, just as preserving the assumption that the pedestrian is observing the Cooperative Principle in Grice’s original petrol case requires assuming that the pedestrian believes that the nearby garage is open, preserving the assumption that the pedestrian is observing the Cooperative Principle in this case requires assuming that the pedestrian believes that the motorist cannot get gas anywhere nearby. Second, the suggestion that the pedestrian implicates that the motorist cannot get gas anywhere nearby easily passes the explicit cancellability test. Consider the following conversation.

Motorist: I’m running out of gas. Is there a garage nearby?

Pedestrian: No, there’s no garage nearby, but there *is* a place nearby where you can get gas. They sell it out at the airfield.

It would be perfectly admissible for the pedestrian to utter ‘there is no garage nearby’ and then add that there *is* a place nearby where the motorist can get gas. By uttering ‘there is no garage nearby, but there is a place nearby where you can get gas,’ the pedestrian successfully conveys to the motorist that there is no garage nearby *without* conveying to the motorist she cannot get gas anywhere nearby. Her reply to the motorist is perfectly felicitous. The proposition that the motorist cannot get gas anywhere nearby seems just as much an explicitly cancellable implicatum of the pedestrian’s utterance of ‘there is no garage nearby’ as the proposition that there is an open garage nearby seems to be an explicitly cancellable implicatum of the pedestrian’s utterance of ‘there is a garage nearby.’

Now consider the suggestion that, by uttering $\neg K$ in High Stakes, Keith implicates *that Rachel and Keith cannot reasonably go straight home*. As we noted above, Rachel’s sole purpose in questioning Keith is to determine the best course of action, and it is transparent what she means to accomplish by asking Keith if he knows that the bank will be open. Just as the motorist asks the pedestrian if there is a garage nearby to find out whether she can get gas nearby, Rachel asks Keith if he knows that the bank will be open to find out whether they can reasonably go straight home. (Conditional on the proposition that the bank will be open, going straight home is clearly the best course of

action.) Given this, however, minimalists can point out that, just as preserving the assumption that the pedestrian is observing the Cooperative Principle in Grice's original petrol case requires assuming that the pedestrian believes that the nearby garage is open, preserving the assumption that Keith is observing the Cooperative Principle in High Stakes requires assuming that Keith believes that he and Rachel cannot reasonably go straight home. This similarity motivates the claim that, by uttering $\neg K$ in High Stakes, Keith implicates that Rachel and Keith cannot reasonably go straight home.

Is this purported implicature explicitly cancellable, then? Yes, for consider the following alternative ending to High Stakes.

... If their paychecks are not deposited before Monday morning, the important check they wrote will bounce, leaving them in a terrible situation. The bank isn't open on Sunday. Rachel reminds Keith of these facts. She then says, "What if our bank discontinued its Saturday hours since your last visit? Do you know that it will be open?" Remaining just as confident as he was before that the bank will be open, Keith replies, "Well, no, I don't know that the bank will be open, but we can still reasonably go straight home. I know the president of the bank. If the bank is closed tomorrow, I will give her a call and she will let us deposit the checks anyway."

It seems perfectly admissible for Keith to utter 'I don't know that the bank will be open' to Rachel and then add that they can still reasonably go straight home, so the suggestion that Keith implicates that they cannot reasonably go straight home by uttering $\neg K$ passes Grice's explicit cancellability test. Here, then, is the proposed WAM.

By uttering $\neg K$, Keith pragmatically conveys that he and Rachel cannot reasonably go straight home. Keith seems to say something true by uttering $\neg K$ because he and Rachel *cannot* reasonably go straight home and we mistake the proposition that they cannot reasonably go straight home for what Keith *says* by uttering $\neg K$ in High Stakes.

Of course, this WAM works just as well for Practical High Stakes as for High Stakes, since it is sensitive to the decision-theoretic nature of both cases. So, to the strong Gricean, or to anyone who thinks the strong Gricean view is even plausible, this WAM will seem a clear improvement over Rysiew's WAM, Brown's WAM, and Davis's WAM. If any WAM will put minimalists in good position to defend their view from High Stakes, this one should.

6. The Limits of WAMs (Gricean or Otherwise)

Sadly, I don't think minimalists can defend their view from High Stakes with my WAM or with any other WAM. In §2, we saw that there is little point in responding to High Stakes with a particular WAM if that WAM cannot handle a slight revision of High Stakes. My WAM and Davis's WAM can both handle revisions of High Stakes like Practical High Stakes, so they have at least that advantage over less general WAMs. But consider the following revision of High Stakes, which is analogous to Fantl and McGrath's high-stakes train case (2002, pp. 67-8), Schroeder's high-stakes bank case (2012, pp. 266-7), Weatherson's high-stakes music case (2012, pp. 82-3), and Ross and Schroeder's high-stakes sandwich case (2014, p. 261), in that it doesn't involve a single utterance of any knowledge-ascribing or knowledge-denying sentence.

Silent High Stakes: Keith is driving to the bank on Friday afternoon when he notices the long lines. He knows that the bank has been open every previous Saturday, but he also knows that there will be disastrous consequences if he fails to deposit his paycheck before Sunday. Given the quality of Keith's evidence that the bank will open, balanced against the severe consequences that Keith knows will follow if he fails to deposit his paycheck before Sunday, Keith cannot rationally go straight home. Yet Keith knows that, conditional on the bank's being open, going straight home is the best of his available options.

Just like Fantl and McGrath's high-stakes train case, Schroeder's high-stakes bank case, Weatherson's high-stakes music case, and Ross and Schroeder's high-stakes sandwich case, Silent High Stakes involves no utterance of 'knows' or any relevantly related word. Yet many philosophers would have the exactly same response to Silent High Stakes as they have to the other "silent" cases in the pragmatic encroachment literature: they would conclude that, in Silent High Stakes, Keith's belief that the bank will be open must fail to satisfy some necessary condition for knowledge. Some of these philosophers would conclude that Keith's belief falls short of knowledge because they find it *prima facie* intuitive that Keith lacks knowledge in Silent High Stakes. Others would conclude that Keith's belief falls short of knowledge because they see that this conclusion follows from principles that they find independently plausible, such as Fantl and McGrath's Knowledge-Justification Principle (2009, p. 66).¹² Indeed, even many *minimalist* would

¹² To give a few more examples, Hawthorne and Stanley's Knowledge-Reasons Principle (2008), and Ross and Schroeder's Knowledge-Action Principle (2014) both entail that Keith does not know that the bank will be open in Silent High Stakes. Also, assuming that one cannot know that *p* if one cannot believe that *p* except by

admit some attraction (at least *prima facie*) to the proposition that Keith fails some necessary condition for knowledge in Silent High Stakes, perhaps because they feel the initial pull of the case, or perhaps because they appreciate the motivation for accepting the relevant principles.¹³ Minimalists interested in responding to High Stakes with just a WAM now have a problem.

Keith's belief that the bank will be open falls short of knowledge in Silent High Stakes just in case it also falls short of knowledge in the original version of High Stakes, since, by hypothesis, there is no epistemically relevant difference between Keith's belief that the bank will be open in High Stakes and his belief that the bank will be open in Silent High Stakes. The intuitions and principles that motivate the conclusion that Keith lacks knowledge in Silent High Stakes therefore provide just as much motivation for the conclusion that he lacks knowledge in High Stakes. But there is no Gricean story to tell about Silent High Stakes. There is no way to apply the concepts, distinctions and principles familiar from Grice to this case. Keith does not utter K or $\neg K$ in Silent High Stakes, or even *mentally token* K or $\neg K$ in this version of the case, so Gricean considerations gain no purchase here. And the same goes for the sentences 'Keith knows that the bank will be open' and 'Keith does not know that the bank will be open.' These sentences make no appearance in Silent High Stakes. Of course, philosophers who find it plausible that Keith lacks knowledge in Silent High Stakes might mentally token one or both of these sentences as they consider the case, or even utter these sentences out loud. But are we supposed to think that, when these philosophers utter or mentally token the sentence 'Keith does not know that the bank will be open in Silent High Stakes,' they thereby implicate *to themselves* some true proposition that they then mistake for what they would say by uttering 'Keith does not know that the bank will be open in Silent High Stakes,' and that *this* is why they find it plausible that Keith's belief that the bank will be open fails some necessary condition for knowledge in Silent High Stakes? This suggestion seems implausible. Surely, very few philosophers would arrive at the belief that Keith lacks knowledge in Silent High Stakes this way. So it seems doubtful that minimalists can adequately argue that Keith knows that the bank will be open in Silent High Stakes by relying on solely the principles and distinctions familiar from Grice. Since Keith lacks knowledge in Silent High Stakes only if he also lacks knowledge in *High Stakes*, it seems equally doubtful that minimalists can adequately argue that Keith knows

being irrationally over-confident that p , Ganson's (2008) and Weatherson's (2012) theories of belief both entail that Keith does not know that the bank will be open in Silent High Stakes.

¹³ Thanks to an anonymous referee for suggesting that I point this out.

that the bank will be open in High Stakes by relying on solely the principles and distinctions familiar from Grice.¹⁴

And the problem here isn't just with *Gricean* pragmatics. The problem generalizes. I have focused exclusively on Grice because the most plausible extant WAM's all rely exclusively on Gricean considerations. But everything I say about Grice in this context applies to the entire field of pragmatics. As Korta and Perry note (2012, p. 1), while semantics deals with the properties of types of expressions, pragmatics deals with the properties of specific *utterances*. This is what distinguishes pragmatics from semantics. Since nobody utters anything in Silent High Stakes, minimalists who want to rely on pragmatic considerations in response to Silent High Stakes must focus on some utterance that is not part of the case itself—either some imaginary utterance of $\neg K$ by Keith in Silent High Stakes, or some real or imaginary utterance of the sentence 'Keith doesn't know that the bank will be open in Silent High Stakes' by the reader herself (or perhaps by someone in her vicinity). Now, plausibly, someone might read Silent High Stakes, imagine herself uttering the sentence 'Keith doesn't know that the bank will be open' in response to Silent High Stakes, think to herself that this would be a perfectly appropriate way to respond to the case, and, on the basis of the apparent propriety of this response, conclude that Keith does not know that the bank will be open in Silent High Stakes. If I happen to arrive at the conclusion that Keith lacks knowledge in Silent High Stakes *this way*, minimalists might be able to tell some story according to which (a) the imaginary utterance would be appropriate because, in the imagined context, it pragmatically conveys something true, and (b) I am misled about Silent High Stakes because I am mistaking the true proposition that my imaginary utterance would pragmatically convey for what I would say by uttering the sentence 'Keith doesn't know that the bank will be open' in the imagined context.

The problem is, even if this story happens to work for *my* conviction that Keith lacks knowledge in Silent High Stakes, many philosophers will arrive at the conclusion that Keith lacks knowledge in Silent High Stakes without evaluating any real or imaginary utterance of any sentence. As I stated above, many philosophers will simply find it *prima facie* intuitive that Keith's belief that the bank will be open fails some necessary condition for knowledge in Silent High Stakes, and others will note that this conclusion follows

¹⁴ I do not mean to suggest that, *in general*, it is impossible to implicate things to ourselves by mentally tokening sentences. As Jason Stanley has pointed out in conversation, if I know that it's raining and it's obvious to me that I know that it's raining, it would be odd for me to *think to myself* 'it might be raining.' I don't want to rule out the possibility that the correct explanation here is at least in part that, by thinking to myself 'it might be raining,' I would implicate to myself that I don't know whether it is raining.

from principles that they find independently plausible. As Brown and Gerken point out (2013, pp. 7-8), neither the intuition that Keith lacks knowledge in Silent High Stakes nor the derivation of this conclusion from some pragmatist-friendly principle must proceed through any evaluation of any real or imaginary utterance of anything. But this means that we haven't simply found a limitation in the work that *Gricean* considerations can do for minimalists. The entire field of pragmatics seems incapable of supplying sufficient resources for defending minimalism from cases like Silent High Stakes. And since Keith lacks knowledge in Silent High Stakes only if he also lacks knowledge in *High Stakes*, it seems that the field of pragmatics cannot supply adequate resources for defending minimalism from cases like High Stakes. Thus, minimalists cannot successfully defend their view from cases like High Stakes by relying solely on WAMs.

7. Conclusion

If minimalists cannot successfully defend their view from cases like High Stakes by relying solely on WAMs, how *should* minimalists respond to cases like High Stakes? I believe they should prioritize the epistemological question whether the subjects in those cases know over linguistic questions about what is pragmatically conveyed by utterances of the relevant knowledge-ascribing and knowledge-denying sentences.

As we just saw, even if minimalists can successfully argue that Keith pragmatically conveys something true by uttering $\neg K$ in High Stakes, this victory will be useless for convincing many philosophers that Keith knows that the bank will be open in Silent High Stakes. To address the sensibilities of these philosophers, minimalists must take a stand on the necessary conditions for knowledge and argue that, in Silent High Stakes, Keith's belief that the bank will be open meets all of them. A successful argument here will require at least two things. First, it will require a successful argument for the conclusion that Keith knows that the bank will be open in Silent High Stakes.¹⁵ Second, it will require an adequate explanation of the fact that, at least initially, Keith seemed to *lack* knowledge in Silent High Stakes. Suppose, then, that minimalists acquire both a convincing argument for the conclusion that Keith knows that the bank will be open in Silent High Stakes and a satisfactory explanation of the initial appearance to the contrary. In this case, what residual need will minimalists have for a WAM?

¹⁵ Among other things, these arguments will have to cast sufficient doubt on the pragmatist-friendly principles defended in Fantl and McGrath (2002), Ganson (2008), Hawthorne and Stanley (2008), Fantl and McGrath (2009), Weatherson (2012), Schroeder (2012), and Ross and Schroeder (2014).

Minimalists formulated WAMs because they needed to explain why Keith seemed to say something true by uttering the sentence ‘I don’t know that the bank will be open’ in High Stakes. But supposing that minimalists now have both convincing arguments that Keith knows that the bank will be open in Silent High Stakes and an adequate explanation of the fact that, at least initially, Keith seemed to *lack* knowledge in Silent High Stakes, why can’t minimalists simply rely on an argument like the one below?

“Now that it’s clear that Keith knows that the bank will be open in Silent High Stakes, and clear that, if Keith knows that the bank will be open in Silent High Stakes, he also knows that the bank will be open in High Stakes, Keith *doesn’t* seem to say something true by uttering the sentence ‘I don’t know that the bank will be open’ in High Stakes. The reason Keith initially seemed to say something true by uttering this sentence in High Stakes is simply that, before we had all of the relevant considerations in view, he seemed to lack knowledge in High Stakes. And the explanation for *this* fact is the same as the explanation for the fact that, before we had all of the relevant considerations in view, he seemed to lack knowledge in *Silent* High Stakes—which is identical to High Stakes on every relevant epistemic dimension. The explanation is that”

Assuming that minimalists have convincing arguments for the conclusion that Keith knows that the bank will be open in Silent High Stakes and a satisfactory explanation for why Keith initial seemed to lack knowledge in Silent High Stakes, I see no reason why this line of reasoning should be unconvincing. But this line of reasoning will not involve a WAM, since it will appeal only to the considerations that explain *Silent* High Stakes. So it seems that, if minimalists can acquire convincing arguments for the conclusion that Keith knows that the bank will be open in Silent High Stakes and a satisfactory explanation for why Keith initial seemed to lack knowledge in Silent High Stakes, they will have no further need for a WAM.¹⁶

* * *

¹⁶ Of course, the reason *Davis* will not need a WAM to respond to High Stakes or Silent High Stakes is that he isn’t a minimalist. On his view, Keith seems to lack knowledge in High Stakes and Silent High Stakes because he *does* lack knowledge in those cases (since his belief isn’t completely justified, as noted above). If Davis faces a challenge analogous to the challenge facing Rysiew, Brown, and other minimalists forwarding WAMs, this challenge will arise from silent versions of the relevant *low stakes* cases, not from Silent High Stakes or any analogous case.

It seems a mistake to expend resources on something that you won't need once you acquire some other thing that you can't get by without. If you can't get by without a computer, for example, and if, having a computer, you won't need a typewriter, you shouldn't waste resources on a typewriter. As we have seen, minimalists can't get by without (a) convincing arguments for the conclusion that Keith knows that the bank will be open in Silent High Stakes and (b) a satisfactory explanation for why Keith initially seemed to lack knowledge in Silent High Stakes. But, once minimalists have these things, responding to High Stakes will not require a WAM. The upshot is that, at best, WAMs will play an insignificant role in any adequate minimalist response to High Stakes. Of course, there is nothing special about *High Stakes*. The same lesson applies to all of the high-stakes cases in the literature that involve knowledge ascriptions or denials. WAMs will play an insignificant role in any adequate minimalist response to *any* of these cases. Questions about the pragmatics of knowledge-ascribing and knowledge-denying sentences are interesting in their own right, so I don't want to suggest that Rysiew, Brown, Davis, and I have wasted our time by asking how pragmatic considerations bear on utterances such as Keith's utterance of $\neg K$ in High Stakes. But I do want to insist that, in response to cases like High Stakes, minimalists should turn their attention away from questions about the implicatures produced by utterances of the relevant sentences and, instead, prioritize the properly epistemological question whether the subjects in those cases know.¹⁷

References

- Bach, Kent (2006). "The Top 10 Misconceptions about Implicature," in *Drawing the Boundaries of Meaning*, Betty Birner and Gregory Ward, Eds. (Amsterdam: John Benjamins Publishing), pp. 21-30.
- Blome-Tillmann (2008), "Conversational Implicature and the Cancellability Test," *Analysis*, Vol. 68, No. 2: pp. 156-160.
- Brown, Jessica (2006). "Contextualism and Warranted Assertability Manoeuvres," *Philosophical Studies*, Vol. 130: pp. 407-35.

¹⁷ Thanks to Matt Benton, Lisa Miracchi, Ernest Sosa, Kurt Sylvan, and an anonymous referee for helpful conversation and comments on this paper.

- Brown, Jessica and Mikkel Gerken (2013). "Knowledge Ascriptions: Their Semantics, Cognitive Bases, and Social Functions," in *New Essays on Knowledge Ascriptions*, eds. J. Brown and M. Gerken (Oxford: Oxford University Press).
- Cohen, Stewart (1988). "How to be a Fallibilist," *Philosophical Perspectives*, Vol. 2, Epistemology: pp. 91-123.
- Davis, Wayne (1998). *Implicature: Intention, Convention, and Principle in the Failure of Gricean Theory* (Cambridge: Cambridge University Press).
- (2007). "Knowledge Claims and Context: Loose Use," *Philosophical Studies*, Vol. 132: pp. 395-438.
- DeRose, Keith (1992). "Contextualism and Knowledge Attributions," *Philosophy and Phenomenological Research*, Vol. 52: pp. 913-29.
- (1995). "Solving the Skeptical Problem," *Philosophical Review*, Vol. 104: pp. 1-52.
- (2002). "Assertion, Knowledge and Context," *Philosophical Review*, Vol. 111: pp. 167-203.
- (2009). *The Case for Contextualism* (Oxford: Oxford University Press).
- Fantl, Jeremy and Matthew McGrath (2002). "Evidence, Pragmatics and Justification," *The Philosophical Review*, Vol. 111: pp. 67-94.
- (2009). *Knowledge in an Uncertain World* (Oxford: Oxford University Press).
- Ganson, Dorit (2008). "Evidentialism and Pragmatic Constraints on Outright Belief," *Philosophical Studies*, Vol. 139: pp. 441-58.
- Grice, Paul (1989). *Studies in the Ways of Words* (Cambridge, MA: Harvard University Press).
- Grimm, Stephen (2011), "On Intellectualism in Epistemology," *Mind*, Vol. 120, No. 479: pp. 705-733.
- Hawthorn, John (2004). *Knowledge and Lotteries* (Oxford: Oxford University Press).
- Hawthorne, John and Jason Stanley (2008), "Knowledge and Action," *The Journal of Philosophy*, Vol. 105, No. 10: 571-90.

- Korta, Kepa and John Perry (2012). "Pragmatics," *The Stanford Encyclopaedia of Philosophy*, <http://plato.stanford.edu/archives/win2012/entries/pragmatics/>.
- Ross, Jacob and Mark Schroeder (2014), "Belief, Credence and Pragmatic Encroachment," *Philosophy and Phenomenological Research*, Vol. 88, No. 2: pp. 259-88.
- Rysiew, Patrick (2001). "The Context Sensitivity of Knowledge Attributions," *Noûs*, Vol. 35: pp. 477-514.
- Schroeder, Mark (2012), "Stakes, Withholding, and Pragmatic Encroachment on Knowledge," *Philosophical Studies*, Vol. 160, No. 2: pp. 265-85.
- Stanley, Jason (2005). *Knowledge and Practical Interests* (Oxford: Oxford University Press).
- Stanley, Jason and Jeffrey C. King (2007). "Semantics, Pragmatics, and the Role of Semantic Content," in *Language in Context*, ed. J. Stanley (Oxford: Oxford University Press), pp. 133-81.
- Weatherson, Brian (2012). "Knowledge, Bets and Interests," in *New Essays on Knowledge Ascriptions*, eds. J. Brown and M. Gerken (Oxford: Oxford University Press).