

Knowledge Attribution, Warranted Assertability Manoeuvre and the Maxim of Relation

JACQUES-HENRI VOLLET

Abstract According to the traditional account of knowledge, what turns a true belief into knowledge is only a matter of truth-relevant factors (evidence, reliability, etc.), and these epistemic standards do not vary across contexts. Still, our knowledge attributions seem to vary depending on practical factors. Some philosophers account for this variation by a shift in the warranted assertability conditions. The main way of fleshing out this proposal is based on the idea of conversational implicature generated by the maxim of relation (Rysiew 2001, 2005, Brown 2006). In this paper, I argue that this is not promising. The proposition that is supposed to be implicated in such cases concerns whether the subject is in a position to eliminate a salient alternative that is not knowledge-destroying. I will claim that in such contexts, this consideration is not more relevant than the question whether the subject knows, even if the stakes are high.

1. Introduction

The bank cases

Suppose that two subjects truly believe that *p*, and are in the same epistemic position with respect to *p*. According to the traditional account of knowledge, both of these subjects know or do not know. The traditional account of knowledge is intellectualist. What turns a true belief into knowledge is only a matter of truth-relevant factors (evidence, reliability, etc.). It is also invariantist. The epistemic standards that a subject must meet in order to count as knowing do not vary with the context.

Against this traditional view, some philosophers have argued that the truth-value of knowledge attributions (utterances of sentences of the form “*S* knows that *p*”) is sensitive to practical factors. There are two main ways of understanding this idea. For epistemic contextualists, “know” is a context-sensitive term. What is semantically expressed by an utterance of this word is determined by the context of utterance, and this context includes practical factors.¹ For subject-sensitive invariantists, there is a pragmatic condition on knowledge that must be satisfied in the subject’s practical environment.² For instance, some philosophers have put forward the following condition: If a subject *S* knows that *p*, then it should not be a problem for *S* to act as if *p*.³

The idea that the truth-value of knowledge attributions is sensitive to practical factors rests in part upon pairs of cases that show that our ordinary knowledge attributions vary according to the stakes and the possibilities of error that are mentioned. One of these pairs of cases is the following DeRose (2009):

Bank Case A (Low Stakes). My wife and I are driving home on a Friday afternoon. We plan to stop at the bank on the way home to deposit our paychecks. But as we drive past the bank, we notice that the lines inside are very long, as they often are on Friday afternoons. Although we generally like to deposit our paychecks as soon as possible, it is not especially important in this case that they be deposited right away, so I suggest that we drive straight home and deposit our paychecks on Saturday morning. My wife says, “Maybe the bank won’t be open tomorrow. Lots of banks are

¹Cohen (1988), Lewis (1996), DeRose (2009).

²Hawthorne (2004), Stanley (2005), Hawthorne & Stanley (2008), Fantl & McGrath (2009).

³Fantl & McGrath (2002, 2007).

closed on Saturdays." I reply, "No, I know it'll be open. I was just there two weeks ago on Saturday. It's open until noon."

Bank Case B (High Stakes). My wife and I drive past the bank on a Friday afternoon, as in Case A, and notice the long lines. I again suggest that we deposit our paychecks on Saturday morning, explaining that I was at the bank on Saturday morning only two weeks ago and discovered that it was open until noon. But in this case, we have just written a very large and important check. If our paychecks are not deposited into our checking account before Monday morning, the important check we wrote will bounce, leaving us in a very bad situation. And, of course, the bank is not open on Sunday. My wife reminds me of these facts. She then says, "Banks do change their hours. Do you know that the bank will be open tomorrow?" Remaining as confident as I was before that the bank will be open then, still, I reply, "Well, no. I'd better go in and make sure."

In case A, Keith says that he knows, which seems true. In case B, he says that he does not know, which also seems true. The only difference between the two cases is that the stakes are high in case B (a large cheque has been written) and a different possibility of error has been raised (banks do change their hours).

At first glance, this is a problem for the intellectualist and invariantist account. Indeed, according to the intellectualist and invariantist view, since Keith is in the same epistemic position in both cases and the epistemic standards do not shift, he knows in case A if and only if he knows in case B. Most intellectualists are also non sceptics, and hence agree that Keith knows in case A.⁴

Thus, according to the intellectualist and invariantist view, it is true that Keith knows in case A and in case B. However, what Keith says in case B, namely that he does not know, seems true. If so, the intellectualist and invariantist account is in contradiction with our intuitions.

Warranted assertability manoeuvres

The main intellectualist and invariantist reply consists in questioning the claim that what Keith says in case B is true. According to this line of thought, in case

⁴Hereafter, I will put aside the sceptical approach to the problem.

B it is right for Keith to say that he does not know, and wrong for him to say that he knows. However, it can be right to say something false.

This line of argument is based on a distinction between the semantic and pragmatic content of an utterance. Whether or not the utterance of a sentence is right depends on two things: the truth-value of the proposition that is the semantic content of the sentence (the literal meaning of the sentence), and the proposition that is communicated by the utterance of the sentence.

Thanks to this distinction, the invariantist intellectualist can make a so-called warranted assertability manoeuvre in order to explain what happens in case B. It is true that Keith knows in case B. However, it is wrong for him to utter the sentence that has this proposition as semantic content in this case. Indeed, given the conversational context of case B, this utterance would communicate a proposition that is false. Furthermore, it is right in this context to say something false since this communicates something true.

To defend this warranted assertability manoeuvre, one needs to show that there is a general rule of conversation which is operative in case B, which explains that saying a truth ("I know") in this context prompts an implicature that is wrong in terms of conversational rationality, and that saying something false ("I do not know") prompts an implicature that is right in terms of conversational rationality. Indeed, the manoeuvre is *ad hoc* if there is no such a general rule at play.⁵

The most influential account along these lines is Rysiew's (2001, 2005). It has been modified by Brown (2006). That is why I will mainly focus on these views. The common basic idea is that what Keith says is right due to a conversational implicature based on the general maxim of Relation (be relevant!). According to this view, on the assumption that Keith makes a relevant speech act, he addresses the worry raised by his wife. Therefore, it is natural to interpret what he says in light of this worry in the following way. In uttering "I do not know that p", Keith implicates that he is not in a position to eliminate his wife's salient alternative, which is true. By contrast, in saying the literal truth, that he knows, he would implicate that he is in a position to eliminate this alternative, which is false.

Plan

I propose a new argument against this strategy, based on the idea that the rule of Relation cannot prompt the pragmatic implication required to explain the

⁵DeRose (2009, 88–9).

intuition that Keith says something right.

My plan will be as follows. I will first expose Rysiew's argument in further details. This will allow me to question one of his presuppositions: that the knowledge attribution or denial must *merely* take the salient alternative into account. By contrast, I will show that in order to be fully rational in conversational terms, Keith must take the general goal of the conversation into account: to decide whether or not he must come back the following day to deposit his pay cheque. Taking this general goal into consideration leads to oppose to the idea that uttering "I do not know that p" might be a relevant speech act in this context due to an implicature regarding the salient alternative. In this context, I will argue, the consideration of a salient possibilities of error is not conversationally *more* relevant than the question whether the subject knows.

It might be replied to my argument that high stakes can explain why it is relevant to consider this salient alternative. This point has been stressed by Brown (2006). I will argue that this reply fails. On the assumption that it is not common knowledge that Keith knows, the question whether he knows he still more relevant.

2. Rysiew's manoeuvre

Rysiew's warranted assertability manoeuvre

I will first present the background assumptions of Rysiew's manoeuvre. Then I will show how they are used to explain our intuitions by claiming that in Bank Case B, there is a true proposition that is pragmatically communicated by Keith's utterance. Then, I will put forward objections.

Background

Rysiew rests upon the idea that in general speakers do not make the distinction between (or are not interested in) the semantic (or literal) meaning and the pragmatic meaning. The semantic meaning of an utterance is the proposition that is literally expressed by this utterance. The pragmatic meaning is the proposition that is imparted, or communicated, by the act of uttering the sentence. For instance, let us suppose that John has three children. If I say "John has two children", what I say is literally true, but misleading. Indeed, due to the rule of Quantity (assert the stronger!), what is pragmatically communicated is that John only has two children, and therefore not three.

According to Rysiew, if there is a confusion between the semantic and pragmatic meaning, one can explain our intuitions about the alleged variation in truth-value in the bank cases. In case A, what Keith says is literally true, and what is pragmatically communicated is identical to what is literally said. In case B, what Keith says is literally false, but what is pragmatically communicated is different, and true. Therefore, the intuition that it is true that Keith does not know in case B relies on this confusion between the semantic and pragmatic meaning of Keith's utterance.⁶

Now, what does a knowledge attribution literally express? Rysiew adopts a relevant alternative theory about knowledge, even though his general view is compatible with other accounts. According to such a view, an utterance of "S knows that p" literally expresses the proposition that S has a true belief that p, and that S's epistemic position is good enough to eliminate all the alternatives to p that are relevant. An alternative is relevant in this epistemic sense if and only if it is knowledge-destroying.⁷

A possibility of error is knowledge-destroying if its likelihood is above a certain threshold. This threshold does not vary depending on conversational and practical factors (such as what is salient in the conversation, what is at stake, etc.) but is set independently. That is why his account is intellectualist and invariantist. The degree of likelihood for a possibility of error is set by what a normal human would think of it in the circumstances in question. Therefore, the set of knowledge-destroying alternatives for a proposition does not vary across contexts.

Rysiew distinguishes these knowledge-destroying possibilities from possibilities of error that are merely salient. A possibility of error is salient if it has been raised in the conversational context. A possibility of error is salient and not knowledge-destroying if and only if it has been raised in the conversational context but does not prevent a subject from knowing, even if he is not in a position to eliminate it. The set of salient alternatives is a variable set.

As regards pragmatic meaning, it is standardly assumed that what an utterance pragmatically imparts depends on the conversational context. Let us suppose that there is a feature, X, that can roughly specify the conversational context, and that is determined by the goal or direction of the conversation, what is salient, etc.⁸ Whether or not a speech act is right in conversational

⁶The idea that speakers are not always aware of, or interested in, the distinction between semantic content and pragmatic implicatures is disputable, though I will put aside this question.

⁷Rysiew uses the expression "relevant alternatives", but to avoid ambiguities, I will use "knowledge-destroying alternatives".

⁸One can think of X as the "conversational score". See Lewis (1979).

terms depends on X. If conversational participants are governed by a cooperative principle, they infer from what is literally said what best rationalizes others' speech acts given X. Therefore, what an utterance pragmatically imparts depends on X.

Given the variety of contexts, one can use an expression to communicate many different things. What is communicated in a given context depends on the semantic content of an utterance, along with the general rules of conversation. Then, according to Rysiew, it is natural to think that an utterance of "S knows" in some conversational contexts can pragmatically impart the proposition that S is in a good enough epistemic position to satisfy the epistemic standards for X. Indeed, in the semantic content of "S knows", there is the idea that S is in a good enough epistemic position. And if in some contexts this is what best rationalizes an utterance of "S knows", then this is part of what will be pragmatically communicated in this context.

The underlying idea is that X, the goal of the conversation, the interests of the speakers, etc., can determine some relevant epistemic standards in a context C. According to this view, an utterance of "S knows" in C can pragmatically impart that S meets these epistemic standards.

Bearing all these ideas in mind, it is possible to reconstruct Rysiew's warranted assertability manoeuvre.

Rysiew's account

Let me present Rysiew's account in detail. I will first explain what a knowledge attribution is supposed to communicate according to Rysiew, in a context where there are salient alternatives, if this attribution satisfies Grice's maxim of Relation (be relevant!). I will show that there are contentious assumptions in this picture. Then, I will focus on my main point.

An utterance is conversationally relevant in a context only if it takes the ingredients of the context into account. Thus, one can agree with Rysiew that a knowledge attribution in a context where some alternatives have been raised is conversationally relevant only if it takes the alternatives that have been raised in this context into account. Moreover, since a knowledge attribution communicates something about the subject's epistemic position, it is also very plausible that it can communicate something concerning the assessment of the subject's epistemic position with respect to the epistemic standards of the context. Rysiew uses both of these ideas to claim that a knowledge attribution in a context where there are salient alternatives can be conversationally

relevant if it communicates something concerning the assessment of the subject's epistemic position with respect to the salient alternatives.

In other words, if you apply the general idea that a conversationally relevant speech act must take the conversational context into account, to a case where the speech act is a knowledge attribution and the conversational context includes salient possibilities of error, you can predict that the knowledge attribution communicates something relative to these salient possibilities of error. And given the semantic content of a knowledge attribution sentence, it is natural to think that what is communicated is something related to the subject's epistemic position relative to these salient possibilities of error.

However, Rysiew goes further. He claims that it is natural to think that a knowledge attribution in this kind of context can be conversationally relevant in communicating that the subject's epistemic position with respect to the operative standards in the context is "good enough".⁹ How does Rysiew support this claim? Here is the explanation:

For it is only if speakers are understood to be intending to communicate information about how the subject fares vis-à-vis the contextually operative standard(s) — hence, i.a., about whether the subject can or cannot rule out any contextually salient alternatives — that they can be seen as striving to be maximally relevantly informative (hence, as conforming to CP) (Rysiew 2005, 48).¹⁰

It is not clear how the argument works. There is a difference between communicating information about the value of a subject's epistemic position with respect to an operative epistemic standard in a context (in Keith's case, a standard which is a function of salient doubts), and communicating that the subject's epistemic position is "good enough" with respect to the operative epistemic standards.

It may seem that Rysiew goes from one idea to the other in assuming that it is a semantic fact that if S knows that p, then S's epistemic position with respect to p is good enough. The idea would be that when you utter "S knows that p", you can communicate an assessment of S's epistemic position with

⁹See for instance Rysiew (2005, 48): "...it's because speakers strive to conform, and are known to so strive, to the maxim of relation ("Be relevant") that, in uttering a sentence of the form, 'S knows that p', the speaker is naturally taken to intend/mean that S's epistemic position with respect to p is 'good enough' given the epistemic standards that are operative in the context in question."

¹⁰See also (Rysiew 2001, 491–492).

respect to X , as well as the idea that the subject's epistemic position is good enough with respect to X .

Let us note that this is not obvious, in particular if one adopts Rysiew's fallibilist framework. If non sceptical fallibilist invariantism is true, the semantic fact is that if S knows that p , then S 's epistemic position with respect to p is good enough *to a certain extent*. Therefore, there is no clear reason to think that an utterance of " S knows that p " would impart the idea that the subject's epistemic position is *good enough* to satisfy the epistemic standards of the context. We will come back to this point below.

Let us also note that the use of the rule of Relation is based on the question of the specific relevance of an utterance. The question is not to recognize the falsity of what is literally said (in which case the conversational rule at work would be the rule of Quality). Therefore, the maxim of Relation can be put at work only if one assumes that the hearer is led to seek for a different (and pragmatic) meaning because the semantic meaning of the uttered sentence lacks specific relevance.

Then, in bank case B, the rule of Relation is operative only if the knowledge attribution or denial in this context lacks specific relevance. If one can show that the denial in case B does not lack specific relevance in this context, one can doubt that a warranted assertability manoeuvre using the rule of relevance is promising. This worry has been raised in particular by DeRose (2009). We will come back to it below.

For the time being, let us apply Rysiew's general idea to Keith's case. When Keith utters "I do not know", the semantic content of the sentence uttered in this context lacks specific relevance. (This would also have been the case if Keith had uttered "I know that the bank is open"). Indeed, this literally means: "I cannot eliminate the knowledge-destroying alternatives", and in the context, the question is about the salient alternative raised by Keith's wife. If the conversational context does not concern knowledge-destroying alternatives, but mere salient alternatives, one needs to find a meaning that is different from the literal meaning, so as to make the utterance specifically relevant (so as to make the speech act in conformity with the maxim of relevance).¹¹ This different meaning is the pragmatic meaning, namely that Keith cannot eliminate the salient alternative.

¹¹See for instance Rysiew (2001, 491–492): "Whether or not what he [Keith] says is strictly speaking true, if the speaker *didn't* think that S 's epistemic position were good enough *in the relevant sense of "good enough"*, whatever that is, he would not say " S knows that p "; and regardless of the truth-value of the sentence itself, he wouldn't say " S doesn't know that p " unless he thought there were salient $\text{no-}p$ possibilities that the speaker could not rule out".

Given this analysis, it is possible to rephrase Rysiew's account using the following claims:

- (1) In general, an utterance of "S knows that p" or "S does not know that p" is specifically relevant in a conversational context only if it means something concerning the assessment of the strength of S's epistemic position with respect to the standards required by X, where X is set by the conversational context C.
- (2) In general, an utterance of "S knows that p" in C, if it is specifically relevant, means that S's epistemic position is *good enough* with respect to the operative standards in C.
- (3) In general, an utterance of "S does not know that p" in C, if it is specifically relevant, means that S's epistemic position is *not good enough* with respect to the operative standards in C.
- (4) When C is the context of case B, X is the salient alternative.

Suppose also that:

- (5) Keith is not in a position to eliminate the salient alternative in case B.

Then, by (1) - (5):

- (6) In case B, Keith's utterance of "I know that p", if it means something that is specifically relevant with respect to the assessment of Keith's epistemic position concerning the salient alternative, would mean something that is false (namely that Keith can eliminate the salient alternative)
- (7) In case B, Keith's utterance of "I do not know that p", if it means something that is specifically relevant with respect to the assessment of Keith's epistemic position concerning the salient alternative, means something that is true (namely that Keith cannot eliminate the salient alternative)

Therefore, supposing that it is conversationally wrong to communicate something false and conversationally right to communicate something true:

- (8) It is conversationally wrong for Keith to say "I know that p" in case B.
- (9) It is conversationally right for Keith to say "I do not know that p" in case B.

Note that this account is supposed not only to explain why it would be improper for Keith to say that he knows, namely that this would prompt a false implicature, but also why it is correct to say that he does not know: this utterance takes the context into account, and communicates something true.

First possible objections

I will now consider the main objection that has been proposed so far, and that has been put forward by DeRose. It states that the semantic content of the knowledge denial is relevant in case B, and hence the hearer cannot be expected to infer a pragmatic meaning from a lack of relevance. In other words (1) can be true even though there is no pragmatic implicature at play in case B. I will argue that Rysiew's account can escape this objection. Then I will consider a different objection that seems more compelling to me. However, my main argument will be more radical.

According to DeRose (2009, 118–124), the semantic content of an utterance of "S does not know that p" does not lack specific relevance in the context of case B. Indeed, if this utterance literally means that S is not in a position to eliminate the knowledge-destroying alternatives, this directly implies that S is not in a position to eliminate the salient alternatives. It may seem that one needs to be in a better epistemic position to eliminate all the salient alternatives than to eliminate the knowledge-destroying ones only. Therefore, not only does this utterance seem specifically relevant, but it also communicates further information: that S is not in a position to eliminate the knowledge-destroying alternatives. Now, suppose that it is actually false that S does not know. Then, this utterance is actually misleading. Since the hearer cannot appeal to a lack of relevance for working out the implicature (since this utterance does not lack any specific relevance), and since the proposition expressed is false, the utterance violates the maxim of Quality. Therefore, the invariantist intellectualist cannot in this way explain why it seems correct.

A possible reply would be to say that it is common knowledge that S does not know. However, this is not a manoeuvre that uses the maxim of Relation, but the maxim of Quality. In addition, it is possible to figure out cases where this common knowledge is absent (DeRose 2009, 122–123).

Still, there is another possible rejoinder available to Rysiew. Indeed, the idea that the semantic content of an utterance is relevant is ambiguous. This can mean that the semantic content helps to reach the proposition that is meant to be communicated, or that it is sufficient to reach the proposition that is meant to be communicated. It does not seem that the knowledge denial in

case B is relevant in the latter sense. Let us assume that the important question is whether Keith can eliminate a salient alternative, and suppose that the semantic content of this knowledge denial does not lack relevance in the sense that it can “settle the question” concerning this alternative (DeRose 2009, 122). It might be argued that the way this semantic content can settle the question supposes a kind of inference that is led to the hearer. This inference may be of the following form: “If Keith does not know that the bank is open, then he is not in a good enough epistemic position to eliminate this salient possibility of error”. Clearly, if the hearer is not able to make this inference, he can be puzzled by the answer (given that the most relevant issue concerns a salient possibility of error). Therefore, the knowledge denial in case B is not sufficient to communicate what is meant to be communicated, and in this sense it is not fully relevant. That is why even if it can be used to settle the question, it can lack relevance.

To illustrate, let us take DeRose’s own analogy, and suppose that we are in a context where “tall” means “at least 6 feet tall”. Suppose that I am interested in whether Sally is tall. If I ask you “Is Sally tall?” and you answer “She is less than 5.5 feet tall”, we can agree that this settles the question in a negative way. In this sense your answer is not irrelevant. However, if I am not able to infer from the fact that Sally is less than 5.5 feet tall, the fact that she is less than 6 feet tall, then I cannot settle the issue whether she is tall. Then, it can be easily accepted that you expect me to make this inference from what you said. So, it can be granted that by saying that, you expect me (the hearer) to reach the relevant conclusion, that is, you aim at communicating this conclusion. (Your intention is to make me, the hearer, believe the conclusion). Therefore, it seems that what you intend to communicate is this conclusion, rather than just the semantic content of the utterance. In that sense, it can be granted that the semantic content is in some way relevant to reach the conclusion, but what is intended to be communicated is the conclusion itself. In this sense, the semantic content lacks relevance, which explains why I make the inference and reach the conclusion.

Nevertheless, pursuing DeRose’s line of thought, one might still wonder whether this explanation is satisfying. Indeed, the main problem is that, according to Rysiew, it is false that Keith knows in case B. And one might think that it is incorrect to say something false to make someone reach the right conclusion. To illustrate, let us follow DeRose and modify Sally’s case. Suppose that she is more than 5.5 feet tall, but less than 6 feet tall. Is it right to falsely say “She is less than 5.5 feet tall” to make the hearer reach the right conclusion that she is not tall? I must say that this does not seem wrong to me. There are

many cases in real life where one tells white lies, and this is acceptable insofar as it is efficient.¹² As a result, it does not seem to me that DeRose's criticism is sufficient to undermine Rysiew's attempt.

However, there are still a problematic assumptions in Rysiew's view at this stage, regarding (2) and (3) above. Recall that it is not clear why the semantic meaning of a knowledge attribution would drive the hearer toward the idea that he is in a good enough epistemic position, rather than an insufficient epistemic position, to meet the epistemic standards of the context. This assumption is needed if Rysiew wants to explain why a knowledge attribution would be inappropriate, while true, in case B. Suppose for instance that Keith says "I know", and by this, literally means that his epistemic position is good enough to eliminate the knowledge-destroying alternative. Contrary to what Rysiew claims, this could drive the hearer to the right piece of information. Indeed, this may pragmatically impart that Keith is not in a position to eliminate the salient alternative. If Keith does not say "I know for sure", but merely "I know", given the rule of Quantity (assert the stronger!), the hearer may infer that Keith's epistemic position is not good enough to eliminate the salient but non knowledge-destroying alternative. The underlying idea is that one cannot just suppose, like Rysiew does, that the pragmatic meaning of an utterance such as "S knows that p" is in general that S's epistemic position with respect to p is good enough for X, where X is fixed by the context.

Rysiew seems to defend his view against this possible objection with the idea that one can base a pragmatic inference upon a mere part of the semantic content. Even if one does not have a clear view of what the truth-conditions of a knowledge attribution are (i.e. being in a good enough epistemic position to eliminate the knowledge-destroying alternatives), one can base the inference upon a part of these truth-conditions (i.e. being in a good enough epistemic position with respect to the proposition).

Thus, Rysiew writes:

All you need is to see that if I didn't think that S did so measure up, it would be odd, indeed uncooperative, of me to say, "S knows that p"; for whatever exactly knowledge is, we know that S's knowing

¹²See Fantl & McGrath (2009, 41–42) for discussion. They insist that it is difficult to explain why a white lie is acceptable in case B. But it may seem that a white lie is efficient. They also raise another worry concerning the first-person perspective. According to them, in case B Keith will be willing to say to himself that he does not know. This cannot be explained easily by a WAM, even assuming that one easily confuses the semantic and pragmatic meanings. However, given Rysiew's fallibilist framework, one might argue that since Keith knows, he will rather be willing to say to himself: "I know that the bank is open but I should not take the risk".

that *p* entails that *S* is in a good epistemic position with respect to *p*. “And why”, the hearer may well ask, “would he say something that means [semantically implies] that if he didn’t in fact think that *S* measured up to the epistemic standards that are in play, and so was able to put to rest any doubts or rule out any not-*p* possibilities that had just been raised?” (Rysiew 2005, 49).

There are two possible replies to this defence. On the one hand, if it is true that “*S* knows that *p*” semantically implies that *S*’s epistemic position is good enough in an absolute sense, then it is true that if *S* knows that *p* then *S*’s epistemic position is good enough no matter the context. Under this interpretation, uttering “*S* knows that *p*” will semantically express that *S*’s epistemic position is good enough for any *X*, where *X* is determined by the context. However, this cannot be what Rysiew means, because then Keith would be in a position to eliminate his wife’s salient alternatives in case B.

On the other hand if Rysiew wants to mean that “*S* knows that *p*” semantically implies that *S*’s epistemic position with respect to *p* is good enough *to a certain extent*, then he cannot escape the previous objection. If in a context the standards are higher than those required for knowledge, to say that someone meets these knowledge-level standards might well communicate that he does not meet those required for more than knowledge.

More generally, it does not seem that “being good enough” is a semantic part of “being good enough to a certain extent”. Indeed, it does not follow from “*X* is good enough to a certain extent” that “*X* is good enough”. Therefore, it does not seem that it is possible to claim that since “*S* knows that *p*” semantically imparts that “*S* is in a good enough epistemic position with respect to *p* to know *p*” then “*S* knows that *p*” semantically imparts that “*S* is in a good enough epistemic position with respect to *p*”. As a result, (2) does not seem true.

It can also be argued in a similar way that (3) may be false. By uttering “I do not know that *p*” in a context with higher epistemic standards, I may convey the idea that I do not *merely* know that *p*. This case is similar to the one where one says that one does not believe *p* in order to convey that one knows *p*.

To recap, the idea that a knowledge attribution (or denial) lack relevance in case B is not sufficient to explain why the hearer is driven to the idea that the speaker is (or is not) in a good enough epistemic position with respect to the salient alternatives. This is one aspect in which Rysiew’s account may seem insufficient.

However, my main argument will focus on (4).

My strategy against Rysiew's warranted assertability manoeuvre

In case B, the salient possibility of error is conversationally relevant because it has been raised in the conversation. One can explain why it has been raised by the fact that the stakes are high. Indeed, following Rysiew, one might think that when the stakes are high, it is natural to think about possibilities of error. Nevertheless, Rysiew's WAM does not essentially appeal to the stakes in order to explain why the knowledge denial is warranted in case B. Indeed, there might be cases where the same manoeuvre is available, even though the possibility of error is raised in a different way.¹³ Brown (2006) has put forward a slightly different view: the presence of high stakes in case B can make a difference as to the extent to which the salient possibility of error is conversationally relevant.

In what follows, I will distinguish the two approaches. First I will focus on Rysiew's account, putting aside the issue of the stakes. I will come back to the second proposal in section 4.

I will develop my criticism of Rysiew with the idea that the conversational relevance of a knowledge attribution in a context must take into account two features:

1. One must take the salient possibilities of error in C
2. One must take these salient possibilities of error in function of the general goal of the conversation in C.¹⁴

In addition, it seems clear that an implicature cannot violate a conversational rule since it results from the need to conform to these rules. Hence, if a proposition is not fully relevant in C, that cannot be an implicature of what is said in C.

To my mind, the problem in Rysiew's explanation is that it assumes that the only way to take the salient alternative into account consists in assessing Keith's epistemic position with respect to it. However, with regard to the

¹³For instance, Hazlett (2009) uses this manoeuvre to account for knowledge denials in sceptical contexts where there are no stakes.

¹⁴See Grice (quoted by Rysiew 2005, 50): "I expect a partner's contribution to be appropriate to the immediate needs at each stage of the transaction" (Grice 1989, 28). This does not imply to deny the general goal of the conversation.

general goal of the conversation, if it is true that Keith knows, it does not seem relevant for him to say that he does not know. Here is my argument.

Let us suppose with Rysiew that if Keith says that he does not know, this pragmatically prompts the true proposition q , where q is that Keith cannot eliminate his wife's salient alternative. Let us also suppose that if Keith says that he knows, this pragmatically prompts the false proposition $\text{not-}q$. Rysiew deduces that since q is true and $\text{not-}q$ is false, Keith must utter what prompts q .

However, I will argue that the question whether q or $\text{not-}q$ lacks of relevance in case B. In a rational conversation, a proposition that lacks of relevance in the context cannot be pragmatically implicated by an utterance in this context. As a result, neither q nor $\text{not-}q$ can be implicatures in this context. If so, whether or not Keith is in a good enough epistemic position to eliminate his wife's alternative cannot explain why Keith's utterance "I do not know that the bank is open" is right, if it is true that he knows.

My main task will be to defend the idea according to which the question whether Keith is in a position to eliminate the salient alternative of his wife lacks relevance given the general course of the conversation. I will use two premises. First, I will claim that the case B is a case where the goal is to take a rational decision. Second, I will claim that in a conversational context where the goal is to take a rational decision, the question whether one is in a position to eliminate a salient but not knowledge-destroying alternative lacks relevance.

It seems to me that the second premise is the most subject to criticism. So I will mainly defend it. I will do it in the following way. First I will show that in general, when the goal of a conversation is to know whether a subject knows, the question whether one can eliminate a salient but not knowledge-destroying alternative is irrelevant. Second, I will argue that if the question has to do with the rationality of a subject acting on a proposition, the question whether the subject knows this proposition is more relevant than the question whether he can eliminate a not knowledge-destroying alternative.

3. Ways of taking the salient possibilities of error that are not knowledge-destroying into account

Context where the question is whether a subject knows a proposition

In a context where what matters is whether or not a subject knows, it is not appropriate to say that he does not know, if he knows, even if this subject is facing a salient possibility of error that is not knowledge-destroying. Indeed, in this context, the question does not turn around the possibility for the subject to eliminate alternatives that are not knowledge-destroying, but concerns only possibility that are knowledge-destroying.

It is clear that it would be wrong for a speaker to say that the subject does not know in such cases semantically and pragmatically. From the semantic point of view, the utterance "S does not know" is literally false. From the pragmatic point of view, it seems to communicate something false: that the salient possibility of error is knowledge-destroying. Therefore, the main way of interpreting such an utterance in this context would be to think that the speaker is mistaken. He thinks that a not knowledge-destroying possibility of error is knowledge-destroying, which is not. So, this utterance is incorrect.

Thus, suppose that in such contexts someone raise a possibility of error which is not knowledge-destroying. It is clear that it is not appropriate to take it into account by saying "S does not know that p". Indeed, given the purpose of the conversation, raising this possibility was not relevant in the first place, and it would be irrational (in conversational terms) to take this possibility into account by communicating, by the utterance "S does not know", that S cannot eliminate it. Indeed, this would communicate that this possibility is knowledge-destroying and that it was right, in the first place, to consider this possibility in order to contemplate whether S knows.

Therefore, uttering "I do not know that p" in front of a salient possibility of error that is not knowledge-destroying, in a context where the point of the conversation is to consider whether I know, seems equivalent to utter "I do not know that p" in front of an epistemically and conversationally unreasonable alternative.

The upshot is simple. In a context where the question is to know whether S knows that p, it is not appropriate to take not-eliminable salient possibilities of errors that are not knowledge-destroying into account by saying "S does not know that p". On the contrary, it is appropriate to take them into account by rejecting them as not relevant.

There are at least two ways of doing it. Rysiew thinks that concessive knowledge attributions (sentences such as "I know that p but it might be that q" (where q obviously entails not-p)) are not contradictory. But then, this would be an appropriate answer. Alternatively, one can directly reject this alternative as irrelevant.

Context where the question is to decide what to do

I have argued that in a context where the question is whether a subject knows, not knowledge-destroying possibilities of error are not conversationally relevant. If one can also show that in the bank case B, the question whether Keith knows is more relevant than the question whether he can eliminate not knowledge-destroying alternatives, then the result will be that there is no implicature to the effect that Keith cannot eliminate his wife's possibility of error (assuming that it is not knowledge-destroying).

I will argue that in general, when the goal of the conversation is to decide whether one should act on a proposition, the question whether one knows this proposition is more conversationally relevant than the question whether one can eliminate a possibility of error that is not knowledge-destroying. (This is a reason to think that in raising a possibility of error, Keith's wife challenges Keith's knowledge.)

First of all, there is no denying that knowledge claims are conversationally relevant when the purpose of the conversation is to decide what to do. For instance, discussing a case similar to the bank cases, Reed (2010, 232) writes:

[Y]ou are trying to decide whether to check if the train stops in Foxboro because it is extremely important that you get there as quickly as possible. You have not yet decided whether it is rational to check if the train makes that stop. One of the relevant factors in your decision would presumably be an answer to the question, do you know the train will stop there?

Similarly, it is very plausible that in the bank cases, one factor that is relevant to decide what to do is an answer to the question: does Keith know that the bank is open on Saturdays? Now, the question is how relevant is this factor. Is an answer to this question more relevant than an answer to the following question: is Keith in a position to eliminate not knowledge-destroying alternatives?

If knowledge questions were sometimes not relevant, then one should expect cases where the relevance of a knowledge claim is appropriately chal-

lenged. However, it seems that in most cases, whether one knows the target proposition settle the issue in a proper way (I will claim that this is also true if the stakes are high). For instance suppose that it is right to come back tomorrow only if the bank is open. Suppose that I say that we should come back tomorrow because I know that the bank is open. One may question whether I really know, but it does not seem appropriate to question the relevance of my contribution to the discussion. For instance, the following reply would be surprising: "Well, the question is not whether you know that the bank is open". It is highly plausible that whether one knows is always a relevant question in such a kind of case. On this assumption, it would not be surprising if Keith's wife challenged Keith's knowledge.

By contrast, it does not seem that in such contexts the question of eliminating not knowledge-destroying possibilities of error is relevant in the same way. Suppose for instance that I say to my wife: "I prefer to go to the bank tomorrow. I know it's open". Suppose that she answers: "But maybe the bank has changed its hours". It is not conversationally irrelevant for me to reply: "Of course, everything is possible, but do you really think that this will happen? What makes you think that?". I ask for a reason to consider this possibility of error, which shows that it is not obvious to me whether raising this possibility of error is relevant to the discussion. In other words, I challenge the relevance of what she has just said.

Thus, the main difference between the knowledge question and the not knowledge-destroying alternative question is the following: in general, the relevance of the first question cannot be easily challenged, while the relevance of the second one can be easily challenged. Does this show that the first question is more relevant? I think that it does, if one also notes that a challenge to a not knowledge-destroying possibility is appropriate precisely if this possibility of error is not obviously not knowledge-destroying nor obviously knowledge-destroying.

To see this, let us compare our reactions to alternatives that are clearly not knowledge-destroying. Suppose for instance that my wife raises the possibility that Martians will destroy the bank tonight. An appropriate answer would be: "Come on! You know this won't happen". If it is obviously not knowledge-destroying, the alternative can be rejected outright. This suggests that it is appropriate to challenge the relevance of a salient possibility of error only if it might be knowledge-destroying.

It is then natural to think that challenging the relevance of an alternative aims at deciding whether or not this alternative is knowledge-destroying, or to what extent it might be knowledge-destroying. Indeed, asking for a rea-

son to think that this alternative might obtain is just asking whether it is knowledge-destroying. After all, if there is a good reason to think that an alternative to *p* might obtain, one's knowledge that *p* is destroyed (unless one can eliminate it). And in cases where there is no good reason to think that this possibility might obtain, one's knowledge is not destroyed. These cases are precisely those in which these alternatives can be rejected as irrelevant for the purposes at hand.

To sum up, our appropriate reactions to alternatives that are raised in contexts where the question is to decide what to do seem in general governed by how these alternatives are supposed to be related to what we know. This shows that in these contexts, the conversational relevance of the question concerning possibilities of error is derived from the conversational relevance of the question concerning knowledge. As a result, the most relevant question in these contexts seems to be the knowledge question.

Applying the analysis to the bank cases

Let us apply this analysis to the bank cases. In the bank cases, what matters is to take a rational decision: either to deposit the pay cheque right now or to come back on Saturday. Even if Keith cannot eliminate his wife's possibility of error, it is also assumed that this salient possibility is not knowledge-destroying. Therefore, according to the line of thought that I have put forward, this possibility is not the most relevant to take the right decision. That does not necessarily mean that it is totally irrelevant. Although this alternative is not knowledge-destroying, it may be relevant in the sense that it is not obviously not knowledge-destroying. But in that case, Keith should challenge it. On the contrary, if one assumes that this possibility of error is obviously not-knowledge destroying, then Keith should reject it outright as irrelevant. As a result, since the fact that Keith cannot eliminate this possibility is not fully relevant (or is totally irrelevant) to the purposes at hand, Keith denial cannot pragmatically impart the proposition that he cannot eliminate this possibility.

Note that there is another independent reason to think that the case B is a case where the question is whether Keith knows. It seems clear that in raising this possibility of error, his wife is interested in whether he knows. She does not ask whether Keith can eliminate this possibility, but whether he knows that the bank is open. This suggests that she thinks that the possibility of error she has raised challenges Keith's knowledge. If so, it would be misleading for Keith to answer negatively just to communicate that he cannot elimi-

nate this alternative. Indeed, he would thereby impart that this alternative is knowledge-destroying, or would be mistaken about what his wife meant.

4. Brown's variation

One possible reply is to say that given the high stakes in case B, whether or not Keith can eliminate his wife's possibility of error is actually totally relevant to decide what to do.

I will argue that even if the stakes are high, this possibility of error is less relevant than whether Keith knows. Then, even if the stakes are high, on the assumption that it is not common knowledge that Keith knows, it is more plausible that Keith's wife raises this alternative to challenge Keith's knowledge.

Let us first develop the idea that this salient alternative has a particular relevance due to the practical situation. Brown (2006) has proposed to fill out Rysiew's view in order to explain why and how the fact the Keith's wife has mentioned a possibility of error can change the goal or the direction of the conversation, so that an utterance of "know" means "can eliminate the salient alternatives" and "not know" means "cannot eliminate the salient alternatives". According to her, the practical importance can explain this. Indeed, it seems that one can more easily resist to possibilities of error raised in contexts where it is not practically important not to be wrong, than to possibilities of error raised in contexts where it is practically important. And a relevant conversational ingredient in Keith's conversational context is the fact that Keith's wife has mentioned that a large cheque has been written.

If this analysis is correct, given that Keith's utterance is relevant, he takes this conversational aspect into account, and one must interpret his claim in light of this consideration. The fact that a large cheque has been written "makes it clear that what's relevant to the conversation is a very strong epistemic position" (Brown 2006). Therefore, our argument does not seem to work against this slightly modified WAM.

However, there are some possible troubles for this rejoinder. First, this leads to reject the idea that if you know that *p*, then your epistemic position with respect to *p* is good enough to act on *p*. Still, one might argue that this claim is false (Brown 2008).

Second, and more importantly, Brown's filling out rests on the idea that it is more difficult to resist the mentioning of not knowledge-destroying possibilities of error when it is important not to be wrong because when the stakes

are high, one needs to be in a strong epistemic position with respect to *p* to be rational to act on *p*.

One can grant this idea, but that does not imply that in case B Keith should impart something relatively to the salient alternative. Indeed, recall that it is assumed that it is not common knowledge that Keith knows. And it does not seem that if Keith can eliminate some specific not knowledge-destroying possibility in which not-*p*, he is in a stronger epistemic position than if he knows *p*. Then, even if a strong epistemic position is required, the more interesting issue in order to take the right decision is still whether he knows that the bank is open, rather than the question whether he is in a position to eliminate the possibility that the bank has changed its hours.

To put in another way, suppose that eliminating a specific not knowledge-destroying possibility of error in which not-*p* does not put you in a better epistemic position with respect to *p* than knowing *p*. Suppose also that the question is to take the right decision and the best epistemic position is required. In this case, what is conversationally relevant, given the general goal of the conversation, is to focus on what one knows rather than on one's epistemic position with respect to some specific far-fetched possibility of error.

I think that one can easily accept that the question whether one knows is more important. Suppose that *p* is relevant to your action, and you have a choice between two options: know whether *p* or eliminate a not knowledge-destroying possibility of error in which not-*p*. It seems clear that you will choose the first option. For instance, consider the following case:

You are a policeman in a country where there is only one gun, and you are about to enter into a house to arrest some dangerous criminal. It is very important that the criminal is not armed. Before entering, you just have enough time to read one of the two following messages on your mobile: (a) is the criminal armed? or (b) does the criminal have a gun?

It seems clear that it is rational to choose (a) if you have to decide what to do. Indeed, suppose that you choose (b) and the answer is positive. Then it is not rational for you to go into the house. Suppose that you choose (b) and the answer is negative. Still, it is not rational for you to enter into the house because the criminal might be armed in a different way. An answer to (b) does not make any difference to what you should rationally do. On the contrary, suppose that you choose (a). If the answer is positive, then it is not rational for you to go into the house. But if the answer is negative, it does not seem problematic for you to enter.

One might object that it might not be rational to enter into the house, if one has not eliminated the possibility that the subject has a gun. Two remarks are in order. First it is important to note that even in high stakes situations, saying that one knows the target proposition is in general considered as sufficient to settle the issue. In high stakes cases, one might be less willing to cite knowledge, but that does not show that citing knowledge does not settle the question.¹⁵ Second, even if some might doubt that it is rational to enter into the house given that there is no answer to (b), it is clear that (a) is more relevant than (b).

Similarly in the bank cases. Suppose that Keith's wife does not know whether Keith knows. Does she raise this possibility of error to challenge Keith's knowledge, or Keith's ability to eliminate this salient alternative? It is implausible that she raises this possibility of error to challenge Keith's ability to eliminate it. Indeed, whether or not the bank has changed its hours is not fully relevant. Maybe the bank has not changed its hours, but actually it is closed on Saturdays except when Keith came the last time. Maybe it has changed its hours but it is still open tomorrow. On the contrary, if she raises this possibility to challenge Keith's knowledge, that seems fully relevant insofar as this possibility of error might be knowledge-destroying.

One can sum up by claiming that knowledge-destroying possibilities of error are more relevant to rational action than salient but not knowledge-destroying possibilities of error even when the stakes are high. Given that they are more relevant, if Keith knows, and on the assumption that it is not common knowledge that Keith knows, Keith should not impart that he cannot eliminate a less important possibility of error but should challenge his wife's alternative, or should say that he knows (if he knows), even if the stakes are very high.

5. Conclusion

The traditional invariantist and intellectualist account of knowledge is challenged by pairs of cases with asymmetrical stakes and different salient possibilities of error. These cases show that the correctness of knowledge attributions varies with the presence or absence of these practical factors. Still, it has been argued that these intuitions does not show that knowledge varies with practical factors, but only that the warranted assertability conditions of knowledge vary. Assertions, and knowledge attributions in particular, prompts

¹⁵See Fantl & McGrath (2007, 562).

conversational implicatures. However, this mechanism can be used to explain our intuition only if a general rule of conversation is operative in the context. The conversational rule that is in general put forward is the rule of relation. I have shown that this rule cannot be used to explain our intuitions about these cases.

6. References

- Brown, J. (2006). Contextualism and warranted assertability manoeuvres. *Philosophical Studies*, 130:407–35.
- Brown, J. (2008). Subject-sensitive invariantism and the knowledge norm for practical reasoning. *Noûs*, 42 (2):167–189.
- Cohen, S. (1988). How to be a fallibilist. *Philosophical Perspectives*, 2:91–123.
- DeRose, K. (2009). *The Case for Contextualism*. Oxford University Press.
- Fantl, J. and McGrath, M. (2002). Evidence, pragmatic and justification. *The Philosophical Review*, 111:67–94.
- Fantl, J. and McGrath, M. (2007). On pragmatic encroachment on epistemology. *Philosophy and Phenomenological Research*, 75:558–89.
- Fantl, J. and McGrath, M. (2009). *Knowledge in an uncertain world*. Oxford University Press.
- Grice, H. P. (1989). *Studies in the Way of Words*. Harvard University Press.
- Hawthorne, J. (2004). *Knowledge and Lottery*. Oxford University Press.
- Hawthorne, J. and Stanley, J. (2008). Knowledge and action. *Journal of Philosophy*, 105 (10):571–590.
- Hazlett, A. (2009). Knowledge and conversation. *Philosophy and Phenomenological Research*, 78 (3):592–620.
- Lewis, D. (1979). Scorekeeping in a language game. *Journal of Philosophical Logic*, 8:339–59.
- Lewis, D. (1996). Elusive knowledge. *Australian Journal of Philosophy*, 74:549–67.
- Reed, B. (2010). A defense of stable invariantism. *Noûs*, 44 (2):224–244.
- Rysiew, P. (2001). The context-sensitivity of knowledge attributions. *Noûs*, 35 (4):477–514.

- Rysiew, P. (2005). Contesting contextualism. *Grazer Philosophische Studien*, 69 (1):51–70.
- Stanley, J. (2005). *Knowledge and Practical Interests*. Oxford University Press.