

A Remark on the Bank Cases

Abstract: Since their formulation by Keith DeRose (1992), the so-called bank cases have played a major role in the discussion about whether knowledge depends on practical factors. According to the proponents of pragmatic encroachment, the proper conclusion to be drawn from the bank cases and similar examples is that knowledge of a proposition p does not supervene on one's evidence for or against p . In my view, this conclusion is ill-founded. The reason is that the bank cases and similar examples suffer from an ambiguity concerning the known proposition — an ambiguity that has so far been overlooked. When this ambiguity is made explicit, it becomes clear that the conclusion does not follow.

Since their formulation by Keith DeRose (1992), the so-called bank cases have played a major role in the discussion about whether knowledge depends on practical factors.

According to the proponents of pragmatic encroachment, the proper conclusion to be drawn from the bank cases and similar examples is that knowledge of a proposition p does not supervene on one's evidence for or against p . In my view, this conclusion is ill-founded. The reason is that the bank cases and similar examples suffer from an ambiguity concerning the known proposition — an ambiguity that has so far been overlooked. When this ambiguity is made explicit, it becomes clear that the conclusion does not follow.

First, let us review the bank cases as they are presented by Jason Stanley (2005, 3-4):

Low Stakes. Hannah and her wife Sarah are driving home on a Friday afternoon. They plan to stop at the bank on the way home to deposit their paychecks. It is not important that they do so, as they have no impending bills. But as they drive past the bank, they notice that the lines inside are very long, as they often are on Friday afternoons. Realizing that it wasn't very important that their paychecks are deposited right away, Hannah says, "I know the bank will be open tomorrow, since I was there just two weeks ago on Saturday morning. So we can deposit our paychecks tomorrow morning".

High Stakes. Hannah and her wife Sarah are driving home on a Friday afternoon. They plan to stop at the bank on the way home to deposit their paychecks. Since they have an impending bill coming due, and very little in their account, it is very important that they deposit their paychecks by Saturday. Hannah notes that she was at the bank two weeks before on a Saturday morning, and it was open. But, as Sarah points out, banks do change their hours. Hannah says, “I guess you’re right. I don’t know that the bank will be open tomorrow”.

It strikes many of us as plausible that although her evidential situation does not vary from Low Stakes to High Stakes, Hannah in Low Stakes *knows* that the bank will be open tomorrow, while she does *not* know that in High Stakes. According to the proponents of pragmatic encroachment, the proper lesson to be drawn from these observations is that it is possible that subjects *S* and *S'* have the same evidence for *p*, and yet *S* knows that *p* while *S'* does not.

Now, notice that in order to be licensed to draw this conclusion it has to be assumed that *the proposition in question does not vary from case to case*. However, that’s exactly the crux of the matter.

In the bank cases, the proposition in question is expressed by Hannah’s utterance “The bank will be open tomorrow”. According to the standard semantics of indexicals, a token of “tomorrow” refers to the day after the day on which the token in question is used by the speaker. Suppose that Hannah utters the sentence in question on Friday, May 15 2020. Then the referent of the “tomorrow”-token in question is Saturday, May 16 2020.

Let us paraphrase the resulting proposition as

[1] The bank will be open on Saturday, May 16.

Now, recall that there are also non-standard, so-called “descriptive” uses of indexicals.

Geoffrey Nunberg (1993; 2002) gives an especially clear example:

Condemned prisoner: “I am traditionally allowed to order whatever I like for my last meal.”

In this case, the pronoun ‘I’ “doesn't refer to the speaker, since obviously there couldn't be any traditions that deal specifically with his last meal; instead it refers to the role he exemplifies” (Nunberg 2002, 266). Thus, the condemned prisoner’s utterance may be paraphrased as “Whoever is in the same situation as I am now, namely a condemned prisoner, is traditionally allowed to order whatever they like for their last meal.”

Similarly, for “tomorrow” (cf. Nunberg 1993):

Drawn from an article in the University of Arizona newspaper that appeared on the Friday before the beginning of classes: “Tomorrow is always the biggest party night of the year.”

Suppose that the sentence in question appeared in the University of Arizona newspaper on Friday, May 15 2020. Then the “tomorrow”-token cannot simply be taken as a device that directly refers to Saturday, May 16 2020. Rather, the “tomorrow”-token functions as a complex description that may be paraphrased as “any Saturday before classes begin”.

Now, if we read Hannah’s use of “tomorrow” in her utterance “The bank will be open tomorrow” in a descriptive sense (and taking the appropriate contextual factors into account), the proposition expressed amounts to something like this:

[2] The bank is usually open on Saturdays.

So, my diagnosis for the bank cases is as follows. Confronted with Low Stakes, we, the reader of the story, interpret Hannah’s utterance in the sense of [2], whereas confronted with High Stakes, we interpret Hannah’s utterance along the lines of [1]. The reason is

that in Low Stakes, nothing depends on the bank being open on a *specific* day. Thus, our interpretation of Hannah's utterance "relaxes", so to speak, allowing the possibility that the bank is closed on May 16. In contrast, in High Stakes everything depends on the bank being open on May 16, and so our interpretation of Hannah's utterance thus focuses on that specific day.

Now, if two different propositions, [1] and [2], are at play in the bank cases, the difference in knowledge can easily be explained without any reference to pragmatic features. Hannah's evidence (which is, *ex hypothesi*, the same in High and Low Stakes) is sufficiently strong to warrant the acceptance of [2], but too weak to warrant the acceptance of [1]. You may know on May 15 that the bank usually opens on Saturdays on the basis of recalling that you were there just two weeks ago on a Saturday morning; but you certainly cannot know on May 15 whether the bank will definitely be open on May 16 on that basis. Thus, the bank cases do not show that two subjects who are the same with respect to their evidence supporting proposition p can differ with respect to their knowledge of p . Rather, the bank cases merely show that a particular piece of evidence may suffice for knowledge of p , while the same piece of evidence does not suffice for knowledge of a different proposition q — which is pretty trivial.

The proponents of pragmatic encroachment might object that the ambiguity diagnosed is simply an artefact of the specific design of the bank cases. One might easily invent other examples supporting the case of pragmatic encroachment which do without sentences involving ambiguous indexicals.

However, this objection misses my point, which is not that the examples supposedly supporting the case of pragmatic encroachment suffer from an ambiguity concerning *indexicals*, but rather that they suffer from an ambiguity concerning *sentences*. Thus, the strategy I propose is not to defuse those examples by disambiguating the meaning of indexicals, but by disambiguating the sentence meanings. In order to get an idea how this strategy might work in other cases, let us take a look at an example taken from Jeremy Fantl and Matthew McGrath (2002):

Jeremy is at Back Bay Station in Boston preparing to take the commuter rail to Providence for vacation. He asks a man, “Does this train make all those little stops in Foxboro, Attleboro, etc.?” It does not matter much to Jeremy whether the train is express or not. The man answers, “Yea, this one makes all those little stops. That's what I was told when I bought the ticket.” Jeremy believes what he says.

It is of dire importance that Matt gets to Foxboro and the sooner the better. While he has a ticket that gets to Foxboro in two hours, which is just in the nick of time, a train rolls into the station and he overhears the conversation above. Matt thinks to himself, “That guy may have misheard. After all, he doesn't care so he probably didn't pay careful attention. I better go check it out myself.”

Again, it strikes many of us as plausible that although their evidential situation is the same, Jeremy *knows* that the incoming train will stop at Foxboro while Matt doesn't. According to the proponents of pragmatic encroachment, the knowledge difference is due to a variation of pragmatic factors: while the stakes for Jeremy are low, they are high for Matt. According to the strategy I propose, the knowledge difference is due to the fact that Jeremy and Matt are related to different propositions: while in Jeremy's case, the proposition in question might be paraphrased as

[3] It is more or less likely that the incoming train will stop at Foxboro,

the proposition in Matt's case is as follows:

[4] It is beyond any reasonable doubt that the incoming train will stop at Foxboro.

Confronted with Jeremy's case, it is natural to interpret the sentence "The train will stop at Foxboro" in the sense of [3], whereas confronted with Matt's case, it is natural to interpret it along the lines of [4]. The reason is that in Jeremy's case, nothing depends on whether or not the incoming train stops at Foxboro. Thus, we can permit ourselves a "lax" interpretation that allows for the possibility that the train won't stop at Foxboro. In Matt's case, however, everything depends on the train stopping at Foxboro. Thus, our interpretation is "strengthened" such that the possibility that the train won't stop at Foxboro is almost ruled out. Again, the difference in knowledge can easily be explained without any reference to pragmatic features. Jeremy's evidence (which is, *ex hypothesi*, the same as Matt's) is sufficiently strong to warrant the acceptance of [3], but too weak to warrant the acceptance of [4]. You may know that it is more or less likely that the incoming train will stop at Foxboro on the basis of a stranger's testimony; but you certainly cannot know that it is beyond any reasonable doubt that the incoming train will stop at Foxboro on that basis.

Regarding my treatment of the train cases, one might object that in contrast to the bank cases, it is unclear where the two different readings [3] and [4] are supposed to be coming from. Unlike the sentence "The bank will be open tomorrow", the sentence "The incoming train will stop at Foxboro" does not contain a word or phrase to which we can attribute different meanings.¹ It therefore looks as if the epistemic operators "it is more or less likely that" and "it is beyond any reasonable doubt that" come from nowhere. The sentence simply contains no linguistic material from which we can derive these

¹ For simplicity, I have ignored the context-sensitivity of "the incoming train".

operators. So it seems that [3] and [4] are quite arbitrary re-interpretations of a proposition which is perfectly fine in itself and does not require any enrichment, namely:

[5] The incoming train will stop at Foxboro.

The only way to avoid the impression of arbitrariness, it seems, is to claim that the sentence “The incoming train will stop at Foxboro”, though syntactically well-formed, is semantically incomplete, rather like “John is tall”: just as the sentence “John is tall” does not express a complete proposition unless a reference class is specified, the sentence “The incoming train will stop at Foxboro” does not express a complete proposition unless the subjective probability of the described fact occurring is specified.

Consequently, [5] is not a full-fledged proposition, but merely a propositional stub, so to speak. However, this is implausible, to put it mildly.

To overrule this objection, recall that there are at least two types of sentential semantic indeterminacy, one of which results from the *absence* of conceptual components required to achieve full-fledged propositionality, and the other from the *abundance* of available interpretations of a sentence that is already complete in its propositional regards. A perfect example of the first type of indeterminacy is “John is tall”: unless a reference class is specified, it does not express a full-fledged proposition. The second type of indeterminacy is nicely illustrated by “You’re not gonna die” (Bach 1994, 134). The sentence “You’re not gonna die” is perfectly fine in propositional regards (given that the referent of “you” is determined). We don’t need to replenish it in order to achieve full-fledged propositionality — for it already expresses a full-fledged proposition.

Nevertheless, it depends on the overall situation exactly *which* proposition is expressed.

To illustrate, consider two short stories:

Zeus has accidentally chopped his head and starts crying terribly. Hera comes by and consoles him: “Hey, it’s not so bad. You’re not gonna die. After all, you’re a god.”

Five year old John has accidentally cut his little finger with a knife and starts crying terribly. His mother comes by and consoles him: “Hey, it’s not so bad. You’re not gonna die. After all, you’re a brave boy.”

Confronted with Zeus’ case, it is natural to interpret “You’re not gonna die” in the sense of

[6] You are immortal.

In contrast, confronted with John’s case, it is natural to interpret “You’re not gonna die” along the lines of

[7] You’re not gonna die from that little cut.

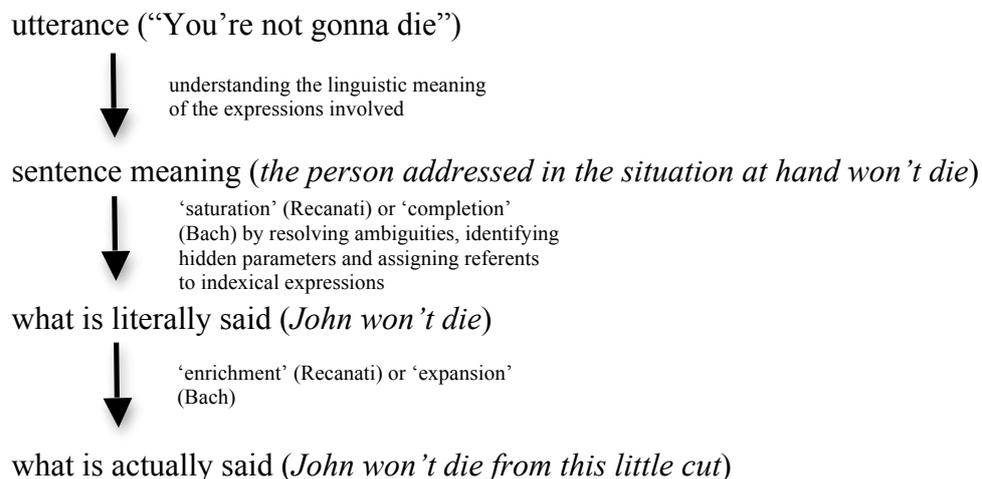
What is happening in John’s case is a phenomenon which goes under various labels in the literature: Bach (1994) calls it ‘expansion’ while Recanati (2003) speaks of ‘enrichment.’ The basic idea is that even if all the operations necessary for achieving full-fledged propositionality (such as resolving ambiguities, identifying hidden parameters, specifying reference classes, and assigning referents to indexical expressions) have been made, what is actually said by an utterance in a given situation is still underdetermined. A further step is necessary.

To be clear about what exactly is happening here, let us look at John’s case in semantic slow motion, so to speak. It is helpful to distinguish between three steps. First, we deal with the uttered sentence containing an indexical expression, “You’re not gonna die.”

This sentence has a conventional meaning that is constant across all utterance contexts,

and might be approximated by *the person addressed in the situation at hand won't die*. Following Kaplan (1989), this might be called the sentence's *character*. However, the sentence's character does not provide us with a complete proposition, a truth-evaluable content, until the referent of the word 'you,' i.e., the person addressed in the situation at hand, is specified. This leads us to the second step. Since we gather from the story that the person addressed in the situation at hand is *John*, we assign John as the referent of the word 'you.' In a certain sense, we have no choice in making this assignment, it is dictated to us by the conventional meaning of the word 'you.' Due to its conventional meaning, the word 'you' clings to the person being addressed by the speaker in the situation at hand, whoever that person may be. In this way, the non-truth-evaluable meaning of the uttered sentence is transformed into a proposition which might be approximated by *John won't die*. This, one might say, is what the mother's utterance literally says. However, it is obviously *not* what she actually says. In the situation described, the mother does not tell her son that he is immortal. Thus, the transformation of the non-truth-evaluable meaning of the sentence "You're not gonna die" into a truth-evaluable proposition by assigning a referent to the expression 'you' is still not sufficient to arrive at what the mother actually says. We have to take one more step, step three. At this step, we are no longer bound by the dictates of the conventional meanings of the expressions occurring in the sentence uttered. It would not help, for example, if we searched the context for more referents which the expressions occurring in the uttered sentence cling to by virtue of their conventional meaning. For except 'you,' no other expressions occur in the uttered sentence to which we could still assign referents. Moreover, it would not help to disambiguate the predicate "x won't die" or to identify a hidden parameter — for the predicate is neither ambiguous nor in need of identifying a hidden parameter. Semantically, everything is fine with the predicate: "x won't die" is true exactly if there is no time t such that x dies at t , period. But nonetheless: the

proposition resulting from the replacement of the predicate's free variable by 'John' does not correspond to what the mother actually says to her son. Thus, we have to depart from what the uttered sentence means in the literal sense and look for something that does not seem to be suggested by the conventional meaning of the expressions involved. We have to instead look at the overall situation from a more general perspective and "freely" consider what the mother might want to communicate with her utterance. That is, we need to take a perspective that does not consider the situation merely as a collection of items from which indexical expressions are assigned their referents. Rather, we need to look at the situation as a meaningful whole, taking into account what the people involved might be thinking and what their intentions might be in light of what has happened, etc. In this way, we arrive at the proposition actually expressed, which might be approximately paraphrased as *John won't die from this little cut*. The three-step procedure described may be simplified as follows:



The philosophically interesting point is that the last step, the transition from what is literally said to what is actually said, seems to be an "inference to the best interpretation" (Atlas & Levinson 1981; Atlas 2005), which, while loosely related to the literal meaning of the utterance, may radically override it. (Thus, Recanati (2003) aptly speaks of 'free

enrichment.’) That the inference is only loosely related to the literal meaning of the utterance and may even override it in a radical way does not, of course, mean that the inference is not guided by principles, but which ones are involved is a matter of controversy. However, the most promising approach is to draw on certain Gricean conversational maxims (or variants thereof). It seems, for example, tempting to say that the transition from what is literally to what is actually said is guided by a general maxim along the lines of “expand what is literally said by finding the most specific interpretation available that fits what you think is the intended point of the speaker” (cf. Levinson 2000, 114-115).

Now, my proposal is that the mechanism described, i.e. expansion/enrichment by an inference to the best interpretation, is also at work in the train cases. We are told a little story about Jeremy and Matt that ends with the narrator asking us if the two characters know that the train will stop at Foxboro. The important factor — the factor that triggers expansion/enrichment — is that we are told that Jeremy and Matt have radically different informational needs: for Jeremy, it doesn’t matter if the train stops at Foxboro, but for Matt it is of paramount importance. From Matt’s perspective, the possibility of the train skipping Foxboro should better be ruled out, while from Jeremy’s perspective, this is not mandatory. So when we are eventually asked to judge whether Jeremy and Matt, respectively, know the train will stop at Foxboro, we understand this question differently, i.e. relativized to Jeremy’s and Matt’s informational needs.

To avoid any misunderstanding, let me add that I am not claiming that Jeremy and Matt understand the man’s utterance (“Yes, this one makes all those little stops”) differently. I do not make any claim about how the story’s characters understand some utterance made within the story. Rather, I am claiming that *we*, the story’s readers, take the content

clause within the narrator's question differently, depending on whether it focuses on Jeremy or Matt: with respect to Jeremy, we interpret the content clause as expressing the "lax" proposition *that it is more or less likely that the train will stop at Foxboro*; with respect to Matt, we interpret the content clause as expressing the "strengthened" proposition *that it is beyond any reasonable doubt that the train will stop at Foxboro*.

Similar to the case considered earlier, these different readings are not the result of resolving ambiguity, identifying hidden parameters, and/or contextually assigning referents to indexical expressions, but rather the conclusions of inferences to the best interpretation of what the narrator's question aims at. Thus, as in the case of the mother comforting her son, it is a case of applying the general maxim of finding the most specific interpretation available that fits what we think is the intended point of the speaker. Since we cannot help but assume, based on what we are told, that Matt needs to be in an epistemic position that allows him to rule out the possibility of the train not stopping at Foxboro, we adjust our understanding of the Matt-directed question to fit Matt's informational needs (which are, as one might say, "acute"). We understand the narrator's question as asking us whether Matt's informational need has been satisfied, that is, whether Matt can rule out the possibility of the train not stopping at Foxboro — for what else could the intended point of the question be? Similarly, for the Jeremy-directed question: since we gather from the story that for him it doesn't matter whether the train stops at Foxboro or not, we adjust our understanding to fit Jeremy's informational needs (which are, in contrast to Matt's, only "moderate"). With these adjustments, we eventually satisfy the maxim that requires us to find the most specific interpretation of the narrator's question.

In a certain sense then, the bank, the train, and other related cases are a kind of philosophical self-deception. Philosophers put a spoke in their own wheel in the way they tell their story, so to speak. By distributing the informational needs of the story's protagonists utterly unequally, philosophers trigger an interpretive process in their readers' minds that inevitably leads to the content clause of the narrator's question — “Does *S* know that *p*?” — being understood differently, depending on the informational needs of the story's protagonists. The alleged evidence for pragmatic encroachment, which lies in the fact that we answer the question differently for subjects whose evidential situation is the same, is thereby nullified. Our different answers do not show that it is possible that subjects *S* and *S'* have the same evidence for *p* and yet *S* knows *p* while *S'* does not. Rather, our different answers merely show that it is possible that subject *S* has the same evidence for *p* that subject *S'* has for *q* (where $p \neq q$), and yet *S* knows *p* while *S'* does not know *q*. This is, however, a rather trivial insight that is insignificant in philosophical terms.

Although I've only dealt with two examples here, it is nevertheless fair to say that the discussion so far makes it seem very likely that *any* example in support of pragmatic encroachment could be defused using the strategy proposed in this paper. This is because all the stories that appear in the literature as evidence of pragmatic encroachment operate with a large imbalance in terms of the protagonists' informational needs.

Note that my strategy also avoids contextualism along the lines of DeRose (2009), i.e. the thesis that the truth-conditions of knowledge attributions vary with pragmatic factors. The variations in meaning diagnosed in this paper are not concerned with the verb “to know,” but rather with the sentences that express what is (putatively) known. According to the approach taken in this paper, it is neither knowledge itself nor the meaning of the

verb “to know” that vary with pragmatic factors. Rather, it is the meaning of the content clause within the question “Does *S* know that *p*?” This variation in meaning, however, is not only found in questions regarding a person’s knowledge. Rather, it is the result of a general interpretative mechanism which is effective almost everywhere in natural language understanding, so we should not be surprised at the variations in meaning that it brings about. At least, epistemology need not be rewritten to accommodate this phenomenon.

My previous reflections had a conservative thrust: knowledge, I concluded, does not depend on practical factors. The impression that this might be the case is due to the particular design of the examples, which triggers different processes of expansion/enrichment of the putatively known proposition, depending on the respective informational needs of the protagonists. Accordingly, the thesis that the pragmatic encroacher wishes to undermine remains untouched: it is impossible for two persons *S* and *S*’ to be in possession of the same evidence for *p*, and yet *S* knows *p* while *S*’ does not. One might argue, however, that this conservatism comes at a high price — for it seems that we have to give up the traditional approach to natural language semantics. According to the traditional approach, the truth-condition of an utterance depends solely on its syntactic structure, the conventional meaning of the involved expressions, and the context of use (in the meagre Kaplanian sense of a collection of items from which the referents of indexical expressions are selected). If what I have just said in the preceding passages is correct, then the truth-condition of an utterance seems to depend on much more: on how the situation evolves, what the people acting in the situation might be thinking, what their communicative intentions are, etc. It therefore looks as if I am simply shifting the buck: from epistemology to the philosophy of language. The defence

against pragmatic encroachment with respect to knowledge thus only succeeds at the price of pragmatic intrusion with respect to semantics, or so it seems.

I think that this objection is rooted in a misunderstanding, however. First, allow me to point out that the fact that I have used terms like ‘expansion,’ ‘enrichment,’ and ‘inference to the best interpretation’ in the foregoing does not mean that I endorse the semantic theories of the people who coined those terms. I have used those terms merely to describe a particular phenomenon that occurs in the course of interpreting utterances, but I did not intend to imply that I agree with everything that Atlas, Bach, Levinson, and Recanati say. Second, and more importantly, the three-step process I described above is a reconstruction of *what interpreters do* when they try to figure out what a speaker is saying by uttering a sentence on a given occasion. It describes a sequence of psychological processes by which the hearer decodes what is said by a given utterance.² In contrast, the traditional approach to natural language semantics does not provide a description of what interpreters do when they interpret an utterance. It does not provide a description of psychological processes. Rather, it is the formulation of a mathematical procedure that takes certain abstract entities, pairs of semantic representations and contexts of use (in the Kaplanian sense) as input, and gives certain other abstract entities, functions from possible worlds to truth-values, as output.³ Note that semantic

² An anonymous reviewer objected that this is not quite correct. The three-step process described earlier is rather a *rational* reconstruction, i.e. a reconstruction that does not describe a blind causal event, but makes clear why what happens in the mind of an utterance interpreter is rationally justified. I concede that, but I do not see how that could affect my intended point. For even a *rational* reconstruction of a psychological process is a reconstruction of a *psychological process*, and is thus different from the formulation of a mathematical procedure that takes certain abstract entities as input and yields certain other abstract entities as output.

³ Here and in the following passages, I have relied heavily on Predelli (2005). However, my formulations differ slightly from Predelli’s. Where he uses the term ‘t-distributions,’ I use the more customary

representations are not the same as utterances. Rather, they are formulas that are completely unambiguous in syntactical terms such as bracketed structures or labelled trees which are used in linguistics. Of course, this does not mean that the mathematical procedure cannot be applied to utterances. On the contrary: on the assumption that the pair of semantic representation r and context c , $\langle r, c \rangle$, is an appropriate representation of an utterance u , the mathematical procedure can be used to assign to u the very function from possible worlds to truth-values which it gives as output for $\langle r, c \rangle$. The crux of the matter, however, is that the mathematical procedure does not tell us *which* pair of semantic representation and context is an appropriate representation of a given utterance. This is something that lies beyond the scope of the mathematical procedure. It would therefore be a mistake to think that the phenomenon of expansion/enrichment poses a threat to the traditional approach to natural language semantics.

To make this last point as clear as possible, let us revisit John's case once again. The mathematical procedure does not give us an answer to the simple question "What does the mother actually say by uttering the sentence 'You're not gonna die'?" Because as previously stated, the mathematical procedure does not take utterances as inputs. However, the mathematical procedure does give us an answer to the conditional question "What does the mother actually say by uttering the sentence 'You're not gonna die', assuming that a particular semantic representation and context pair is an appropriate representation of her utterance?" Thus, in order to apply the mathematical procedure to the mother's utterance, we must first have already decided which semantic representation and context pair is an appropriate representation of it. This decision cannot be brought about by the mathematical procedure itself — it is something that must be brought about

'functions from possible worlds to truth-values.' Similarly, I prefer to use 'semantic representation' and 'context of use (in the Kaplanian sense)' instead of 'clause' and 'index.'

before the mathematical procedure is put into operation. This is precisely where the three-step process described above comes into play: the decision which semantic representation and context pair is an appropriate representation of the mother's utterance is the result of an interpretive process that begins with understanding the literal meaning of the words uttered, continues with completion/saturation, and finally culminates in expansion/enrichment. Thus, from the viewpoint of the mathematical procedure, expansion/enrichment is located on the same level as other preparatory measures such as the resolution of ambiguities or unpacking of ellipses. Like those other preparatory measures, they play no role in assigning a function from possible worlds to truth-values to the mother's utterance. This assignment is done by the mathematical procedure quite automatically — without taking the overall situation or the speaker's beliefs and intentions into consideration — once it has been decided which semantic representation and context pair is an appropriate representation of the mother's utterance. In short, the phenomenon of expansion/enrichment and the traditional approach to natural language semantics can peacefully coexist. The former describes a psychological process that occurs in the course of utterance understanding, the latter a mathematical procedure that operates on abstract entities. Thus, the objection raised above is unfounded. I am not shifting the buck from epistemology to the philosophy of language. My defence against pragmatic encroachment with respect to knowledge does not entail pragmatic intrusion with respect to semantics.⁴

⁴ Thanks to Andy Müller and the participants of my graduate seminar at Goethe University for helpful discussions. Thanks are also due to an anonymous reviewer of this journal who drew my attention to interesting points that I had not seen before.

References

- Atlas, J. D. (2005), *Logic, Meaning, and Conversation*, Oxford University Press.
- Atlas, J. D. & Levinson, J. (1981), "It-Clefts, Informativeness and Logical Form: Radical Pragmatics," in: P. Cole (ed.), *Radical pragmatics*, Academic Press.
- Bach, K. (1994), "Conversational Implicature," *Mind & Language* 9, 124-162.
- DeRose, K. (1992), "Contextualism and Knowledge Attributions," *Philosophy and Phenomenological Research* 52, 913-929.
- DeRose, K. (2009), *The Case for Contextualism*, Oxford University Press.
- Fantl, J., and McGrath, M. (2002), "Evidence, pragmatics, and justification," *The Philosophical Review* 111, 67-94.
- Kaplan, D. (1989), "Demonstratives," in: J. Almog, J. Perry, and H. Wettstein (eds.), *Themes from Kaplan*, Oxford University Press.
- Levinson, S. C. (2000), *Presumptive Meanings. The Theory of Generalized Conversational Implicature*, MIT Press.
- Nunberg, G. (1993), "Indexicality and Deixis," *Linguistics and Philosophy* 16, 1-43.

Nunberg, G. (2002), “Descriptive Indexicals and Indexical Descriptions,” in M. Reimer and A. Bezuidenhout (eds.), *Descriptions and Beyond: An Interdisciplinary Collection of Essays on Definite and Indefinite Descriptions and Other Related Phenomena*, Oxford University Press.

Predelli, S. (2005), *Contexts. Meaning, Truth, and the Use of Language*, Oxford University Press.

Recanati, F. (2003), *Literal Meaning*, MIT Press.

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