Protagonist Projection

Andreas Stokke
andreas.stokke@gmail.com

- forthcoming in Mind & Language -

Abstract: This paper provides a semantic analysis of Protagonist Projection, the phenomenon by which things are described from a point of view different from that of the speaker. Against what has been argued by some, the account vindicates the intuitive idea that Protagonist Projection does not give rise to counterexamples to factivity, and similar plausible principles. A pragmatics is sketched that explains the attitude attributions generated by Protagonist Projection. Further, the phenomenon is compared to Free Indirect Discourse, and the proposed account is seen to preserve the relation between them.

1. Introduction

In 1977 Lauri Karttunen (1977) made a famous observation. The observation was that what Partee (1973a) had called *verbs of communication*, examples of which include *tell*, *show*, *indicate*, *inform*, *disclose*, behave differently depending on whether they take a that-complement or an indirect question. Here is his example:

(1) a. John told Mary that Bill and Susan passed the test.
b. John told Mary who passed the test.

Karttunen noted that (1b) implies that (the speaker believes that) John told Mary the truth, whereas (1a) does not. More generally, verbs of communication imply truth-telling when they take an indirect question, but do not when they take a that-complement. Call this *Karttunen's Thesis*.

Almost twenty years on, Karttunen's Thesis was challenged by Tsohatzidis (1993) who pointed to cases in which the thesis appears to be violated, as in these examples:

(2) a. John told the voters what he intended to do for them once elected, but, as usual, he was lying to them.
b. Old John told us whom he saw in the fog, but it turned out that he was mistaken.

I would like to thank Nicholas Allott, Pål Antonsen, Herman Cappelen, Josh Dever, Guðmundur Andri Hjalmarsson, Ole Hjortland, Torfinn Huvenes, Dilip Ninan, Philippe Schlenker, Anders Schoubye, and François Recanati for very valuable and extensive discussion. This paper was presented at the workshop *Perspective in Discourse* at the Centre for Advanced Studies, Norwegian Academy of Science and Letters, in March 2011. I would like to thank the audience, and in particular, Emar Maier, Arnim von Stechow, Kjell Johan Sæbø, and Henk Zeevat for discussion at this event. Thanks are also due to two anonymous reviewers for *Mind and Language* for helpful suggestions.

Address for correspondence: CSMN, Postbox 1020, Blindern, 0315 Oslo, Norway.
Email: andreas.stokke@ifikk.uio.no

c. John told them where he had been between 4 and 5 p.m., but he was certainly lying since nobody was at the place he said during that time.

For instance, since (2a) clearly does not imply that (the speaker believes that) John told the voters the truth, (2a) looks like a counterexample to Karttunen's Thesis.

However, as Holton (1997) subsequently pointed out, there is a different option. Holton argued that Tsohatzidis's examples involve what he called *Protagonist Projection*, which he characterized as a technique whereby things are described using words that the protagonist of the story would use, although they do not necessarily correspond to what the speaker herself believes.[1] For example in (2a), things are described from the point of view of the voters who falsely believed that John was telling the truth, and it is this reorientation that licenses the wh-construction.

Protagonist Projection, Holton observed, is quite widespread. Here are some of his further examples:[2]

(3) a. He gave her a ring studded with diamonds, but they turned out to be glass.
b. She knew that he would never let her down, but, like all the others, he did.
c. I saw a shooting star last night. I wished on it, but it was just a satellite.
d. She sold him a pig in a bag. When he got home he discovered it was really a cat.

All of these sentences are perfectly acceptable, although they too have a 'literary' feel to them. According to Holton, they are just more examples of Protagonist Projection, and hence if one endorses Tsohatzidis's argument, then one should also endorse the analogous arguments with respect to all of these cases. For example, one should conclude from (3a) that some diamonds are made of glass, and from (3b) that *know* is not factive, etc. But since this is clearly an undesirable strategy, Holton's line looks more attractive.

Many have thought that Holton's observation was a good one. But recently Hazlett (2010) has claimed that, although valid for cases like (3a), the argument does not work for examples involving factive verbs, such as (3b). This claim is a key premise in Hazlett's case for the controversial conclusion that *there are no* factive verbs.

Neither Holton nor Hazlett is explicit about what is going on in the examples, however. So, to evaluate these arguments, we need an understanding of the relevant aspects of the style of discourse they involve. This paper tries to provide the basics of such an understanding.

I begin in Section 2 by responding to Hazlett's arguments. I then go on in Sections 3 and 4 to present an analysis of Protagonist Projection, which maintains Holton's overall strategy.

[1] Note that by the label 'Protagonist Projection' we are not alluding to other phenomena also sometimes referred to with the term 'projection', e.g., the way presuppositions of compound sentences are computed from those of their parts. I choose to retain the label to indicate continuity with Holton's project.

[2] See Holton (1997) for sources. An anonymous reviewer points out that Protagonist Projection can also arise from relative clauses, as in 'He gave her a ring studded with diamonds that turned out to be glass'.

2. Are there Counterexamples to Factivity?

As mentioned above, Hazlett agrees with Holton regarding the examples not involving factive verbs. So, for instance, he accepts (as seems undeniable) that (3a) is not a counterexample to the thesis that no diamonds are made of glass. Yet, he argues that examples like (3b) *are* counterexamples to factivity. As Hazlett recognizes, for this claim to be tenable, it needs to be shown that there is a sufficient degree of relevant difference between the two kinds of examples. In this section I refute Hazlett's arguments to the effect that there is such a difference.

2.1 Hazlett on Factive Verbs

Here is how Hazlett (2010, 499) formulates the thesis of factivity:

> **Factivity**
> Certain two-place predicates, including 'knows', 'learns', 'remembers', and 'realizes', which denote relations between persons and propositions, are *factive* in this sense: an utterance of 'S knows p' is true only if p, an utterance of 'S learned p' is true only if p, and so on.

Hazlett argues that examples such as Holton's (3b) falsify Factivity. Here are the cases he focuses on:[3]

(4) a. Everyone knew that stress caused ulcers, before two Australian doctors in the early 80s proved that ulcers are actually caused by bacterial infection.[4]
b. He figures anything big enough to sink the ship they're going to see in time to turn. But the ship's too big, with too small a rudder ... it can't corner worth shit. Everything he knows is wrong.
c. In school we learned that World War I was a war to 'make the world safe for democracy,' when it was really a war to make the world safe for the Western imperial powers.
d. I had trouble breathing, sharp pains in my side, several broken ribs and a partially collapsed lung, and I was in the middle of nowhere without any real rescue assets - it was then that I realized I was going to die out there.

Noting examples like these, Hazlett asks what a traditional theory of these expressions, that is, a theory that accepts Factivity, can say about them. He concedes that the 'most promising route' is Holton's, and he grants that 'The idea seems basically right about the case of the 'glass diamonds'.' (2010, 526) But he then argues that Holton's argument does not apply to the cases involving (allegedly) factive verbs. So Hazlett concludes that the traditional theory has no way of accommodating these

[3] See Hazlett (2010, 516, fn. 13-16) for sources.
[4] It might be argued that for this example to work in the way Hazlett wants it to, it needs to be assumed that *prove* is a factive verb. But for the purpose of discussion, we ignore this complication.

examples, and hence they should be taken as refuting that theory; that is, they are counterexamples to Factivity.

The important point to note is that, on Hazlett's view, sentences like the diamond sentence (3a) are *false*. And for him, this is the reason they are not counterexamples. E.g., (3a) is not a counterexample to the thesis that no diamonds are made of glass. But on the other hand, he thinks that the sentences involving factive verbs are all *true*, and therefore, they are counterexamples to Factivity. In other words, he assumes that, for all the examples, *if* they are true, they are counterexamples.

Let φ be a sentence of the relevant kind, e.g., from (2)-(4). And let *T* be the relevant thesis - e.g., Factivity, Karttunen's Thesis, the thesis that no diamonds are made of glass, etc. We can then spell out Hazlett's two assumptions as follows:

H1. φ entails not-*T*.
H2. Therefore, if φ is true, φ is a counterexample to *T*.

The notion of entailment in H1 is the familiar one:

φ entails ψ =df If φ is true, ψ is true.

And we are using 'counterexample' here in the following sense:

φ is a counterexample to *T* iff φ entails not-*T* and φ is true.

So, according to Hazlett, since (3a) is false, it does not motivate the conclusion that some diamonds are made of glass; but since (3b) is true, (3b) is a counterexample to Factivity.

In contrast to this, the account I will present in this paper rejects H1, and consequently also rejects H2. In other words, on my account, whether the cases are counterexamples or not does not rest on whether they are true or false. Even if they are true, they are not counterexamples. So I will not be insisting, as Hazlett does, that the cases that are obviously not counterexamples, such as the diamond case, are false.

To motivate his own line, Hazlett needs to show that there are relevant differences between these kinds of cases and those not involving factive verbs. Hazlett argues for three such differences. I respond to each in turn below.

2.1.1 Difference in Responses? Hazlett's first claim is that

> the speaker who says 'He gave her a ring studded with diamonds, but they turned out to be glass' will respond negatively to the question of whether the ring was really studded with diamonds. It's not clear how the ordinary person who utters [(4a) or (4b)] would respond to such a question (e.g., of whether people really knew stress caused ulcers). Non-factive uses of the other 'factive verbs' don't resemble the case of 'glass diamonds' in this respect. And someone who uttered [(4c)], for instance, would most likely respond *affirmatively* to the question of whether the falsehoods were 'really' learned. (Hazlett, 2010, p. 516)

Although it is not entirely clear what the argument is that Hazlett has in mind at this point, I think the most charitable way of construing it is as follows. Hazlett assumes that if it can be shown that speakers will respond to the *really*-questions in the way he claims they will, this will be evidence for the conclusion that examples like the diamond case are false while the examples involving factive verbs are true.

The first thing to say in response is that it is not clear that if the data goes the way Hazlett thinks it does, this will indeed motivate the conclusion he wants to draw. That is, even if it is found that speakers respond negatively to the question of whether he really gave her a ring studded with diamonds, it is not obvious that this is evidence that the original sentence is false.

For one thing, it will depend crucially on how *really* works in the relevant questions. Although I will not go into the details of this, it might be argued that *really* serves to shift significant parameters of the context in a way that blocks Hazlett's conclusion. Here is an analogy. Suppose A and B are talking about their friends, and are comparing who is tall and who is not. At one point, A says,

(5) A: John is tall.

This utterance may be obviously true - suppose that John is clearly tall by normal standards. But now suppose that A and B are also sometimes interested in who is tall by Guiness Book of Records standards. So B may feel like asking,

(6) B: Is John *really* tall?

To which A will be compelled to reply negatively. But of course, we would not conclude on this basis that (5) is false. Rather, it is clear that *really* here indicates the context-shift that the felicity of B's question relies on.

I will not undertake an investigation of whether the analogous point could be justified with respect to the dialogues Hazlett is imagining. I bring this out in order to suggest that even if the data patterns the way Hazlett expects, it will require further argument before one can take it as motivating his conclusion.

The more important observation, and the reason for not pursuing the former, is that it is far from clear that the data actually does pattern the way Hazlett expects it to. In particular, I think it is clear that for the cases involving *know*, a positive response to a *really*-question will be significantly less felicitous than a negative response:

(7) A: Everyone knew that stress caused ulcers, before two Australian doctors in the early 80s proved that ulcers are actually caused by bacterial infection.
B: Wait, did people *really* know that?
A: No, not really, of course - they just thought they knew that./??Yes they did - they were sure of it.

It is clear that similar dialogues are available for the other cases involving *know*, although they may not play out in a completely parallel fashion.

For example, take Hazlett's (4b). I think it is fair to say that the use of *know* that this case exemplifies is less smooth than the sort we have just looked at. But I do not want to focus on this issue, which might very well turn out to be explainable in some independent way. The point to note is that there are easily imaginable situations

in which speakers will respond negatively to a *really*-question of the relevant sort. As in this dialogue:

(8) A: He figures anything big enough to sink the ship they're going to see in time to turn. But the ship's too big, with too small a rudder ... it can't corner worth shit. Everything he knows is wrong.
B: Wait, you mean he *really* knew that they would spot anything big enough to sink them in time?
C: No, of course not - he just thought he knew that./??Yes, he was sure of it.

It seems, then, that there is little ground for claiming that speakers will respond differently in the cases involving *know*.

What about (4c)? I agree with Hazlett's intuition that an affirmative answer is very likely to be the response with respect to this case. For example, the following seems to be a plausible exchange:

(9) A: In school we learned that World War I was a war to 'make the world safe for democracy,' when it was really a war to make the world safe for the Western imperial powers.
B: Wait, did you *really* learn that?
A: Yeah, we did - those hypocrites!

One possible conclusion to draw from this is that Hazlett is right about *learn*, although he is not right about *know*. That is, one may want to concede that *learn* has non-factive uses. I am not unsympathetic to this conclusion. However, given that there are reasons to doubt that these dialogues can be taken as providing the relevant evidence at all, I do not think that this conclusion is warranted by data like (9).

So, with respect to Hazlett's first point of disanalogy, I conclude that it has not been shown that the cases involving *know* are significantly different from cases like the diamond example, although there may be some cause for doubt regarding other factive verbs. For that reason, I will focus exclusively on *know* when giving my own account of Protagonist Projection. But we must now turn to Hazlett's two further arguments for disanalogy.

2.1.2 Deliberateness Here is Hazlett's second argument:

> It is crucial to Holton's account that the false utterances be *deliberate* - i.e. the speaker must know that her utterance is false. [...] Speakers who properly use the sentence 'He gave her a ring studded with diamonds, but they turned out to be glass' believe that no diamonds are made of glass - i.e. that 'He gave her a ring studded with diamonds' is false. But it is not plausible - at least not *as clearly plausible* as in the case of the 'glass diamonds' - to suppose that speakers who properly use sentences like [(4a)] and [(4b)] believe that nothing false can be known, or that 'Everyone knew that stress caused ulcers' is false. (Hazlett, 2010, pp. 516-517)

As is clear, this argument relies on the assumption that examples like the diamond sentence are false. As I said, I will reject this assumption, and hence this argument can be regarded as irrelevant.

However, a few things are important to note here. First, I take it that everyone agrees that there is a sense in which the examples are 'deliberate'. Namely, they are marked, literary, stylized, or whichever characterization one prefers. As such, it is clear that speakers who engage in Protagonist Projection do so deliberately. They deliberately present things from the point of view of someone else. But one does not have to accept that the examples are false for this to be so.

Second, one can accept deliberateness in this sense without being committed to the claim that speakers who use (4a) or (4b) believe (or know) that nothing false can be known. So I agree with Hazlett that

> It is compatible with the fact that 'Nothing false can be known' is not obvious to most people that nothing false can be known. People can be wrong about such things. (Hazlett, 2010, p. 517)

Further, we can accept that, in this sense, there is a difference from the diamond case in that it is more likely that speakers know that no diamonds are made of glass than that they know that nothing false can be known. The latter is a philosophical thesis, the former is everyday knowledge. But this makes no difference for deliberateness, construed in the way I suggested.

2.1.3 Independent Evidence Finally, then, here is Hazlett's third observation:

> we have independent scientific reason to believe that no diamonds are made of glass. It's on the basis of this prior knowledge that we conclude that 'He gave her a ring studded with diamonds, but they turned out to be glass' must be false. [...] I don't think we have anything like scientific grounds on which to believe that nothing false can be known. (Hazlett, 2010, p. 517)

As with the former point, this argument relies on the premise that the examples are false. Again, I will reject this premise, and therefore this argument is irrelevant.

We may point out, however, that the argument Hazlett gives here concerns a difference in the kind of evidence we have for theses like the thesis that no diamonds are made of glass vs. the thesis of Factivity. But, even if there is a such a difference, it is hard to see how this could be relevant to the question of whether the examples on the table are true or false. Surely, whether a claim of the kind in question is true or false does not depend on what kind of evidence speakers and hearers have for believing it.

3. How to Handle Protagonist Projection

We have seen that Hazlett fails to motivate that there is a difference in the examples. In the rest of this paper, therefore, I will pursue the task of providing an analysis of Protagonist Projection that preserves Holton's conclusion for cases like the diamond

sentence, for the cases involving *know*, as well as for the putative counterexamples to Karttunen's Thesis. In this section, I first spell out what I take to be the chief points that need to be explained, and I then outline my overall strategy for doing so.

3.1 What Needs to be Explained

What are the desiderata of an account of the examples on the table? The first is that a suitable account needs to preserve Holton's conclusion that the examples are not counterexamples to such attractive theses as Karttunen's Thesis, Factivity, the thesis that no diamonds are made of glass, etc. The second is to explain how the effect of Protagonist Projection arises.

But what more precisely is this effect? Holton wrote:

> I suggest that these sentences work by projecting us into the point of view of the protagonist [...]. In each case the point of view into which we are projected involves a false belief. We describe the false belief using words that the protagonists might use themselves, words that embody their mistake. (Holton, 1997, p. 626)

Following this, a slightly more precise conjecture is the following: A sentence using Protagonist Projection conveys an *attribution* of a belief, which is easily seen to be false. For example, the thought would be that (3a) conveys that someone falsely believed that the ring was studded with diamonds. The context would then be required to supply the subject of this attitude. It might be the giver of the ring who believed they were real diamonds, or it might be the receiver, or it might be the speaker, or some combination of the above.

However, this conjecture is in fact too narrow, as can be seen from an observation made by Recanati (2010). He considers the following example:

> (10) *John to Bill*: Ok, I am stupid and I don't understand the matter. Why do you ask me for advice, then?

I take it to be a plausible suggestion that this is an example of what we have followed Holton in calling Protagonist Projection. John's utterance of 'I am stupid and I don't understand the matter' parallels our other examples in the relevant respect of expressing a point of view different from the speaker's. As Recanati says,

> What the sentence expresses is something that John puts in the mouth of the addressee, Bill. It is Bill who is supposed to think or say that John is stupid and does not understand the matter; and his so saying or thinking provides John with a reason for asking the question in the second part of the utterance: 'Why do you ask me for advice, then?' (Recanati, 2010, p. 198)

Recanati points out that it would be wrong to describe this phenomenon narrowly as that of conveying a belief-attribution. The example is consistent with a reading on which John's utterance does not attribute a belief to Bill:

> Bill need not really think, or have thought, the thought that is in question (to the effect that John is stupid and does not understand the matter). We can

> imagine that Bill was insincere, and perhaps overtly so, when he said, or implied that John was stupid and incompetent. Still John can use [(10)] to describe the situation. (Recanati, 2010, p. 201)

I take it to be obvious that the same observation applies, *mutatis mutandis*, to the other cases. Indeed, Holton himself acknowledges this point, and gives the following two examples:

> (11) a. She sold him a pig in a bag, but they both knew that it was really a cat.
> b. When they parted they exchanged photos and other keepsakes that would keep their love alive forever. But they both knew that it was hopeless.

But contrary to Recanati, Holton thinks cases of this kind merely show that

> We do not in general need a protagonist who has actually been fooled. It is good enough to have someone who is simply a *possible* location for the salient false belief, especially if they are the target of a pretence. (Holton, 1997, p. 627)

However, it is hard to see how one could spell out this suggestion in a satisfactory way. Suppose it was claimed that Protagonist Projection conveys belief-attributions of the form:

> Possibly, someone believes that...

So for example, with respect to (11a), the utterance would be seen as conveying that possibly someone believed that it was a pig. Again, it would then be up to the context to supply the relevant individual(s). In this case, the context would most likely supply both the referent of *she* and of *him*. So the utterance would be claimed to convey that they could have, or that it is possible that they believed that it was a real pig.

For this claim to be true in any interesting way, the notion of possibility must be construed rather narrowly, otherwise the claim becomes trivial. But on the other hand, it is clear that on such a narrow notion of possibility, there are plenty of cases of pretense for which it is not true that the protagonist could have believed the content in question. People can pretend things that they could not have believed (in the narrow sense), or at least for which the circumstances would have to be extreme for them to do so.

I will agree with Recanati, then, that we need an analysis of Protagonist Projection that is capable of accounting for the fact that sentences in this mode are consistent with readings on which they convey belief-attributions but also with weaker pretense-like readings. This, then, is the second feature that needs to be explained.

Let us sum up these desiderata as follows:

> D1. Sentences involving Protagonist Projection are not counterexamples to Karttunen's Thesis, Factivity, etc.

D2. Sentences involving Protagonist Projection convey that one or more protagonists believe/pretend that what is said is true.

The challenge is to sketch an analysis of Protagonist Projection which will explain both these two features.

3.2 Division of Labor: Semantics and Pragmatics

Holton's response to Tsohatzidis is a note of four pages. As such, it does not offer an account of Protagonist Projection, but merely makes the (important) observation that by classifying the examples in this way, we may avoid having to reject Karttunen's Thesis, Factivity, and other convictions that we make good use of, and for which we have good evidence. I will not speculate as to the details of the particular analysis that Holton may or may not have had in mind. Rather, I will present an analysis that I myself believe to be plausible and well motivated, and which straightforwardly preserves the spirit of Holton's proposal.

Let me make a few brief remarks here about the overall strategy of this approach. The account I will offer involves a particular view of Protagonist Projection as regards the semantics-pragmatics distinction. On my account, the semantics provides the explanation for D1, while pragmatics is responsible for the explanation for D2. In particular, the fact that sentences involving Protagonist Projection are not counterexamples to the theses in question will be seen to be a direct consequence of their truth conditions. On the other hand, the fact that Protagonist Projection communicates attributions of belief or pretense, depending on the context, will be explained in terms of an uncontroversial type of pragmatic inference.

This means that I will be endorsing the following argument as regards D1:

L1. φ does not entail not-*T*.

L2. Therefore, φ is not a counterexample to *T*.

The semantic analysis of Protagonist Projection I will offer will give rise to a notion of entailment that will preserve this argument. But the point to note up front is that - in contrast to Hazlett's line - the account will have the consequence that the sentences are not counterexamples regardless of whether they are true or false, or neither.

With respect to the pragmatics that will be invoked to explain D2, the strategy will be to appeal to the following inference:

P1. The speaker said that *p*.
P2. Therefore, the speaker believes that *p*.

I take it to be undeniable that hearers undergo this inference in normal circumstances. To have a fully spelled out account of this fact, though, a number of qualifications would be required. Two points are particularly important. First, the inference is defeasible. Second, it does not always take place. Yet, it will not be necessary to provide the details of these points here. Nor do we need to rely on any assumptions

concerning *why* this inference is drawn.[5] All we will need is to appeal to the undeniable fact that hearers do very often infer in the manner of P1-P2.

4. A Shifting Account of Protagonist Projection

Having introduced the strategy to be pursued above, this section details my account of Protagonist Projection. The chief goal is to articulate the analysis of the truth conditions of sentences involving Protagonist Projection that will motivate L1.

The central claim of my analysis is that Protagonist Projection involves a particular kind of *context-shift*. I therefore call it a *shifting* account of Protagonist Projection. This kind of context-shift has been argued by Recanati (2000), (2010) and Schlenker (2004) to be the chief mechanism involved in a closely related form known as *Free Indirect Discourse*. I begin by detailing this connection between Protagonist Projection and Free Indirect Discourse below.

4.1 Protagonist Projection and Free Indirect Discourse

Free Indirect Discourse is a style of report that blends Direct Discourse (quotation) and Indirect Discourse (that-reports). These three ways of reporting are illustrated here:

(12) a. Napoleon was relaxed. He thought, 'Those sloppy Englishmen I will meet tomorrow will cower before my strategic genius.' (*Direct Discourse*)
b. Napoleon was relaxed. He thought that the English troops he would meet the next day were disorganized and inferior to his own strategic skills. (*Indirect Discourse*)
c. Napoleon was relaxed. Those sloppy Englishmen he would meet tomorrow would cower before his strategic genius. (*Free Indirect Discourse*)

In all of these cases, we are given a report of Napoleon's thought. In the Free Indirect Discourse case in (12c), we are given a report that preserves some elements of the Direct Discourse report (12a), while maintaining the third person perspective of the Indirect Discourse report (12b).

The standard observation about Free Indirect Discourse is that its chief characteristic is to be found in the behavior of tenses, pronouns, and indexicals. Schlenker (2004) provides a useful summary:

> Free Indirect Discourse behaves as a mix of direct and of indirect discourse: tenses and pronouns take the form that they would have in an attitude report

[5] The inference from P1 to P2 is sometimes seen as a conversational implicature arising from Grice's (1975) First Maxim of Quality, 'Do not say what you believe to be false'. See for example Levinson (1983, 105). It should be noted that this kind of implicature is different from the perhaps more familiar type on which a maxim is flouted. So if one adopts this line, one is not committed to the view that Protagonist Projection is an example of saying something false to implicate something true, a category in which Grice himself placed other 'literary' phenomena, most notably irony and metaphor. This account of these phenomena is highly controversial. On irony, see e.g., Sperber and Wilson (1981), Recanati (2004), (2010), Wilson (2006), Camp (forthcoming). On metaphor, see e.g., Stern (2000), Bezuidenhout (2001), Camp (2006). On both, see e.g., Carston (2002).

> [...], while everything else - including *here, now, today, yesterday* and the demonstratives (e.g. *this*) - behaves as in direct discourse. (Schlenker, 2004, p. 283)

Or more concisely:

> In Free Indirect Discourse *everything* except pronouns and tenses is read *de dicto*, i.e. from the character's perspective. (Schlenker, 2004, p. 284)

So in (12c) the third person pronoun *he* is used to refer to Napoleon as it is in the Indirect Discourse (12b). Whereas *tomorrow* is used, not to refer to the day after the day of utterance, but is shifted so as to refer to the day after Napoleon did the thinking, as it is in the Direct Discourse (12a). As further illustration, here are two more examples:

(13) a. John looked at my picture. Yes(, he thought,) he wanted to marry me today.[6]
b. Tomorrow was Monday, Monday, the beginning of another school week![7]

Again, while pronouns and tenses behave as in Indirect Discourse, everything else behaves as in Direct Discourse. Note also that, as shown by (13a), one can often facilitate the Free Indirect Discourse reading by adding a parenthetical.

There is an intuitive sense in which Free Indirect Discourse and Protagonist Projection are very closely related. Both involve something that may be described as speaking from the point of view of someone else. To be more precise about the relation between them, I want to make two observations.

The first observation is that Free Indirect Discourse forces Protagonist Projection. For instance, the following example shows that inside Free Indirect Discourse one has to reorient *diamonds* in the same way as it was reoriented in Holton's original example:

> (14) When she saw it, Ann was convinced that the ring was studded with diamonds, although it was really only studded with cheap glass beads. Wow, she thought, how beautiful those diamonds/#glass beads were!

Inside the Free Indirect Discourse environment, the ring has to be described from the point of view of the protagonist. This agrees with the point noted above that, in Free Indirect Discourse, everything except pronouns and tenses is read *de dicto*. That is, things are described from the point of view of the protagonist of the story. Free Indirect Discourse necessarily involves Protagonist Projection.

The second observation is that Protagonist Projection does not force Free Indirect Discourse. The most direct evidence for this comes from the following kind of example:

> (15) A week ago, Ann was pacing around after coming home from the jeweler, disappointed and angry with John. The day before/#yesterday he gave her a ring studded with diamonds, but they turned out to be glass.

[6] From Sharvit (2008, 354).
[7] From Lawrence (1920, 185), cited in Banfield (1982, 98), Doron (1991, 53), Schlenker (2004, 280).

If *yesterday* could refer to the day before the day the story is about, then using it here would naturally explain why Ann was angry, and so *yesterday* should be acceptable. But the fact that *yesterday* is unacceptable in this case shows that it can only refer to the day before the day of utterance. So even though one element of a sentence is described *de dicto*, does not mean that everything can be.

I conclude that although Free Indirect Discourse always involves Protagonist Projection, the converse is not the case. One can use Protagonist Projection outside Free Indirect Discourse, but inside Free Indirect Discourse, one has to use Protagonist Projection. Free Indirect Discourse, then, can be thought of as a more complete type of projection of point of view. Whereas in Protagonist Projection the shift merely targets particular non-indexical elements, the additional feature of Free Indirect Discourse that indexicals like *yesterday* and *today* shift effectuates a more thorough identification with the point of view of the protagonist.

We need an analysis of our Protagonist Projection examples that preserves the connection with Free Indirect Discourse, but does not conflate the two phenomena. Below I present one way of doing so that draws on Schlenker's (2004) semantics for Free Indirect Discourse.

4.2 Schlenker's Semantics for Free Indirect Discourse

In standard semantic systems in the tradition from Kaplan (1989), Lewis (1980), and others, semantic evaluation is relativized to a context and an index. Yet, more recently, many semanticists - e.g., Banfield (1982), Doron (1991), Schlenker (1999), (2004), Sharvit (2004), (2008) - have agreed that what is needed to handle Free Indirect Discourse, and related phenomena, is to relativize semantic evaluation to *two* contexts.

In the system Schlenker (2004) develops, semantic evaluation is relativized to a *context of utterance* ν, a *context of thought* θ, and an assignment of values to variables *g*. The context of utterance includes a speaker a_ν, an addressee h_ν, a location l_ν, a time t_ν, and a possible world w_ν. The context of thought includes a time t_θ, a location l_θ, and a world w_θ. In the normal case, these two contexts overlap. That is, the time, location, and world of the context of thought are those of the context of utterance. But in non-standard modes of speech such as Free Indirect Discourse, they come apart.

Further, we define some context-sensitive expressions as depending on the context of utterance, and others as depending on the context of thought. The former are those that are not shifted - i.e., are read *de re* - in Free Indirect Discourse. But by defining some expressions as taking their values from the context of thought, these will be interpreted as shifted - i.e., as read *de dicto* - in cases, like Free Indirect Discourse, where the two contexts do not coincide. Hence, the pronouns and tenses are taken as depending on the context of utterance, whereas everything else - including *here, now, today, tomorrow, yesterday* - depend on the context of thought.

To illustrate this, consider Schlenker's Free Indirect Discourse example (16).

(16) Now he$_i$ was$_j$ rich(, John thought).

Here are Schlenker's entries for *now, he*, and the past tense:[8]

(17) a. $[[Now]]^{v,\theta,g} = t_\theta$.
b. $[[He_i]]^{v,\theta,g} = \#$ iff in w_v, $g(i)$ is not a male individual distinct from a_v and h_v. Otherwise, $[[He_i]]^{v,\theta,g} = g(i)$.
c. $[[past_i]]^{v,\theta,g} = \#$ iff $g(i)$ is not before t_v. Otherwise, $[[past_i]]^{v,\theta,g} = g(i)$.

This means that the past tense and *he* depend on the context of utterance, while *now* depends on the context of thought.[9]

Given this, (16) receives the following truth conditions:

(18) a. $[[(16)]]^{v,\theta,g} = \#$ iff either
(i) in w_v, $g(i)$ is not a male individual distinct from a_v and h_v, or
(ii) t_θ is not before t_v.
b. If $\neq \#$, $[[(16)]]^{v,\theta,g} = 1$ iff in w_θ, $g(i)$ is rich at t_θ.

This clause contains two parts. (18a) specifies when the sentence is defined, that is, when it has a truth value. According to (i) and (ii) the sentence has a truth value if and only if the individual assigned to *he* is (in the world of utterance) a male individual distinct from the speaker and addressee of the context of utterance, and the time of the context of thought (denoted by *now*) is before the time of the context of utterance. This reflects the intuition that if either of these requirements is not satisfied, (16) is infelicitous, in a strong sense. In turn (18b) states that, if defined, (16) is true if and only if the referent of *he* is rich at the time of thought in the world of thought.

4.3 Application to Protagonist Projection

Here is my proposal concerning Protagonist Projection: As with Free Indirect Discourse, a sentence involving Protagonist Projection is evaluated with respect to a non-actual context of thought. However, as opposed to what happens in Free Indirect Discourse, the context of thought coincides with the context of utterance *except* for the world-parameter w_θ which is shifted. Hence, in Protagonist Projection cases, not only tenses and pronouns, but also indexicals like *yesterday* will be read *de re*, even though evaluation for truth and falsity will still proceed with respect to a non-actual world of thought.

---

[8] It is assumed here, as opposed to the indexical treatment in Kaplan (1989), that the phi-features - i.e., person, gender, and number - of pronouns semantically function as preconditions on definedness, that is, as presuppositions. This kind of treatment of pronouns goes back to Cooper (1983). For a more recent instance, see Heim (2008). Similarly, in the tradition from Partee (1973b), tenses are here treated as pronouns - that is, as variables awaiting an assignment and governed by presuppositions.

[9] There is a potential problem here regarding the genders of pronouns. Sharvit (2008) argues that Schlenker's proposal makes wrong predictions because, she claims, genders in Free Indirect Discourse are read *de dicto*, i.e., from the subject's point of view. If so, then Schlenker's clause in (17b) needs to be revised (as does the corresponding clause for *she*). An obvious proposal is to let the genders rely on the world of thought, rather than the world of utterance. This would give the right results in cases where genders are read *de dicto* in Free Indirect Discourse. However, doing so would also make the proposal made in this paper for Protagonist Projection incorrect, since in these cases, as with indexicals, genders are read *de re*. There are two considerations that suggest that this difficulty can be overcome. First, Schlenker (2004, 290) notes some data to the effect that, at least in some cases, genders are read *de re* in Free Indirect Discourse, and so the *de dicto* effects may not be grammaticalized. Second, he therefore (2004, 291) suggests a strategy according to which *de dicto* pronouns in Free Indirect Discourse are treated as covert descriptions. I take it therefore that it is not ruled out that one may treat the gender features as depending on the world of utterance.

Let me illustrate this. Consider the indexed representation of the first conjunct of (3a), call it (3a'):

(3a') He$_i$ gave her$_j$ a ring studded with diamonds.

The truth conditions of this conjunct are:

(19) a. $[[(3a')]]^{v,\theta,g} = \#$ iff either
(i) in w_v, $g(i)$ is not a male individual distinct from a_v and h_v, or
(ii) in w_v, $g(j)$ is not a female individual distinct from a_v and h_v.
b. If $\neq \#$, $[[(3a')]]^{v,\theta,g} = 1$ iff in w_θ, $g(i)$ gave $g(j)$ a ring studded with diamonds.

Note that we have ignored the past tense here. It is clear, though, that what we will predict is that (3a') requires that the time of evaluation be before t_v, which is what we want. (I continue to ignore tenses in what follows.)

(19) states that (3a') is defined if and only if the referent of *he* is male in the actual world and is distinct from the speaker and hearer of the context of utterance, and the referent of *her* is female in the actual world and distinct from the speaker and hearer of the context of utterance. If defined, (3a') is true if and only if in the world of thought, the referent of *he* gave the referent of *her* a ring studded with diamonds.

The claim is, then, that whereas in the world of utterance w_v, the ring is studded with glass beads, in the world of thought w_θ, it is studded with real diamonds. So (3a') is true, and hence so is (3a) as a whole.

Next, consider (3b):

(3b) She$_i$ knew that he$_j$ would never let her$_j$ down, but, like all the others, he$_j$ did.

The analysis of this case is completely parallel to that of (3a). Call the first conjunct of this sentence (3b'). Then we calculate:

(20) a. $[[(3b')]]^{v,\theta,g} = \#$ iff either
(i) in w_v, $g(i)$ is not a female individual distinct from a_v and h_v, or
(ii) in w_v, $g(j)$ is not a male individual distinct from a_v and h_v.
b. If $\neq \#$, $[[(3b')]]^{v,\theta,g} = 1$ iff in w_θ, $g(i)$ knew that $g(j)$ would never let $g(i)$ down.

Consequently, (3b') is true if and only if the referent of *she* knew in w_θ that the referent of *he* would never let her down. As before, we claim that the world parameters of the two contexts come apart, and correspondingly she does know in w_θ, although she does not know in w_v. So (3b') is true, and hence so is (3b) as a whole.

Finally, this proposal correctly handles indexicals in these environments. We used (15) to demonstrate that Protagonist Projection does not force the indexical shifts characteristic of Free Indirect Discourse:

(15) A week ago, Ann was pacing around after coming home from the jeweler, disappointed and angry with John. The day before/#yesterday he gave her a ring studded with diamonds, but they turned out to be glass.

Given Schlenkerian semantics, *yesterday* will be associated with the following meaning:

(21) $[[\textit{Yesterday}]]^{v,\theta,g}$ = the day before the day of t_θ.

According to our proposal, sentences involving Protagonist Projection are evaluated with respect to a context of thought that coincides with the context of utterance except for the world parameter that is shifted. So, the second sentence of (15) is evaluated with respect to a context of thought for which the time is the same as the time of the context of utterance ($t_v = t_\theta$), although the world parameter is shifted ($w_v \neq w_\theta$). Consequently, *yesterday*, like the other shiftable indexicals, is evaluated with respect to the time of the context of thought. But in the case of Protagonist Projection the time of the context of thought *is* the time of the context of utterance, and hence *yesterday* in (15) would refer to the day before the day of utterance, in accordance with our intuitions about the case. And this explains why *yesterday* cannot be used here.

As this shows, the formal analysis preserves the relation between Free Indirect Discourse and Protagonist Projection described earlier. Since in the former case all the parameters of the context of though shift, this means that indexicals will be read *de dicto*. By contrast, in Protagonist Projection, the shift is confined to non-indexical elements, and this is accounted for by the semantics by just letting the world parameter of the context of thought shift.

4.4 A Problem and its Solution

Having illustrated the main idea of this proposal, we must now attend to a problem. As it stands, the above framework in fact does not assign adequate truth conditions to sentences involving Protagonist Projection. Our desideratum was to be able to maintain that (3a) may be true. Intuitively, the sentence is true just in case (i) the male character gave the female character a ring studded with what the relevant protagonist believed were diamonds and (ii) the ring was actually only studded with glass beads.

But given what we have so far, all we can say is that (3a) is true because the first conjunct is true at the context of thought of the protagonist who falsely believes that the ring is studded with diamonds and the second conjunct is true at the actual context of thought. Call the first of these contexts of thought θ_D. And in general let us agree to call the actual context of thought θ_v. Intuitively, the latter is the context of thought of the speaker, i.e., the context of thought that is intended to be the topic of talk for the vast range of normal assertions. So given the semantics so far, what we will be claiming is this

$[[\text{He gave her a ring studded with diamonds}]]^{v,\theta_D,g} = 1$, and
$[[\text{They turned out to be glass}]]^{v,\theta_v,g} = 1$.
Therefore, $[[(3a)]]^{v,\theta_v,g} = 1$.

However, at least in the absence of some non-trivial remedy, this means taking the questionable step of giving up the principle that a conjunction is true at a context of thought θ if and only if both its conjuncts are true at θ.[10] Abandoning this principle

[10] An anonymous reviewer suggests that a possibility is to preserve the principle as the 'official' rule, while supplementing it with an operation that shifts the contexts between the two conjuncts. I do not wish to suggest that

leads to obvious problems. Bluntly, one will not be able to infer from the truth of a conjunction that both conjuncts are true, nor vice versa.

In other words, we still do not have a satisfactory grasp of the truth conditions of sentences involving Protagonist Projection. In particular, we have not yet achieved an analysis that will allow us to *compositionally* determine the truth conditions of complex expressions containing projected constituents.

It is important to note that this is also a problem for the Free Indirect Discourse cases, given Schlenker's analysis, which we have been following. Let me explain why. One of Schlenker's key observations is that conjunctions like the following are not contradictory, i.e., they may be true (or false):

(22) Now, he was rich(, John thought), but in fact he wasn't rich.

As Schlenker points out, this is a feature that Free Indirect Discourse shares with both Direct and Indirect Discourse, but not with standard assertions. Here are his examples illustrating this:[11]

(23) a. [In a novel:] 'Tomorrow was Monday, Monday, the beginning of another school week!' (As it turned out, this wasn't true. The following day was Sunday).
b. [In a non-fictional context:] Tomorrow is Monday, the beginning of another school week. #This isn't true.
c. Tomorrow was Monday, Monday, the beginning of another school week! (As it turned out, this wasn't true.)
d. John thought that the following day was Monday, the beginning of another school week. (As it turned out, this wasn't true.)

But now the same problem arises as we noted above. Namely, we are faced with the problematic claim that (22) is true because the first conjunct is true at some context of thought (John's) while the second conjunct is true at another one.

We need a way, then, of preserving the principle that a conjunction is true at a context of thought if and only if both its conjuncts are true at the same context of thought. While at the same time retaining the spirit of the overall proposal, i.e., that the truth conditions of Free Indirect Discourse and Protagonist Projection essentially involve reference to a non-actual context of thought.

My solution to this problem is this: We assume that the context of utterance contains as a parameter an *accessibility relation* on contexts of thought, which we denote as R_v. Intuitively, R_v represents the range of contexts of thought that one can talk about in the context of utterance v. Whether one can successfully speak of non-actual contexts of thought is a feature of the context of utterance. Some contexts of utterance allow projection of this sort, others do not. These differences will be represented by a variation in R_v across different contexts of utterance. Since one can always speak of the actual context of thought, we will require that R_v be *reflexive*. I will say more about R_v below (in 4.7). For now, note that taking on board this idea, we can redefine truth in the following fashion:

---

there is no possible remedy to this problem. However, I think the solution I propose below is preferable to one that introduces supplementary operations.

[11] Schlenker's (2004, 285) text has the present tense in (23c), but this is obviously a typo, and I therefore change it to the past tense. Cf. (13b).

$[[\varphi]]^{\nu,\theta,g} = 1$ iff there is a θ' s.t. $R_\nu(\theta,\theta')$ and φ is the case in θ'.
$[[\varphi]]^{\nu,\theta,g} = 0$ iff for all θ' s.t. $R_\nu(\theta,\theta')$, φ is not the case in θ'.

This is a rough definition. In particular, we are here ignoring all dependencies on both the context of utterance and the context of thought induced by indexicals, persons, genders, tenses, etc. But I take it to be clear that these can be incorporated here in the same way as we have seen earlier.

Given this conception of truth, we can now maintain that

$[[\text{He gave her a ring studded with diamonds}]]^{\nu,\theta_\nu,g} = 1$, and
$[[\text{They turned out to be glass}]]^{\nu,\theta_\nu,g} = 1$.
Therefore, $[[(3a)]]^{\nu,\theta_\nu,g} = 1$.

The reason for this is that we will claim that, in this case, R_ν relates the actual context of thought θ_ν to two contexts of thought, namely itself (by reflexivity) and θ_D. The first conjunct is true because θ_D is accessible, and the second conjunct is true because θ_ν is accessible.

This carries over in the obvious way to the cases of Free Indirect Discourse. For example, (22) will be true if and only if there is an accessible context of thought where John was rich at the relevant time and there is an accessible context of thought where he was not. And this is so: John's context of thought verifies the first conjunct, the actual context of thought verifies the second.

4.5 Entailments

The crucial consequence of the amendment just proposed is that it allows us to vindicate the desired entailments. Our goal was to develop a semantic analysis of Protagonist Projection that could provide the theoretical underpinnings of Holton's conclusion. That is, we want to be able to maintain that the diamond sentence does not entail that some diamonds are made of glass. But moreover, we also want to preserve other intuitive entailment relations. For example, we want to be able to maintain that the first conjunct of the diamond sentence entails that something is studded with diamonds.

Both of these goals can now be achieved. First, we define entailment as follows:

$\varphi \models \psi$ =df for all ν,θ,g: if $[[\varphi]]^{\nu,\theta,g} = 1$, then $[[\psi]]^{\nu,\theta,g} = 1$.

This is the analogue of the familiar conception of entailment as truth preservation that we introduced earlier.[12] It can then be seen clearly that, given this notion of entailment, we will validate the following:

(24) a. (3a) $\not\models$ Some diamonds are made of glass.
b. (3a') $\models$ Something is studded with diamonds.

[12] As such, it corresponds to the usual conception of entailment in intensional semantics in this tradition. For example, Kaplan (1977, 547) defined validity (logical truth) as truth in all contexts, where truth in a context *c* is in turn defined as truth with respect to *c* and the index of *c*, i.e., the index obtained by plugging in the parameters of *c*. Derivatively, entailment is then understood as truth preservation in all contexts (in this technical sense).

Consider first (24a). It is easy to see that there are contexts where (3a) is true, although it is false that some diamonds are made of glass. That is, it is obvious that although there are accessible contexts of thought where the ring is studded with real diamonds, this does not mean that there are accessible contexts of thought where some diamonds are made of glass. Indeed, in the scenario we have been imagining, neither θ_D nor θ_v is like that. Although the protagonist falsely believes that the ring is studded with diamonds, there is no reason to assume that he or she believes that some diamonds are made of glass. And nor does the speaker believe that, obviously. So, we maintain that (3a) is true because each of its conjuncts is verified by an accessible context of thought, but since none of these verify that some diamonds are made of glass, the entailment fails.

Next, here is why (24b) holds. All contexts of thought in which someone is given a ring studded with diamonds are contexts of thought in which something is studded with diamonds. So for all contexts, if there is an accessible context that verifies (3a'), there is an accessible context (the same one) that verifies 'Something is studded with diamonds'. Hence, the entailment holds.

The fact that (24b) holds has a further important consequence. According to (24b), the first conjunct of (3a) entails that something is studded with diamonds. In general, if *A* entails *C*, then *A&B* entails *C*. And accordingly, we also predict that (3a) entails that something is studded with diamonds:

(25) (3a) |= Something is studded with diamonds.

Again, it is not hard to see why. (3a) is true at the actual context of thought θ_v because for each of its conjuncts there is an accessible context of thought that verifies it. And by the same token, there is an accessible context of thought that verifies that something is studded with diamonds, namely the one that verifies the first conjunct, i.e., θ_D.

A word of caution is in order here. The result in (25) can seem counterintuitive. I.e., one might think that since the whole point of (3a) is that the ring was not really studded with diamonds, (3a) should not entail that something is studded with diamonds. But this line of thought fails to recognize that it does *not* follow from the actual truth of 'Something is studded with diamonds' that something is actually studded with diamonds. Such is the nature of our proposal. One may speak of other contexts of thought than the actual one, if the setting is right. (3a) does just that. And so the reason that (3a) entails that something is studded with diamonds is that, in the *non-actual* context of thought we are projected into, something is studded with diamonds.

Now let us turn to the cases concerning Factivity. Consider again (3b).

(3b) She knew he would never let her down, but, like all the others, he did.

As with the diamond case, this conjunction is true at the actual context of thought θ_v because each conjunct is true at an accessible context of thought. The first conjunct is verified by the context of thought of the protagonist who falsely believes that the subject knew that he would never let her down. The second conjunct is verified by the actual context of thought itself, which is accessible by reflexivity. And hence the conjunction is true.

We can now see that (3b) does not threaten Factivity. Assume that Factivity holds. This means that

(26) She knew he would never let her down |= He would never let her down.

Indeed, that this should hold is intuitively plausible. Clearly, the contexts of thought in which she knew that he would never let her down are contexts of thought in which he would never let her down. The protagonist who believes the former obviously also believes the latter.

Further, note that, as we saw above, our system preserves the principle that if *A* entails *C*, then *A&B* entails *C*. Hence, we predict that (3b) entails that he would never let her down:

(27) (3b) |= He would never let her down.

As with the diamond case, it is important to note that, even though (3b) is true at the actual context of thought, the entailment in (27) does not license the conclusion that he actually would never let her down. The right hand side of (27) is true at the actual context of thought because there is an accessible (non-actual) context of thought at which he would never let her down. Namely the same (non-actual) context of thought that verifies the first conjunct of (3b).

So since (3b) entails the complement of the knowledge-ascription, it is not a counterexample to Factivity. In general, it is obvious that we predict that any context of thought at which a knowledge-ascription is true is a context of thought at which the complement is true. Hence, Factivity holds.

To sum up, then, we have seen that the shifting account of Protagonist Projection vindicates Holton's conclusion with respect to the diamond case and the cases of Factivity. I take it to be relatively obvious how this applies, *mutatis mutandis*, to the putative counterexamples to Karttunen's Thesis, and I will refrain from going through more examples.[13] So we have provided an intuitively adequate way of satisfying our first desideratum D1.

4.6 Explaining Attributions

Our second desideratum D2 was to explain why cases of Protagonist Projection convey attributions of belief or pretense. The same question arises for cases of Free Indirect Discourse. Intuitively, utterances involving Free Indirect Discourse convey reports about the subject's beliefs. But Schlenker's semantics itself does not provide

[13] There are aspects of the phenomenon we have not addressed. One concerns constraints on shifting. An anonymous reviewer points to the infelicity of the following, even given a fair amount of contextual priming: 'His lawyer knew that OJ was innocent, but really he was certain that OJ was guilty.' Arguably, the pragmatic constraints on the possibilities for shifting are influenced by what participants are interested in. Since we are often interested in detecting and revealing illusory knowledge, it might be conjectured that *know* tends to project into the point of view of the subject of the verb. In the example above that projection pattern threatens to generate an inconsistency. One cannot normally believe one knows something while being certain of its negation. Many of these features will be explained by constraints on the accessibility relation between contexts we have introduced. (See also 4.7 below.) There are other features of Protagonist Projection that will need to be explained by a complete account. E.g., there may be a limited range of possible relations between the 'literal' meaning and the shifted meaning. Some evidence for this is that, in each case, the shifted meaning preserves a remnant of original meaning; i.e., in cases involving *tell*, a speech act was actually made, in cases involving *know*, a belief was actually entertained, etc.

an explanation of this fact. Indeed, Schlenker is explicit that, on his analysis, 'a sentence in Free Indirect Discourse is not strictly speaking a report, but rather expresses a thought through someone else's mouth [...].' (2004, 295) So what more precisely is meant by expressing a thought through someone else's mouth, and how can one account for the attitude attributions along these lines?

The first thing to consider is what proposition, or thought, *is* expressed by a Free Indirect Discourse sentence on this view. In semantic systems in this tradition, one usually defines the proposition expressed by a sentence as a function from indices to truth values.[14] Similarly, one can define a notion of the proposition expressed in Schlenker's two-context system as follows:

The proposition expressed by φ relative to ν and g =df $\lambda\theta.[[\varphi]]^{\nu,\theta,g}$

Schlenker takes this to be a way 'to recover the *thought* attributed to the agent.'[15] (2004, 293) But he does not tell us how more precisely to understand this claim. We can agree that the above captures a familiar notion of the proposition expressed, or 'what is said', but how do we go from there to the fact that this proposition is attributed as the object of an attitude had by the relevant agent?

To be sure, this is not a problem for the semantics itself, but more an incompleteness of the overall picture. Although the account gives us a way of identifying which proposition is attributed to the subject in question, it does not tell us anything about why this proposition is attributed. There is no explanation of why hearers infer that the speaker is intending to attribute the proposition defined above to the relevant protagonist(s).

As I said earlier, I want to advocate a strategy according to which this fact is to be explained pragmatically. That is, on this view, the semantics tells us the truth conditions of sentences involving Free Indirect Discourse and Protagonist Projection, and derivatively, it gives us a proposition expressed, defined as above. The further fact that these cases also communicate attributions of belief or pretense is to be explained by an additional pragmatic component. Below, I provide a rough sketch of what I take to be the most plausible direction for such an account.

The central part of the proposal is that, as mentioned earlier, hearers typically infer from the fact that a speaker said that p to the conclusion that the speaker believes that p. Given our framework, we can state this general inference pattern roughly as follows:

(i) The speaker said that p is the case in θ.
(ii) The speaker believes that p is the case in θ.

Next, we assume that, for standard utterance situations, hearers entertain a general assumption to the effect that the speaker believes that the world of thought is actual. This corresponds to what we assumed above. When someone says something, it is usually intended to be true in the actual context of utterance, that is, in the world of thought corresponding to that context, and the accessibility relation R_ν will be

[14] Following Carnap (1956), this kind of function is often called an *intension*. Kaplan (1989) identified intensions with *what is said*, or the *content* of, or *proposition expressed* by an utterance, relative to a context.
[15] Schlenker cautions that more complicated cases require more than this, i.e., cases involving 'essential' indexicals such as Kaplan's (1989) well known pants-on-fire example. Schlenker provides a sophisticated way of analyzing these based on Kaplan (1969). These kinds of cases can be safely ignored for the purposes of this paper.

correspondingly restricted. The claim is, then, that in normal cases, hearers infer as follows:

> (iii) Assumption: The speaker believes that θ is actual.
> (iv) Therefore, the speaker believes that *p* (is actually the case).

In the case of Protagonist Projection and Free Indirect Discourse, however, hearers realize that (iii) is not in play for the relevant sentences. This is a simple way of spelling out the intuition that hearers know how to interpret sentences like the diamond sentence so as not to infer that the speaker is proposing that some diamonds are made of glass, and similar implausible claims. In the terminology of Hazlett, this explains the 'deliberateness' of the examples. Rather, in these cases, another assumption is in place, and hearers therefore infer a different conclusion:

> (iii') Assumption: The speaker believes that someone *x* believes/pretends that θ is actual.
> (iv') Therefore, the speaker believes that *x* believes/pretends that *p* (is actually the case).

So the idea here is that if you know that the speaker believes that *x* believes or pretends that θ is actual, and you hear the speaker say that *p* is the case in θ, you infer that the speaker believes that *x* believes or pretends that *p*. This kind of pragmatic account is a promising way of explaining the attributions of belief or pretense conveyed by sentences involving Free Indirect Discourse or Protagonist Projection.[16] I will not try to spell this out in more detail here, although to be sure, a lot more needs to be said. Yet, I assume that with this we have good reasons to believe that we can satisfy our second desideratum.

4.7 Free Shifting vs. Accessibility

The final feature of the proposal I want to turn to is a foundational one. It concerns the question of how to understand the very idea of evaluation relative to non-actual contexts of thought. As mentioned, traditional semantics in the Kaplan-Lewis tradition relativizes semantic evaluation to a context and an index. The context supplies referents for indexicals, demonstratives, etc. in order to determine a content, or a proposition expressed, which is then evaluated for truth or falsity at a particular index.[17]

---

[16] The claim that, in cases of Protagonist Projection, (iii') supplants the more usual (iii) is an, albeit idealized, empirical observation about the cases. Correspondingly the inference terminating in (iv') is intended to represent the pragmatic inference that hearers undergo leading them to realize that the speaker intended to attribute the thought in question to someone other than herself. We have not given an account of the factors that determine whether something like (iii') is in place. Importantly, examples like (3a) threaten to be unintelligible if read as standard assertions; i.e., the speaker cannot reasonably be taken to believe that the ring was studded with things that are both diamonds and made of glass. By contrast, the following type of example, mentioned by an anonymous reviewer, is clearly best read as a standard assertion: 'John sold me a ring of diamonds, but/even though he wasn't sure it wasn't made of glass.' Since reasoning along the lines of the standard (iii), in this case, yields a reasonable interpretation, a projected reading is not readily available.

[17] To what extent this contextual determination is a semantic matter, and which expressions count as indexicals, are controversial issues. For some discussion, see Stanley (2000), Carston (2002), Cappelen and Lepore (2004), Recanati (2004), Predelli (2005), Stokke (2010).

It is useful to think of Schlenker's notion of a context of thought as an index in this sense in that truth values are determined by the context of thought. One major difference from more traditional systems is then that, for some context-sensitive expressions, their values are determined not by the context but by the index (context of thought). In particular, all the expressions that are read *de dicto* in Free Indirect Discourse (e.g., *now*, *yesterday*) take their values from the context of thought.

What I want to focus on here, though, is another point of divergence from the mainstream way of doing semantics in this tradition. In the original system presented in Schlenker (2004), the truth of a sentence sometimes depends on a non-actual context of thought. Characteristically, this happens with Free Indirect Discourse. In traditionalist terms, this amounts to the claim that sometimes, truth (that is, actual truth) depends on evaluation at a *shifted* index. There is nothing peculiar about this claim. What is peculiar is that, in Schlenker's system, this shift is *free*, that is, it is not triggered by the presence of an operator in the syntax.

According to other proposals, e.g., Schlenker (1999), Sharvit (2008), Free Indirect Discourse is to be analyzed as generated by a covert operator, which functions as an attitude verb, and hence Free Indirect Discourse is seen as a special kind of direct discourse. By contrast, one of Schlenker's (2004) main objectives is to argue that 'Free Indirect Discourse can be analyzed without recourse to any modal operator.' (2004, 293) This line was originally advocated by Banfield (1982) who provides a range of arguments to the effect that Free Indirect Discourse is never embedded.[18]

As we have seen, the crucial mechanism of Schlenker's (2004) account of Free Indirect Discourse is that the context of utterance and context of thought may come apart. But in accordance with the general line just described, there is no special syntax associated with Free Indirect Discourse. In other words, there is no operator that triggers the shift. The fact that, in cases of Free Indirect Discourse, the context of thought is different from the context of utterance is not accountable for by any syntactically driven process. In Kaplanian terms, then, the proposal here is that the index of evaluation can shift freely, without this shift being triggered by an intensional operator.

From a traditionalist point of view, this idea can seem puzzling. Orthodoxy accepts, as Lewis says, that

> Often the truth (-in-English) of a sentence in a context depends on the truth of some related sentence when some feature of the original context is shifted. (Lewis, 1980, p. 27)

But orthodoxy also assumes that shifting occurs only when an operator in the syntax triggers it. Examples of such operators are *possibly, necessarily, it was the case that, it will be the case that*, and more controversially, *in Paris*, *somewhere*, or *strictly speaking*.[19] By contrast, Schlenker's (2004) semantics assumes that the world of the

[18] This debate is far from settled. Although I have been following Schlenker (2004), I have not argued against the operator-approach to Free Indirect Discourse (or analogous approaches to Protagonist Projection). It remains an open question which of these approaches is ultimately correct.

[19] For discussion, see Lewis (1980), Cappelen and Hawthorne (2009). It should be noted here that extending the system I propose in the paper to take into account intensional operators is likely to be non-trivial. In particular, one will probably want to appeal to other notions of accessibility than the accessibility relation between contexts of thought invoked here, and further, there may be some kinds of modality that one will want to treat in other ways than by appealing to accessibility at all. Similarly, there is a question to be asked about mixed cases, although it is not so clear what the range of felicitous mixed cases is.

index - or the context of thought - can be shifted without this shift being triggered by an operator.

There are other kinds of examples than Free Indirect Discourse, for which this kind of free shift has been posited. For example, Recanati (2000) discusses the following case:

(28) There is a great film showing at the Piazza. A giant spider swallows New York City.

As Recanati observes, the second sentence in (28) is intended to be evaluated with respect to a non-actual world, i.e., the world of the film. That is, we do not take the second sentence to be obviously false (because there actually is no giant spider swallowing New York City). Rather, we take the sentence to be true or false depending on whether it is an adequate description of the film. But there seems to be no operator, such as *in the film*, or the like, which triggers this shift.[20]

Motivated by the problems involving the truth conditions of conjunctions, I proposed a modification of Schlenker's (2004) semantics, on which truth is defined with reference to the contextually determined accessibility relation on contexts of thought. This account can be conceptualized as occupying a middle ground between the operator-driven semantics of Schlenker (1999) or Sharvit (2008) and the free accounts of Recanati (2000) and Schlenker (2004). Although, on the view I have proposed, the shift to a non-actual context of thought involved in Protagonist Projection and Free Indirect Discourse is not triggered by a (covert) operator in the syntax, there is a sense in which it is more controlled than on the liberal proposals of Recanati (2000) and Schlenker (2004).

Let me explain what I mean by looking at an example. Suppose that, during the course of a conversation about cosmology, I say:

(29) The Earth is flat.

I take it to be obvious that we want to say that this utterance is false. That is, we assume that the context is one in which the speaker cannot achieve projection into a point of view that makes the utterance true. Of course in the right setting, I *could* be speaking in the mode of Protagonist Projection and thus be describing false beliefs, to use Holton's phrase. But the situation has to be a certain way for this to be possible.

The contrast between the free-shift theory and the accessibility theory now becomes clear. In Schlenker's (2004) framework, the blocking of the shift does not show up in the semantics *per se*. However, given the accessibility theory, we can say that the reason (29) is false in a normal setting is that R_v is such that no suitable contexts of thought are accessible. The only accessible context of thought is the actual context of thought, at which the Earth is not flat.

Given that R_v is a parameter of the context of utterance, the accessibility theorist has built into her system the feature of utterance contexts that she takes to be the one controlling the possibility of non-standard modes of speech such as Protagonist Projection. One may very well feel that this provides a better handle on the restrictions on shifting.

[20] By contrast, Lewis (1978) famously argued that sentences about fictions are pre-fixed by an operator of this sort. Others, though, do things differently. For instance, Predelli (2005) argues for a treatment of fictional sentences where the world of evaluation is shifted to a non-actual, fictional world, without the mediation of a fiction-operator.

That said, the main argument for the accessibility theory against Schlenker's free approach remains the one we brought out earlier, namely that the latter type of account is less straightforwardly in a position to handle compound sentences.

To be sure, it might be objected that we have given no account of *when* the accessibility relation allows for Protagonist Projection, and when it does not. That is, we have given no account of what the utterance situation has to be like for Protagonist Projection to be possible. All we have said is that the possibility of Protagonist Projection depends on the context of utterance.

I take it to be fairly uncontroversial that for the purpose of providing a descriptive semantic account of a particular class of phenomena the theorist is allowed to assume that this kind of foundational issue can be settled independently. Here is an analogy. It is widely agreed that quantificational determiners like *all*, *some*, etc. are highly context-dependent in the sense that what domain they quantify over depends on the context of utterance. Yet, it is safe to say that it will be a very complicated matter to explicate precisely *how* their domains are determined in context.

This kind of issue is sometimes referred to as a *metasemantic* issue.[21] And so, the reply to the objection that I am gesturing at is that the account of how the value of R_v is determined by the context is a metasemantic matter, just like the account of how the domain of a quantifier is determined. Hence, for the purpose of the descriptive semantics, we should be allowed to relegate it to another area of inquiry.

5. Conclusion

We have seen that, despite Hazlett's arguments, Holton's suggestion for handling putative counterexamples to well established theses such as Karttunen's Thesis and Factivity can be underwritten by a plausible semantic analysis of the examples in question. Further, this analysis provided us with a way of explaining the connection between Protagonist Projection and Free Indirect Discourse, as well as the differences between them.

References

Banfield, A. 1982. *Unspeakable sentences*. Boston: Routledge & Kegan Paul.

Bezuidenhout, A. 2001. Metaphor and what is said. *Midwest Studies in Philosophy*, *25*, 156–186.

Camp, E. 2006. Contextualism, metaphor, and what is said. *Mind and Language*, *21*3, 280–309.

[21] For relevant discussion, see Stanley and Szabó (2000), Glanzberg (2007), Ichikawa (2010).

Camp, E. forthcoming. *Sarcasm, pretense, and the semantics/pragmatics distinction.* Forthcoming in *Noûs*.

Cappelen, H., & Hawthorne, J. 2009. *Relativism and monadic truth*. Oxford and New York: Oxford University Press.

Cappelen, H., & Lepore, E. 2004. *Insensitive semantics – a defense of semantic minimalism and speech act pluralism*. Oxford: Blackwell.

Carnap, R. 1956. *Meaning and necessity: A study in semantics and modal logic. 2nd ed. with supplementary essays.* Chicago: University of Chicago Press.

Carston, R. 2002. *Thoughts and utterances*. Oxford: Blackwell.

Cooper, R. 1983. *Quantification and syntactic theory*. Dordrecht: Reidel.

Doron, E. 1991. Point of view as a factor of content. *Proceedings of SALT*, *1*, 51–64.

Glanzberg, M. 2007. Context, content, and relativism. *Philosophical Studies*, *136*, 1–29.

Grice, H. 1975. Logic and conversation. In *Studies in the way of words*. Cambridge, MA: Harvard University Press, 1989.

Hazlett, A. 2010. The myth of factive verbs. *Philosophy and Phenomenological Research*, *80*, 497–522.

Heim, I. 2008. Features on bound pronouns. In D. Harbor, D. Adger, & S. Bèjar Eds., *Phi-theory: Phi-features across modules and interfaces*. Oxford and New York: Oxford University Press.

Holton, R. 1997. Some telling examples: A reply to Tsohatzidis. *Journal of Pragmatics*, *28*, 625–628.

Ichikawa, J. J. 2010. Quantifiers and epistemic contextualism. *Philosophical Studies Online First*.

Kaplan, D. 1969. Quantifying in. In D. Davidson & J. Hintikka (Eds., *Words and objections: Essays on the work of W. V. O. Quine*. Dordrecht, Holland: Reidel.

Kaplan, D. 1989. Demonstratives. In J. Almog, J. Perry, & H. Wettstein (Eds., *Themes from Kaplan*. Oxford and New York: Oxford University Press.

Karttunen, L. 1977. Syntax and semantics of questions. *Linguistics and Philosophy*, *1*, 3–44.

Lawrence, D. 1920. *Women in love*. London: Heinemann, 1970.

Levinson, S. 1983. *Pragmatics*. Cambridge: Cambridge University Press.

Lewis, D. 1978. Truth in fiction. *American Philosophical Quarterly*, *15*, 37–46.

Lewis, D. 1980. Index, context, and content. In *Papers in philosophical logic*. Cambridge: Cambridge University Press, 1998.

Partee, B. 1973a. The semantics of belief sentences. In J. Hintikka, J. Moravcsik, & P. Suppes Eds., *Approaches to natural language*. Dordrecht, Holland: Reidel.

Partee, B. 1973b. Some structural analogies between tenses and pronouns in english. *Journal of Philosophy*, *70*, 601–609.

Predelli, S. 2005. *Contexts – meaning, truth and the use of language*. Oxford: Clarendon Press.

Recanati, F. 2000. *Oratio obliqua, oratio recta*. Cambridge, MA: MIT Press.

Recanati, F. 2004. *Literal meaning*. Cambridge: Cambridge University Press.

Recanati, F. 2010. *Truth-conditional pragmatics*. Oxford and New York: Oxford University Press.

Schlenker, P. 1999. *Propositional attitudes and indexicality.* PhD dissertation, Massachusetts Institute of Technology.

Schlenker, P. 2004. Context of thought and context of utterance: A note on free indirect discourse and the historical present. *Mind and Language*, *19*, 279–304.

Sharvit, Y. 2004. Free indirect discourse and 'de re' pronouns. *Proceedings of SALT*, *14*.

Sharvit, Y. 2008. The puzzle of free indirect discourse. *Linguistics and Philosophy*, *31*, 353–395.

Sperber, D., & Wilson, D. 1981. Irony and the use-mention distinction. In P. Cole Ed., *Radical pragmatics*. New York: Academic Press.

Stanley, J. 2000. Context and logical form. In *Language in context – selected essays*. Oxford: Clarendon Press, 2007.

Stanley, J., & Szabó, Z. G. 2000. On quantifier domain restriction. In *Language in context – selected essays*. Oxford: Clarendon Press, 2007.

Stern, J. 2000. *Metaphor in context*. Cambridge, MA: MIT Press.

Stokke, A. 2010. Intention-sensitive semantics. *Synthese*, *175*, 383–404.

Tsohatzidis, S. 1993. Speaking of truth-telling: The view from wh-complements. *Journal of Pragmatics*, *19*, 271–279.

Wilson, D. 2006. The pragmatics of irony: Echo or pretence? *Lingua*, *116*, 1722–1743.