

Hazel Pearson
Harvard University

Forthcoming in *Journal of Semantics*.

A Judge-Free Semantics for Predicates of Personal Taste

Abstract

We offer a new account of the semantics of predicates of personal taste (PPTs) like *tasty* and *fun* which, unlike recent proposals (Lasersohn 2005, Stephenson 2007a,b), does not appeal to a judge parameter as a component of the evaluation index. We identify empirical shortcomings of previous proposals, arguing that PPTs have a first person oriented meaning component even in cases that seem to involve an exocentric interpretation. We propose that the interpretation of PPTs involves first person oriented genericity in the sense of Moltmann (2006, 2010). When I say *The cake is tasty*, I say roughly that for all worlds w and all individuals x such that x is relevant in w and I identify with x , the cake is tasty to x in w . We explain the shifting of the first person orientation from speaker to attitude holder in attitude reports by taking both matrix and embedded sentences to express properties rather than propositions (Stojanovic 2011). In both cases, an abstraction operator in the left periphery of the clause binds the variable responsible for the first person oriented interpretation of the sentence. The paper closes with a discussion of the implications of the proposal for the analysis of attitudes de se.

1. Introduction

Consider the following dialogue.

- 1a. John: This cake is tasty.
- 1b. Mary: No, it's not tasty.

This seemingly innocuous exchange is the departure point for a family of proposals in the formal semantics and philosophy of language literature that constitute a radical revision of the traditional view of the considerations determining whether a given sentence is true or false (on the semantic side: Lasersohn 2005, Stephenson 2007a, 2007b; on the philosophy of language side: Koelbel 2002, 2003, Egan 2010 *inter alia*). In these works, it is proposed that in addition to truth being relativized to a world and a time, it is also determined relative to an additional parameter – the ‘judge’ whose tastes or opinions are relevant, perhaps, or the standard of taste against which an evaluation of tastiness, fun, beauty etc. is being made.

Most versions of these proposals start with the observation that a dialogue like (1) constitutes a case of so-called ‘subjective disagreement’ (Stephenson 2007a, 2007b). The

claim is that use of predicates of personal taste (PPTs) such as *tasty* and *fun* may give rise to a subjective disagreement – that is, a disagreement with the following properties:¹

2. *Subjective disagreement*

(I) Intuitively, the interlocutors disagree with one another.

(II) There is a sense in which both speakers have said something true, so long as each was sincere in her expression of her opinion.

(III) For this reason, the disagreement does not seem to be one that can be resolved.²

The puzzle is this: how can it be that two speakers disagree with each other – indeed, that they apparently make assertions with incompatible contents, as indicated by the felicitous use of *No...* and negation in (1b) – and yet both of them speak truly? The target of this paper is recent attempts in the formal semantic literature to solve the puzzle by adding to the evaluation index an individual type parameter corresponding to the ‘judge’ (Lasersohn 2005, Stephenson 2007a, 2007b). Such a move complicates the semantics by requiring that each predicate takes an additional individual argument as well as a world (and time) argument, with implications across the inventory of semantic types: the intension of a 1-place predicate like *table* or *run* is now of type $\langle s, \langle e, \langle e, t \rangle \rangle \rangle$, that of a 2-place predicate like *love* is of type $\langle s, \langle e, \langle e, \langle e, t \rangle \rangle \rangle \rangle$, and so on.³ All things being equal, a solution to the puzzle of subjective disagreement that does not necessitate such a complication should be preferred. We shall see that all things are not equal, however: we identify a series of empirical challenges faced by theories that appeal to the judge parameter, and develop a semantics for predicates of personal taste that meets these challenges without complicating the semantics. The resulting judge-free semantics for PPTs is we think to be preferred.

The paper is structured as follows. In section 2 we examine Lasersohn’s and Stephenson’s proposals and identify drawbacks of the two accounts. In section 3 we propose a new semantics for predicates of personal taste. In section 4 we show how the treatment meets the empirical challenges identified for Lasersohn and Stephenson, explain how it accounts for subjective disagreement, and highlight some correct predictions. To anticipate, we claim that the semantics of PPTs involves a peculiar species of genericity called ‘first-person oriented genericity’ (Moltmann 2006, 2010a); we therefore make one or two forays into the realm of generics along the way. Our findings also have implications for the semantics of attitude reports and of *de se* expressions, and we conclude with a discussion of these.

2. Lasersohn (2005) and Stephenson (2007a, 2007b)

Intuitively, PPTs are predicates that give rise to statements whose truth or otherwise is a matter of opinion rather than a matter of fact. Lasersohn takes it that examples like those in (3) fall into the former category, while the cases in (4) fall into the latter one.

¹ One also finds the term ‘faultless disagreement’ (Koelbel 2002, 2003) being applied to essentially the same phenomenon, the word ‘faultless’ being chosen because apparently neither party to the dispute is at fault.

² Because of this, one might question whether it is rational to engage in such an argument, which as Stephenson (2007a, 2007b) notes is a separate issue.

³ For simplicity, I ignore reference to times and tenses throughout the paper.

- 3a. This cake is tasty.
 3b. Rollercoasters are fun.

- 4a. John is a doctor.
 4b. Paris is in France.

Only sentences like those in (3) can participate in subjective disagreements: if you assert that John is a doctor, and I reply, *No, he's not!* then one of us has said something false. In this paper, we focus fairly narrowly on the predicates in (3). In principle we think that the account proposed is at least applicable to other predicates that (i) can take an overt Experiencer argument as in *tasty to Mary* or *fun for John*; (ii) can occur without such an overt argument as in (3); and (iii) express statements whose truth is a matter of opinion when they occur without an overt Experiencer argument. Examples include *attractive*, *boring* and *funny*. Additionally, it may be desirable to extend the analysis to predicates such as *smart* and *silly* that satisfy (ii) and (iii) but not (i). We leave it to future research to determine how successful our analysis, or a variant of it, might be for this class of predicates.

Here is a first attempt to explain why PPTs give rise to subjective disagreements. Lasersohn notes that if I say that this cake is tasty, I seem to be saying that I find it tasty, that it is tasty to me. One might therefore think that PPTs have a covert indexical argument referring to the speaker. Assuming a Kaplanian framework, the content of (1a) would therefore be ‘this cake is tasty to John’, while the content of (1b) would be ‘this cake is not tasty to Mary’.⁴ Lasersohn shows that by assigning these contents to (1a) and (1b), we fail to predict that their use gives rise to the kind of disagreement that is attested, as is revealed by comparison with (5).⁵

- 5a. John: This cake is tasty to me.
 5b. # Mary: No, it's not tasty to me!

Lasersohn takes this as evidence that assigning PPTs a covert indexical as their internal argument is the wrong move. His diagnosis of the problem is as follows. That an assertion of the form ‘x is P’ can be countered with ‘No, x is not P’ (where P is a PPT) shows that both occurrences of ‘x is P’ in the dialogue express the same content, even though they are uttered by different people. This is why treating PPTs as having an indexical argument will not do. Nonetheless, we must capture the fact that *This cake is tasty* may be true in my mouth but false in yours – all it would take would be for me to find the cake tasty and you to dislike it. Lasersohn’s solution is to let the truth of sentences containing PPTs be evaluated relative to a judge parameter, in addition to the usual world parameter.

6. $\llbracket \text{‘This cake is tasty’} \rrbracket^{c, w, j} = 1$ iff this cake is tasty to j in w.

⁴ We ignore the role of the demonstrative in these sentences, whose presence is harmless as long as one ensures that *this cake* has the same content in both John and Mary’s utterances.

⁵ (5b) is improved by replacing *no* with *well*, as shown in (i).

(i) John: This cake is tasty to me.
 Mary: Well it isn’t tasty to me!

One might even argue that the dialogue in (i) constitutes a disagreement. The point is that it is not the same kind of disagreement as in (1), as shown by the infelicity of the use of *no* in (5b).

The range of possible values for j is left unconstrained. Typically, it will be the speaker, since we generally take what Lasersohn calls an ‘autocentric’ view, but it may be another individual if we adopt an ‘exocentric’ view. An example of the latter is an attitude report. In (7a), for example, Lasersohn argues that it is plausible to treat the individual whose tastes are relevant as John rather than the speaker, a view that is supported by the felicity of the continuation in (7b).⁶

7a. John thinks that this cake is tasty.

7b. ...but I don’t like it.

Stephenson (2007a, 2007b) proposes an analysis in the spirit of Lasersohn’s, but with the following modifications.

8a. PPTs always have an internal argument, which may be either overt (rendered by a PP headed by *to* or *for*), or covert.

8b. If covert, the internal argument is either the judge-dependent element PRO_j , or an indexical element *pro* whose value is supplied by some salient individual in the discourse.

8c. A norm of assertion permits a speaker to assert a sentence of form x is P , for P a predicate of personal taste, just in case x is P to the speaker. In addition, the judge parameter is bound in the scope of an attitude predicate, so that in effect, the value of the judge parameter is the speaker or the attitude holder, depending on the position of the PPT.

8d. Since propositions are sets of world-individual pairs, a common ground is also modelled as such a set. Where a PPT takes PRO_j as its internal argument, the content of the resulting sentence is only added to the common ground if for every conversational participant j , setting the value of the judge parameter to j yields truth.

Taste predicates, then, are taken to be two place predicates. (9) is a sample lexical entry.

9. $\llbracket \text{tasty} \rrbracket^{c; w; j} = \lambda x \lambda y. y$ is tasty to x in w

(9) does not introduce judge-dependency in the strict sense of dependency upon the value of the judge parameter. Instead, judge dependency is implemented by appeal to an element that is present in the syntactic structure, PRO_j , which has the following semantics.

10. $\llbracket PRO_j \rrbracket^{c; w; j} = j$

⁶ This test is based on the assumption that one cannot consistently describe something as tasty and also say that one dislikes it. An anonymous reviewer points out that *attractive* does not appear to behave this way, as shown by the felicity of (i).

(i) Mary is attractive but I don’t like her.

We take it that this divergence does not reflect a difference in the semantics of *tasty* and *attractive*, but rather a difference in the circumstances under which people are prepared to say that they like a foodstuff versus a person: typically I will say that I like a foodstuff if I find it tasty, but I may find someone attractive but nonetheless dislike her as a person.

The motivation for providing the option of letting the internal argument of a PPT be *pro* rather than PRO_j is the existence of cases where the individual whose tastes are relevant to the value of the taste statement is not the speaker or attitude holder. (11) is such a case.⁷

11. *Context: John sees his cat tucking into a new brand of cat food.*

John: I think the cat food is tasty. / The cat food must be tasty.

John's utterance in this context does not report the unlikely scenario where John has tried the cat food and found it tasty, but rather the plausible one where he observes on the basis of his cat's behavior that she must like the food. Given the cat's salience in the discourse, it is a suitable internal argument for *tasty*, and hence we derive (11)'s intuitive meaning:

12. $\llbracket \text{I think the cat food is tasty } pro_i \rrbracket^{ci, w, j} = 1$ iff
AUTHOR(c) thinks the cat food is tasty to $g_c(i) = 1$ iff
John thinks the cat food is tasty to the cat.

Let's evaluate Stephenson's proposal first, before returning to Lasersohn's. Arguably, one advantage of her system over Lasersohn's is that insofar as it assumes that an autocentric perspective is always taken, it is more constrained than its predecessor.⁸ That is, wherever a PPT receives a judge-dependent interpretation (whenever it takes PRO_j as its internal argument), there is only one option for what the judge can be, as opposed to the array of options made available in Lasersohn's system. However, a consequence of this move is that Stephenson must allow for the possibility of a PPT taking *pro* as its internal argument, lest examples such as (11) fail to be generated. Since *pro* is a null pronominal, it will only be available when the context supplies a salient individual to serve as a referent. Beyond this, however, Stephenson has relatively little to say about how the distribution of PRO_j and *pro* is determined, a problem to which we shall return.

A second respect in which Stephenson attempts to improve upon Lasersohn's proposal is by accounting for what she calls the 'immediateness requirement', illustrated in (13). Where PRO_j is embedded below more than one attitude predicate, it must denote the immediately higher attitude holder. Intuitively, (13) describes Mary's belief state as entailing that John finds the cake tasty. It does not mean that Mary thinks that John thinks that she finds the cake tasty.

13. Mary thinks that John thinks that this cake is tasty.

To show how this might be accounted for within Stephenson's system, we need a semantics for *think*. Stephenson follows Lewis (1979) in treating attitude predicates like *think* as quantifiers over centred worlds – world-individual pairs $\langle w, x \rangle$ such that it is compatible with what the attitude holder believes for her to be x in w .⁹ An attitude predicate must

⁷ The example is due to Kai von Fintel (p.c. to Stephenson).

⁸ Because in Stephenson's system the judge parameter is bound by the attitude holder, the behavior of PPTs in the scope of attitude predicates does not pose a problem for the view that an autocentric perspective is always taken. This is a departure from Lasersohn's view that the behavior of PPTs below attitude predicates arises as a consequence of their being interpreted relative to an exocentric context.

⁹ Lewis' proposal was intended to account for attitudes de se. We discuss the relationship between attitudes de se and the semantics of PPTs in the concluding section.

therefore compose with a constituent that denotes a set of centred worlds, type $\langle s, \langle e, t \rangle \rangle$. Stephenson points out that a consequence of the addition of the judge parameter is that clauses now have just this type; the clausal complement of *think*, for example, expresses a set of world-judge pairs. *Think* itself is a universal quantifier over such pairs, so that the world and judge parameters are effectively bound. (14) is the lexical entry.

14. $\llbracket \text{think} \rrbracket^{c; w, j} =$

$\lambda p_{\langle s, \langle e, t \rangle \rangle} \lambda x_e$. For every $\langle w', j' \rangle$ such that it is compatible with what x believes at $\langle w, j \rangle$ for x to be j' in w' , $p(w')(j') = 1$

Let's give the truth conditions for (13), working upwards from the lowest clause to the highest.

15a. $\llbracket \text{This cake is tasty PRO}_j \rrbracket^{c; w, j} = 1$ iff this cake is tasty to j in w .

15b. $\llbracket \text{John thinks that this cake is tasty PRO}_j \rrbracket^{c; w, j} = 1$ iff

for every $\langle w', j' \rangle$ such that it is compatible with what John believes at $\langle w, j \rangle$ for John to be j' in w' , $\llbracket \text{this cake is tasty PRO}_j \rrbracket^{c; w', j'} = 1$ iff

for every $\langle w', j' \rangle$ such that it is compatible with what John believes at $\langle w, j \rangle$ for John to be j' in w' , this cake is tasty to j' in w' .

15c. $\llbracket \text{Mary thinks that John thinks that this cake is tasty PRO}_j \rrbracket^{c; w, j} = 1$ iff

for every $\langle w', j' \rangle$ such that it is compatible with what Mary believes at $\langle w, j \rangle$ for Mary to be j' in w' , $\llbracket \text{John thinks that this cake is tasty PRO}_j \rrbracket^{c; w', j'} = 1$ iff

for every $\langle w', j' \rangle$ such that it is compatible with what Mary believes at $\langle w, j \rangle$ for Mary to be j' in w' , it is the case that for every $\langle w'', j'' \rangle$ such that it is compatible with what John believes at $\langle w', j' \rangle$ for John to be j'' in w'' , this cake is tasty to j'' in w'' .

(15c) says roughly that Mary thinks that John thinks that the cake is tasty to him, which is what we want. There is no possibility of PRO_j resolving to the higher attitude holder Mary, because the clause in which PRO_j occurs denotes a set of centred worlds that are included in John's belief state, not Mary's. In effect, PRO_j – or more precisely the judge parameter from which it gets its value – is bound by the lower attitude predicate.

One problem with Stephenson's account is that the availability of *pro* means that the reading on which *tasty* is construed as 'tasty to Mary' is not in fact ruled out. Mary is clearly highly salient in (13); hence where the internal argument of *tasty* is an overt pronoun, it can easily be construed as referring to Mary (16).

16. Mary thinks that John thinks that this cake is tasty to her.

Moreover, while (13) makes a convincing case that there is an immediacy requirement at work in the interpretation of PPTs, one can construct examples where the requirement appears not to hold. In (17) we employ Stephenson's tactic of using an item that the attitude holder is unlikely to have tasted – cat food once again.

17. The cat thinks that John thinks that the cat food is tasty (since he keeps buying it for her).

Suppose that cats are capable of having beliefs, of constructing theories about other minds etc. (17) can be construed as meaning that the cat thinks that John thinks that the cat finds the cat food tasty. Given the way we have set things up so far, there are three possible ways of interpreting the data in (13) and (17). One is to deny that the immediacy requirement holds, and to say that PRO_j need not refer to the immediately higher attitude holder. The second is to say that the immediacy requirement holds, but in (17) the internal argument of *tasty* is *pro*, whose index is assigned as its value the individual denoted by *the cat*. The third is to say that in the unmarked case, PRO_j and not *pro* is the internal argument of a predicate of personal taste. Moreover, the immediacy requirement holds. Where contextual considerations render it implausible that the judge should be the internal argument of the PPT, *pro* is inserted as a last resort. Let's consider each of these three options in turn.

The first option, to abandon the immediacy requirement, is not very promising. Consider (13) again, repeated as (18).

18. Mary thinks that John thinks that this cake is tasty.

Not only is this sentence most naturally construed with *tasty* meaning *tasty to John*, it is very difficult to construe it as meaning *tasty to Mary*. We have argued that (17) constitutes an example where the internal argument of the PPT can be construed as a higher attitude holder, but abandoning the immediacy requirement wholesale would fail to explain why such cases are exceptional, and why they only reveal themselves when contextual considerations exclude the possibility of the closest attitude holder serving as the internal argument.

We can quickly rule out the second possibility – that the immediacy requirement holds, but *pro* is also available – for similar reasons. Unless something further is said about why the option of inserting *pro* is not readily available, the system will overgenerate.

The third option is the most promising. The immediacy requirement holds and the unmarked option is to insert PRO_j rather than *pro*. In fact, since we are discovering that the postulation of *pro* alongside PRO_j leads to an overgeneration problem, let us state this in somewhat more neutral terms: the immediacy requirement holds, but an exocentric interpretation is available as the marked option, leading to what look like violations of the immediacy requirement. Without constraining the system further, Stephenson's proposal cannot explain why this should be, but the proposal to be developed in section 3 will offer a principled account of this observation.

In the meantime, let's consider more data that call into question the notion that as long as a suitable referent for *pro* is salient, either PRO_j or *pro* can be freely inserted. Firstly, we consider cases in which *pro* does not appear to be available. According to Stephenson, wherever *pro* is inserted, the Experiencer argument of the PPT is some contextually salient individual who need not be the speaker. When the option of inserting *pro* is available, ...*but I don't like it/ ...but I wouldn't like it* ought to be felicitous continuations. In these examples we modify the subject of the PPT with the relative clause *that Mary and I ate* in order to make salient an individual, Mary, who might serve as a value for *pro*.

19. The cake that Mary and I ate was tasty. # But I didn't like it.
20. I found the cake that Mary and I ate tasty. # But I didn't like it.
- 21a. I am glad that the cake that Mary and I ate was tasty. # But I didn't like it.
- 21b. I said honestly that the cake that Mary and I ate was tasty. # But I didn't like it.
- 21c. I know that the cake that Mary and I ate was tasty. # But I didn't like it.
- 21d. I remember that the cake that Mary and I ate was tasty. # But I didn't like it.
- 21e. I think that the cake that Mary and I ate was tasty. # But I didn't like it.
- 21f. I pretended that the cake that Mary and I ate was tasty. # But at least *I* liked it.

Stephenson considers examples like (20) involving *find* and argues that the impossibility of inserting *pro* as the internal argument of the PPT is due to *find* introducing an entailment that the subject is basing her opinion on direct evidence.¹⁰ We find her semantics for *find* convincing, however a similar case cannot be made for matrix sentences like (19). Nor would such an argument seem plausible for the cases in (21), which involve attitude predicates that do not carry the evidential meaning component associated with *find*.

One might respond that the explicit statement that the speaker ate the cake in (19-21) makes it difficult to exclude the speaker from the internal argument of *tasty*. But if PRO_j and *pro* are both freely available (given a context that makes salient a suitable referent for *pro*), then we should expect it to be just as difficult for the speaker to deny that Mary liked the cake as it is for the speaker to deny that *she* liked the cake. This expectation is not borne out.

22. The cake that Mary and I ate was tasty. But she didn't like it.
23. I found the cake that Mary and I ate tasty. But she didn't like it.
- 24a. I am glad that the cake that Mary and I ate was tasty. But she didn't like it.
- 24b. I said honestly that the cake that Mary and I ate was tasty. But she didn't like it.
- 24c. I know that the cake that Mary and I ate was tasty. But she didn't like it.
- 24d. I remember that the cake that Mary and I ate was tasty. But she didn't like it.
- 24e. I think that the cake that Mary and I ate was tasty. But she didn't like it.
- 24f. I pretended that the cake that Mary and I ate was tasty. But at least *she* liked it.

A challenge for the theory of predicates of personal taste is therefore to account for the asymmetry between the infelicity of the continuations in (19-21) on the one hand, and the felicity of the continuations in (22-24) on the other. Simply assuming that either PRO_j or *pro* can be freely inserted (*modulo* salience) is not sufficient.

Now let's think about cases that give rise to what Lasersohn and Stephenson would call exocentric interpretations. One way of constructing such a case is to set up a context where it seems plausible that the opinion of an individual other than the speaker or attitude holder is being reported.¹¹ This seems to be true of (25a) below, which out of the blue conveys that

¹⁰ For more on the semantics of *find* and its interaction with predicates of personal taste, see Saebo (2009).

¹¹ We henceforth use 'agent' as a catch-all term for the speaker and the (most local) attitude holder. It should not of course be confused with the thematic role, 'Agent'.

Mary ate the cake and the speaker did not. On the other hand, it does not commit the speaker to not having eaten any of the cake, as shown by the felicity of (25b).

In a context where Mary ate the cake and the speaker did not, the internal argument of *tasty* in (25a) must be *pro* by Stephenson's lights. But suppose that after asserting (25a), the speaker says something that commits her to having also eaten some of the cake. If the internal argument of *tasty* is *pro*, no contradiction should arise if the speaker goes on to say that she didn't find the cake tasty. But this is not what we find: consider the infelicity of (25c). Moreover, suppose the speaker didn't try any of the cake, but she considers how things would have been if she had. Suppose further that she thinks that, unlike Mary, she wouldn't have liked the cake. Stephenson's theory predicts that a continuation expressing this should be felicitous, but this is again contrary to what we find (25d). (26) enables us to sharpen our intuitions by comparing them with cases where *tasty* has an overt internal argument. If *pro* was inserted in (25a), with its index assigned the individual Mary as its value, there should be no difference between our judgments of (25c) and (26c), and (25d) and (26d), but this is not the case.

- 25a. The cake that I made and Mary ate was tasty.
- 25b. ...we each ate half of it.
- 25c. #...but I didn't like it.
- 25d. #...but I wouldn't have liked it.

- 26a. The cake that I made and Mary ate was tasty to her.
- 26b. ...we each ate half of it.
- 26c. ... but I didn't like it.
- 26d. ... but I wouldn't have liked it.

Another domain in which we might look for evidence for or against Stephenson's analysis concerns factive attitude predicates. (27) asserts that John knew that the cake was tasty, and presupposes that the cake was tasty. As ever, this begs the question, tasty to whom?

- 27. John knew that the cake was tasty.

Let's assume that the internal argument of *tasty* is the same in both the presupposed and asserted components. Suppose also that this internal argument is PRO_j , which refers to John. Then (27) asserts that John discovered that the cake was tasty to him, and presupposes that the cake was tasty to him. It should be consistent to follow up with (28), but it is not.

- 28. # ...but I didn't like it.

Evidently, the factive presupposition of (27) commits the speaker to finding the cake tasty. One might explain this by saying that in determining the truth conditions, PRO_j refers to the attitude holder, but in determining the presupposition, it resolves to the speaker. This might be reasonable, since presuppositions are speaker commitments, and there is no attitude predicate in the presupposed component to bind the judge parameter. But then one must explain (29), which surely does not presuppose that the speaker finds the cat food tasty.

- 29. The cat discovered that the cat food was tasty.

The obvious move is to let the internal argument of *tasty* be *pro*. But this brings us back to the question why *pro* cannot be the internal argument of *tasty* in (27). The conclusion is that exocentric readings are only available as a last resort, say if the autocentric reading is implausible. Stephenson's system does not account for this in a principled manner.

At this point, one might suspect that the challenges identified for Stephenson's account can be met by adding constraints on the distribution of PRO_j and *pro*. We counter this by highlighting a case that suggests that the very attempt to formulate a crisp distinction between autocentric and exocentric interpretations is on the wrong track. Our evidence involves epistemic *must*.

Suppose that I am watching Mary enjoying a piece of cake, and infer from her behavior that the cake must be tasty. I might say,

30. The cake must be tasty, because Mary seems to be enjoying it.

What exactly is the conclusion that I express in (30)? According to Stephenson's proposal, there are two possibilities: either I have expressed the expectation that I would find the cake tasty if I tried it, or I have observed that the cake is tasty to Mary. The two possibilities arise via insertion of PRO_j and *pro* respectively. Suppose that in (30) *tasty* takes *pro* as its internal argument. This would fail to predict the existence of a reading which entails that in worlds in which the speaker eats the cake, she finds it tasty. That there is such a reading is shown by the incoherence of (31a), whose oddness can be compared to the acceptability of the minimally different (31b), where an overt PP referring to Mary is introduced.

31a. # The cake must be tasty, but I wouldn't like it.

31b. The cake must be tasty to Mary, but I wouldn't like it.

Suppose therefore that, for some reason, *tasty* in (30) cannot take *pro* as its internal argument, which is instead PRO_j. Then what of (32)?

32. The cake must be tasty, but I wouldn't like it because I don't like chocolate.

My consultants find this example significantly better than (31a), which for the defender of Stephenson's view must indicate that the interpretation of (31a) just proposed is at least not the only available one. The only other option is to let *pro* be the internal argument of *tasty*, but as we have just seen it is then mysterious why this option is not available to rescue (31a). There seems to me to be no way out of this quandary but to conclude that (30) receives a 'mixed' autocentric and exocentric interpretation, which carries commitments about both the speaker's and Mary's tastes. (32) shows that the speaker's tastes can be excluded from the mix only by stating explicitly why her tastes are not relevant. One of our tasks is to account for the 'mixed' interpretation, and to give a principled explanation of why the continuation '*...because I don't like chocolate*' improves (31a). The proposal described in section 3 meets these challenges.

Until now, our arguments have mostly focused on Stephenson's proposal. Let's step back and consider how Lasersohn's theory fares against the objections raised. Lasersohn says that in the unmarked case, we adopt an autocentric perspective; that is, sentences are evaluated

with respect to indices where the judge is the speaker. One might argue, therefore, that it is unsurprising that exocentric interpretations are only available given a special context, such as one in which we are talking about food items that humans don't eat. In other cases, Lasersohn's proposal suffers from the same problems as Stephenson's does. In particular, it would have difficulty accounting for the existence of 'mixed' autocentric and exocentric readings that we have discussed in the case of PPTs embedded below epistemic modals. Moreover, to revert from Stephenson's system to one in which autocentric or exocentric readings arise as a consequence of pragmatic factors (what kind of context we are in) rather than structural factors (which covert element is inserted) comes at a price: we would have little hope of providing an account of the behavior of taste predicates in attitude reports. These cases require us to (i) rule out the possibility of the judge resolving to the speaker (unless the speaker happens to be the attitude holder), and (ii) rule out the possibility of the judge resolving to a non-local attitude holder. Here we think Stephenson's insight is essentially correct. However, the problem is not merely that a structural account of how the judge is determined is to be preferred to a pragmatic one; we cannot rule out (i) and (ii) by brute stipulation, since such interpretations are available in exceptional cases such as (17) above.

We have described the problems we think are associated with existing relativist accounts of the semantics of taste predicates. One could of course investigate whether it is possible to have a theory that invokes the judge parameter but that does not suffer from these drawbacks. This is not the strategy that we will take. Rather, we shall offer a diagnosis of what has gone wrong with previous analyses, and take it as a starting point for a novel proposal. Happily, this proposal does not require complicating the inventory of semantic types by adding an individual parameter to the evaluation index.

3. A new account

3.1 Diagnosis: Part 1

A critical problem of Stephenson's theory is that it overgenerates. We encountered a series of cases where one might expect that *pro* could serve as the internal argument of a PPT, but our intuitions suggest that an exocentric reading is in fact unavailable.¹² Relatedly, Stephenson's account fails to explain why in the unmarked case, the PPT receives a first person oriented interpretation. By this we mean that taste statements typically carry entailments about the agent's opinion. Focusing on the case of *tasty*, this requirement can be overridden if (i) the agent has not tasted the subject of *tasty* – in which cases the agent's opinion in those counterfactual worlds in which she has tried the food item are relevant; or (ii) there is something about the agent that makes her an unsuitable candidate for being an individual to whom the item is tasty – say because she doesn't like chocolate if she is talking about chocolate cake, or because she is not a member of the appropriate species in the case of cat food. In general, we can say that the agent's tastes are relevant unless something about the context renders them irrelevant – not having tasted the item, perhaps, or not being among its 'target audience'. We might know that we are in such a context because of

¹² For continuity, we use the term 'exocentric', but strictly speaking this kind of talk should be confined to theories that employ the judge parameter – the very move that we propose to reject. The intended meaning should be intuitively clear: 'exocentric' interpretations involve the opinion of a third party rather than the agent.

linguistic information such as use of epistemic *must*, or because of extralinguistic information such as the knowledge that people don't eat cat food.

The view that PPTs receive a first person oriented interpretation unless the agent's tastes are known to be irrelevant accounts for the problematic data discussed in the last section.

Firstly, it captures the observation that an example like (25c), repeated here, is contradictory.¹³

33. # The cake that I made and Mary ate was tasty, but I didn't like it.

If we know that I ate the cake, and there is no reason to think that I am excluded from its target audience, then the first conjunct of (33) entails that I liked the cake. Moreover, in a context in which I have not tried the cake, the first person oriented interpretation of a PPT persists in modal form: (34) (repeated from (31a)) entails that the speaker likes the cake in those counterfactual worlds in which she tries it.

34. # The cake must be tasty, but I wouldn't like it.

Yet if we know that the speaker is not among the target audience, this entailment goes away. Hence (34) is improved by the *because*-clause in (35a), while our knowledge about the diets of humans is sufficient to make (35b) acceptable.

35a. The cake must be tasty, but I wouldn't like it because I don't like chocolate.

35b. The cat food must be tasty, but of course I wouldn't like it.

Here is another way in which the first person orientedness of PPTs is manifested. In order to assert that *x* is *P* for some taste predicate *P*, one typically must have direct sensory experience of the relevant kind on the basis of which to judge whether *x* is *P*. For *tasty*, for example, I must have tasted the object I am talking about, or at least have tasted an instance of the same kind of object. To assert that shortbread is tasty I must have tasted shortbread. If I have good reason to believe that shortbread is tasty, say because a reliable expert has told me so, I might say, *Apparently, shortbread is tasty*, but not, *Shortbread is tasty*. If I have good reason to believe that shortbread is tasty because I like all of the ingredients in it, I might say, *Shortbread must be tasty*, but not *Shortbread is tasty*. Someone who heard me describe shortbread as tasty and later learned that I had never tried it would be entitled to feel misled. This constraint on the assertibility of taste statements is strikingly strong: it seems that the conditions under which I am entitled to form the belief that something is tasty (such as when I have read the ingredients list on the packet) are typically weaker than those under

¹³ An anonymous reviewer notes that (33) may not strictly speaking be a contradiction, but is perhaps better described as pragmatically odd. This intuition presumably reflects the fact that contextual factors enter into determining whether the PPT receives a first person oriented interpretation. In 3.3 we capture this by positing in the semantic representation of taste statements free variables whose value is contextually supplied. As we shall see, the consequence of supplying values for these variables is typically a first person orientation, but not always. We shall stay with the term 'contradiction' because it reflects our own intuitions better than 'pragmatically odd', but the line of argument pursued here will not be undermined if the reader chooses to substitute one term for the other.

which I am entitled to assert that it is tasty.¹⁴ ¹⁵ This is not only a property of predicates of taste: I can only say that John seemed tired yesterday if yesterday I had some contact with him that gave me this impression. If my information comes from Mary's description of an encounter with him, I should say, *Apparently John seemed tired yesterday*. By contrast, there are a wide range of circumstances, including some in which I have had nothing to do with John personally, in which I may say *John was tired yesterday*. By the same token, I need not have had direct experience with shortbread in order to be able to say that it is a Scottish foodstuff.¹⁶

This parallel between the behavior of predicates of personal taste and *seem* suggests that the usage condition we have identified applies to predicates that take Experiencer arguments. We might conjecture, for example, that PPTs and *seem* take covert Experiencer arguments that in the unmarked case include the speaker, and that there is a presupposition of direct perceptual experience built into the meaning of the Experiencer theta role. This is not quite enough though: we need to explain why this meaning component apparently disappears in the presence of an epistemic modal even though it is a hole for presupposition projection. An advantage of our account is that it provides an explanation for these data, at least as far as predicates of personal taste are concerned. Section 4.4 provides this explanation, but first let's move on to the second part of our diagnosis.

3.2 Diagnosis: Part II

One aspect of Lasersohn's and Stephenson's accounts that we have not considered is their failure to account for the possibility that a PPT may take a generic internal argument, so that *tasty* is construed as meaning roughly 'tasty to people in general'. We argue that such meanings need to be taken seriously, however this is not a direct criticism of Lasersohn and Stephenson's analyses, since both are compatible with generic interpretations being available in principle. In Lasersohn's system, one would let the judge parameter be set to an arbitrary individual, and in Stephenson's, one would admit the possibility that the covert internal argument of the PPT could be generic, as she herself acknowledges (Stephenson 2007a: 56, fn. 24). However, Lasersohn and Stephenson do not consider the possibility of generic readings as seriously as they might, and consequently a potential solution to the puzzle is

¹⁴ The application of this constraint is pragmatically conditioned. Irene Heim (p.c.) points out that if I am learned about a culture in which an exotic insect is eaten, but have never sampled the culture (nor the insect) for myself, I might still say, 'The insects they eat are sweet and delicious'. We take it that in this example the context is rich enough to exclude the possibility that the speaker has tasted the insect, without employing overt linguistic material in order to do so.

¹⁵ An anonymous reviewer speculates that the usage condition may not apply to *attractive*, pointing out that (i) is acceptable in a context in which the speaker has not seen John's wife.

(i) John married an attractive woman.

It may be that this is because the PPT is used attributively rather than predicatively here: we think that (ii) is infelicitous in the same context.

(ii) John's wife is attractive.

We leave it to future research to investigate the behavior of PPTs *qua* modifiers. One observation that will need to be accounted for is that in attributive uses below attitude predicates, PPTs express the speaker's opinion rather than the attitude holder's (Jacopo Romoli, p.c.); hence (iii) is a coherent statement.

(iii) John thinks that this tasty cake is disgusting.

Presumably in (iii) *this tasty cake* is construed de re, and hence outside the scope of the attitude predicate. The details will have to be worked out on another occasion, however.

¹⁶ We are grateful to an anonymous reviewer for prompting us to think about the domain of this usage condition beyond predicates of personal taste.

overlooked. We propose that an analysis based upon a particular type of genericity can handle the problematic cases we have discussed, capturing our observations about the first person oriented interpretation of PPTs, and liberating us from the problem of overgeneration.¹⁷

Firstly, let's consider why one might think that a generic internal argument for a PPT is available. By 'generic internal argument' we mean a generic individual meaning something like *people (in general)*, or an arbitrary individual. Part of the work of the next section is to make this notion precise. We are also being somewhat sloppy in using the term 'internal argument'. We employ it to describe the role of *x* in a paraphrase of a taste predicate *P* as meaning 'P to *x*'. In a moment we make a commitment to the view that PPTs are always two place predicates, regardless of whether the internal argument is overt, but we have not yet argued for this.

Our first case is an example from Keshet (2005).

36a. *John*: The chilli is tasty.

36b. *Mary*: No, the chilli is not tasty.

36c. *John*: Wait, how can you say that? The beans are bursting with flavor and the meat is cooked to perfection!

36d. *Mary*: Well, it's too spicy, for one thing.

36e. *John*: Let's ask someone else...

[Keshet 2005: 2, ex 5]

Keshet points out that a dialogue such as (36) is perfectly rational. That is, it is rational to give evidence on which to base one's opinion about whether the chilli is tasty, and it is rational to ask a third party for theirs. This suggests that an attempt is being made to establish whether the chilli is tasty to people in general, rather than to simply exchange information about whether one's interlocutors like it.

Our second example concerns arbitrary PRO. Where an infinitive or gerund whose subject is PRO_{arb} is embedded below a PPT, the value of the internal argument of the taste predicate covaries with the value of PRO_{arb}. That is, the internal argument is obligatorily construed generically.

37. It is fun [PRO_{arb} to ride the rollercoaster].

We conclude that PPTs can be construed as having a generic internal argument. However, even in cases where this seems clear, PPTs nonetheless receive a first person oriented interpretation. For instance, the examples in (38) are not felicitous continuations of (37).

38a. #...but I wouldn't like it.

38b. #...but I don't like it.

38c. #...but I didn't like it.

¹⁷ This account owes much to Moltmann's (2006, 2010a) work on generic *one*. Keshet (2005) noticed the parallel between this expression and PPTs, and Laurence Horn appears to have had a related idea (p.c. to Stephenson).

Crucially, there is an asymmetry between first person and third person here: (39) provides acceptable variants on (38).

- 39a. ...but Mary wouldn't like it.
- 39b. ...but Mary doesn't like it.
- 39c. ...but Mary didn't like it.

Notice that this asymmetry is not apparent with garden variety generic statements. Out of the blue, (40a) conveys that the speaker is not a linguist, and (40b) that Mary is not a linguist. In a context in which it is known that the speaker/Mary are linguists, these sentence are somewhat degraded with respect to (41) and (42).

- 40a. (?) Linguists write their papers using LaTeX, but I use Microsoft Word.
- 40b. (?) Linguists write their papers using LaTeX, but Mary uses Microsoft Word.

- 41a. Most linguists write their papers using LaTeX, but I use Microsoft Word.
- 41b. Most linguists write their papers using LaTeX, but Mary uses Microsoft Word.

- 42a. Linguists usually/typically/generally write their papers using LaTeX, but I use Microsoft Word.
- 42b. Linguists usually/typically/generally write their papers using LaTeX, but Mary uses Microsoft Word.

We revisit the significance of the contrast between (40) and (41-2) in section 4.4, but the important point for now is that the status of the examples in (40) is not altered by reference to the third person instead of the first person. By contrast, we take it that the first/third person asymmetry exhibited by predicates of personal taste constitutes evidence that they receive a first person oriented interpretation. This is the heart of the puzzle: a tension between the first person oriented nature of the interpretation of PPTs, and the possibility that they can be construed as making claims that extend beyond the agent. In the next section, we argue that we can make sense of this if we assume a peculiar type of genericity whose existence is already attested elsewhere in natural language.

3.3 The core of the proposal

(43) states how we propose to treat predicates of personal tastes.

43. *Proposal* Predicates of personal taste such as *tasty* are used to make statements about whether something is tasty to people in general, based on first person experience.

The idea is this: when I say *This cake is tasty*, I commit myself to finding the cake tasty. That's why I cannot coherently follow up with, ...*but I don't like it*. But this is not all I do. I also generalize beyond my own experience to the likely experience of anyone with whom I

empathize who might eat the cake, and claim that they would find it tasty too. This is why PPTs are interpreted generically, at least when there is no overt PP.¹⁸

Our first step is to consider whether taste predicates always take an internal argument. An advantage of supposing that they do is that it gives us a means to explain the generalization discussed in 3.1 that typically one may only describe something as tasty if one has tasted it. To see this, consider what happens when there is an overt internal argument. Some examples:

- 44a. This cake is tasty to John. # But he hasn't tried it.
- 44b. This cake is tasty to me. # But I haven't tried it.
- 44c. This cake is tasty to you. # But you haven't tried it.
- 44d. This cake is tasty to us. # But we haven't tried it.

The infelicity of the continuations in (44) shows that use of a sentence of form *x is tasty to y* is felicitous only if *y* has the relevant kind of direct perceptual experience of *x* - that is, if she has tasted it. We can model this requirement as a presupposition on the internal argument of the PPT. (45) is the lexical entry for *tasty*.¹⁹

45. $\llbracket \text{tasty} \rrbracket^{c:w} =$

$\lambda x \lambda y: x$ has direct perceptual experience of the relevant kind of y in w . y is tasty to x in w

Projection tests verify that a presuppositional treatment is appropriate. The presupposition that John has tasted the cake survives in the environments displayed in (46).

- 46a. The cake isn't tasty to John.
- 46b. Is the cake tasty to John?
- 46c. If the cake is tasty to John, I'm glad.

Suppose that when the internal argument of the PPT is covert, it includes the agent. This naturally accounts for the observed constraint on when one is entitled to describe something as tasty, which can be seen as a consequence of the presupposition associated with the PPT, along with an auxiliary assumption (to be justified and refined) that where a taste predicate does not take an overt internal argument, it takes a covert one that includes the agent.

¹⁸ We confine our attention to predicative uses of PPTs. Moltmann (2010b) claims that *tasty* is not interpreted generically in examples such as *John considers the cake tasty*. We postpone discussion of how our analysis might handle this case until another occasion.

¹⁹ Of course, for any given object there are multiple kinds of direct perceptual experience that we may have of it: we can see, touch, or smell a cake in addition to tasting it. In (45), the modifier *of the relevant kind* is intended to limit attention to the perceptual experience that in ideal conditions constitutes the most reliable basis for determining whether something is tasty: so long as our tastebuds are in good working order, this is the experience of tasting. It seems that for any given PPT *P* we have stable intuitions about what the most reliable way of determining whether something is *P*: for *tasty*, it is eating the object, for *fun*, it is participating in the activity, for *beautiful*, it is seeing the individual, and so on. We remain neutral as to whether these intuitions reflect world knowledge or information encoded in the semantics of each predicate. Depending on which view one adopts, one might either leave it to context to determine what the relevant kind of direct perceptual experience is, or alternatively flesh out the meaning of *of the relevant kind* in the lexical entry of each PPT.

The question that now arises is how this presupposition can apparently be defeated via an expression that signals that the agent lacks direct evidence such as epistemic *must*.²⁰ We saw in 3.1 that a speaker who announces, *Shortbread must be tasty* need not have tasted shortbread. The problem is sharpened by comparing the felicity of this statement with the infelicity of (47) when uttered by someone who has never tried shortbread.

47. # Shortbread must be tasty to me!

The reason for the infelicity of (47) is clear: we know that the speaker has not tasted shortbread. Yet we have said that with PPTs in matrix position with no overt internal argument, there is a covert internal argument that includes the speaker. One of the challenges that lies ahead of us is therefore to account for the difference between (47) and its counterpart with no overt Experiencer argument. We show how this can be done in section 4.4.

The basic components of our analysis are now in place. The problem with previous accounts is their failure to recognise that the interpretation of PPTs (i) is first person oriented and (ii) involves genericity. Putting the pieces together, we can say that the interpretation of PPTs involves a peculiar type of genericity that is first person oriented. In fact, it has been argued that this species of genericity is needed in order to account for the semantics of generic *one* in English (Moltmann 2006, 2010a), exemplified in (48).

48. One can see the picture from the entrance. [Moltmann 2006: 258, ex 1]

For our purposes, the crucial components of Moltmann's analysis of generic *one* are (i) the notion of genericity, implemented via the generic operator, and (ii) an *identify with* relation *I*, where *I*(x,y) just in case x identifies with y. In 3.4 we show that the presence of the generic operator in the LF of a sentence containing a PPT is independently required. In the meantime, let's consider the *identify with* relation.

The idea is roughly this. To say that the cake is tasty is to say that the cake is tasty to every (contextually restricted) individual with whom I identify. The *identify with* relation is intended to model a notion of empathy, and is therefore reflexive: I always empathize with myself. In normal conversation we empathize with our interlocutors, but we may also empathize with other individuals, such as our pet cat when we consider whether or not it likes the new cat food we have bought it, or even cats in general. There is more to say on this topic, but this suffices for our purposes. In the next section, we provide a principled explanation of why PPTs with no overt internal argument receive a generic interpretation.

3.4 Accounting for genericity

In this section, we show that predicates of personal taste are individual level predicates (ILPs) in the sense of Carlson (1977), and claim that it follows from this that the variable introduced by a PPT's covert internal argument is bound by the generic operator, GEN. The generic interpretation of PPTs is derived without any additional assumptions.

²⁰ We assume that von Stechow and Gillies (2010) are correct that epistemic *must* signals that a claim is being made on the basis of indirect evidence.

We begin by considering how PPTs fare against the diagnostics that have been proposed to determine whether a given predicate is individual level.²¹ Intuitively, an individual level predicate is one that expresses a permanent property of individuals, while a stage level predicate expresses a property that one might hold at some times but not at others. While I might be sick today but not next week, if I am tall then I will likely remain tall. I employ *sick* as a prototypical stage level predicate (SLP), and *tall* as a prototypical ILP.

It is known since Milsark (1974) that SLPs can appear in postnominal position in existential *there* constructions, but ILPs cannot. (49) provides the relevant contrast.

- 49a. There were people sick.
- 49b. *There were people tall.

Tasty and *fun* both pattern like *tall* in this respect:

- 50a. *There were cakes tasty.
- 50b. *There were games fun.

Secondly, with bare plurals and subjects of form *A NP*, SLPs give rise to existential interpretations, while ILPs yield (quasi-) universal interpretations (Milsark 1974, Carlson (1977)). Whereas all it takes for (51a,b) to be true is for there to be a sick man and some sick people respectively, (52a,b) attribute tallness to men in general and to people in general.

- 51a. A man is sick.
- 51b. People are sick.

- 52a. A man is tall.
- 52b. People are tall.

Tasty and *fun* also give rise to universal readings, lending further support to the view that they are individual level predicates.

- 53a. A cake is tasty.
- 53b. Cakes are tasty.

- 54a. A game is fun.
- 54b. Games are fun.

Third, Carlson (1977) observed that a verb of perception such as *see* taking a small clause complement licenses a postnominal predicate within the small clause only if it is stage level.

- 55a. John has seen Mary sick.
- 55b. *John has seen Mary tall.

Like *tall*, *fun* and *tasty* are excluded from this position.

²¹ We employ six of the best-known diagnostics for the SLP/ILP distinction. Readers familiar with other tests, such as Stump's (1985) tests with free adjuncts, can verify that they also suggest that taste predicates are ILPs.

- 56a. *John has seen the cake tasty.
56b. *John has seen the game fun.

Now consider the series of diagnostics involving quantificational readings. *When*-adjuncts with a proper name as subject permit an SLP to be predicated of the subject, but not an ILP.

- 57a. When Mary is sick, everyone is unhappy.
57b. *When Mary is tall, everyone is unhappy.

Inserting *tasty* also produces ungrammaticality (58a).

- 58a. *When Nutella is tasty, Mary is happy.
58b. ?When Monopoly is fun, Mary is happy.

(58b) is in fact somewhat better than (58a). *Tasty* and *fun* also part company in another quantificational environment: that involving quantificational adverbs such as *always*. ILPs cannot be modified by a quantificational adverb when the subject is a proper name; compare (59a) and (59b). *Tasty* patterns with *tall* in this respect (60a), but *fun* patterns with *sick* (60b). We have an argument that *tasty* is an ILP, then, but we are on less certain territory with *fun*. I return to this issue shortly.

- 59a. Mary is always sick.
59b. *Mary is always tall.

- 60a. *Nutella is always tasty.²²
60b. Monopoly is always fun.

Carlson (1982) and Kratzer (1995) point out that SLPs generally tolerate modification by temporal modifiers (61a), but ILPs do not (61b). Turning to PPTs, we again find a split between *tasty*, which patterns like *tall*, and *fun* which is marginally acceptable in this environment. Let's add this to our list of problematic cases, discussion of which we postpone until a little later.

- 61a. Mary is sick this morning.
61b. *Mary is tall this morning.

- 62a. *Nutella is tasty this morning.
62b. ?Monopoly is fun this morning.

Since we can be confident that *tasty* is individual level, we initially focus on this predicate. Chierchia (1995) assumes that like all predicates, ILPs have a Davidsonian argument ranging over worlds or situations, which for ILPs is bound by the generic operator, GEN. Any syntactic structure involving an ILP must therefore include GEN, the covert counterpart of

²² There is perhaps a kind-oriented interpretation of *Nutella* that renders (60a) acceptable. This is unsurprising given familiar observations about the behavior of ILPs when predicated of kind-level subjects.

quantificational adverbs such as *always* and *usually*. As such, it is an unselective binder of free variables in its scope. We adopt the orthodox assumptions that the operator has universal force, and that its restrictor includes a two place predicate C, the ‘contextual restriction’. *John is tall* receives the LF in (63b) and the interpretation in (63c).

63a. John is tall.

63b. $[[\text{John}_i [\text{GEN } [t_i \text{ is tall}]]]]$

63c. $[[(63b)]]^{c,w} = \forall_{w'} [\text{Acc}(w,w') \ \& \ C(\text{John},w')] [\text{tall}(\text{John},w')]$

Where (i) for any two worlds w, w' , $\text{Acc}(w,w')$ iff w' is accessible from w and
(ii) for any individual x and world w , $C(x,w)$ iff w is inhabited by x and x is relevant in w .

(63c) says that in every accessible world w such that John is relevant in w , John is tall.

Returning to PPTs, we are assuming that where these lack an overt internal argument, there is a covert one. Let this argument be a restricted free variable ranging over individuals with whom the speaker identifies, much like indefinite nominals are taken to be free variables restricted by the predicate denoted by the noun phrase (Lewis (1975), Kamp (1981), Heim 1982)). As with indefinites, the free variable may be bound by the generic operator, in which case a generic interpretation is assigned. We illustrate firstly with an example involving the indefinite noun phrase *a big dog* (64).²³ We define a 3-place variant of C, C_3 such that $C_3(x,y,w)$ iff w is inhabited by x and y and x and y are relevant in w . The sentence is true just in case for every accessible world w and every individual x such that (i) John and x inhabit w , (ii) John and x are relevant in w and (iii) x is a big dog in w , John likes x in w . The restriction of the generic quantification to relevant triples allows the sentence to be true even though there might be big dogs that John doesn’t like, say because they are ferocious. C_3 excludes worlds w and individuals x such that x is ferocious in w .

64a. John likes a big dog.

64b. $[[\text{John}_i [\text{GEN } [t_i \text{ likes } \lambda x. \text{big-dog}(x,w)]]]]$

64c. $[[(64b)]]^{c,w} = \forall_{x,w'} [\text{Acc}(w,w') \ \& \ C_3(\text{John},x,w') \ \& \ \text{big-dog}(x,w')] [\text{likes}(\text{John},x,w')]$

PPTs work the same way, except that the predicate associated with the free variable in internal argument position is the predicate $\lambda x. I(\text{speaker}, x)$, where I is the relation *identify with*.

65a. This cake is tasty.

65b. $[[\text{This cake}_i [\text{GEN } [t_i \text{ is tasty } \lambda x. I(\text{speaker}, x)]]]]$

65c. $[[(65b)]]^{c,w} =$

$\forall_{x,w'} [\text{Acc}(w,w') \ \& \ C_3(\text{this cake},x,w') \ \& \ I(\text{speaker}, x)] [\text{tasty}(\text{this cake},x,w')]$

Formula (65c) says that for all accessible worlds w and all individuals x such that (i) w is inhabited by this cake and x , (ii) this cake and x are relevant in w , and (iii) the speaker identifies with x , this cake is tasty to x in w . This is the meaning of *This cake is tasty*. The accessibility relation excludes counterfactual worlds in which the relevant individuals’ tastes

²³ We consider only the generic interpretation of this sentence.

in foodstuffs are different from what they are in the actual world, in much the same way as in (64a) it excludes worlds in which John's tastes in animals differ from his actual tastes.

A feature of the analysis is that the generic interpretation of the PPT's internal argument comes for free since the responsible operator is already present in the syntax given Chierchia's view of the licensing conditions for ILPs. To apply our semantics to *fun* as well as *tasty*, therefore, we must say something about the ILP diagnostics that *fun* apparently fails. We repeat the relevant examples in (66); they involve *when*-adjuncts (66a), quantificational adverbs (66b) and temporal modifiers (66c). Each example is either grammatical or marginally acceptable, contrary to what is expected if *fun* is an ILP. On the other hand, *fun* passes the diagnostics involving existential *there*-constructions (50b), bare plural and indefinite subjects (54) and small clause complements of perception verbs (56b). Given this predicament, we must either claim that *fun* is an ILP and explain away the first class of data, or claim that it is an SLP and explain away the second class. Suppose we choose the former option.

66a. ?When Monopoly is fun, Mary is happy.

66b. Monopoly is always fun.

66c. ?Monopoly is fun this morning.

We focus on (66b), although the argument is applicable to the other two cases. Suppose, as seems intuitively correct, that (66b) can be paraphrased as *Playing Monopoly is always fun*. This would suggest that *always* quantifies over events of playing Monopoly, and that (66b) says that for every occurrence *e* of this event, *e* is fun.²⁴ For reasons discussed in Kratzer (1995), it is crucial for this test that the subject should be a rigid designator; that is, that it should have the same extension in different situations. This is the motivation for using proper names. It seems, however, that under some circumstances, *Monopoly* is interpreted roughly as *playing Monopoly*, an expression with different extensions in different situations – an observation that can be captured by treating it as denoting an event kind. This behavior reveals itself in constructions involving quantification over situations (66a,b) and in those involving reference to particular times (66c). When *fun* is predicated of a particular game of Monopoly, it behaves like *tasty* once again.²⁵

67a. *When the game of Monopoly that we are playing is fun, Mary is happy.

67b. *The game of Monopoly that we are playing is always fun.

67c. *The game of Monopoly that we are playing is fun this morning.

²⁴ We think that it is no coincidence that *fun* can be predicated of infinitival or gerund subjects such as this but *tasty* cannot:

(i) Playing/To play Monopoly is fun.

(ii) *Eating/To eat Nutella is tasty.

This is expected if the infinitives and gerunds in these examples denote event kinds and *fun*, but not *tasty*, can be predicated of such objects.

²⁵ These examples are significantly improved if we imagine that the game is taking place over an extended period of time; one thinks of correspondence games, for example. It is as though in such a context, there are multiple events of playing the present game of Monopoly, even though only one game is being played.

We conclude that we can set aside these diagnostics and maintain on the basis of the tests where the behavior of *fun* patterns with *tasty* that it too is an ILP, and that the generic interpretation follows naturally from this.

3.5 Accounting for the behavior of PPTs under attitude predicates.

So far we have said nothing about how to account for the behavior of PPTs under attitude predicates. As it stands, our proposal wrongly predicts that (68a) has the meaning in (68b).

68a. John thinks that this cake is tasty.

68b. John's belief state entails that for all worlds w and all individuals x such that (i) w is an accessible world inhabited by this cake and x , (ii) this cake and x are relevant in w and (iii) the speaker identifies with x , this cake is tasty to x in w .

We have made a mistake in our formulation of the restriction on the variable that occupies the internal argument of the PPT, repeated in (69).

69. $\lambda x. I(\text{speaker}, x)$

Let us mention one possible solution to this problem, and explain why we decide against it. We might conclude that although we disagree with aspects of Stephenson's proposal, the idea that PPTs motivate the incorporation of a judge parameter into the index is correct. We would amend the predicate in (69) accordingly to mean *is identified with by the judge* rather than *is identified with by the speaker*.

Suppose we adopt Stephenson's view that indices are world-individual pairs, and attitude predicates quantify over these pairs. *Think* would have the semantics in (14), repeated below.

70. $\llbracket \text{think} \rrbracket^{c; w; j} =$

$\lambda p_{\langle s, \langle e, t \rangle \rangle} \lambda x. \text{For every } \langle w', j' \rangle \text{ such that it is compatible with what } x \text{ believes at } \langle w, j \rangle \text{ for } x \text{ to be } j' \text{ in } w', p(w')(j') = 1$

Suppose further that the covert internal argument of a PPT contributes a variable restricted by the predicate in (71). (We borrow PRO_j from Stephenson for notational convenience.)

71. $\lambda x. I(\text{PRO}_j, x)$

The truth conditions for *This cake is tasty* and *John thinks this cake is tasty* would be as follows.

72a. $\llbracket \text{This cake is tasty} \rrbracket^{c; w; j} = 1$ iff

$\forall_{x, w'} [\text{Acc}(w, w') \ \& \ C_3(\text{this cake}, x, w') \ \& \ I(j, x)] [\text{tasty}(\text{this cake}, x, w')]$

72b. $\llbracket \text{John thinks this cake is tasty} \rrbracket^{c; w; j} = 1$ iff

For all $\langle w', j' \rangle$ such that it is compatible with what John thinks at $\langle w, j \rangle$ for John to be j' at $\langle w', j' \rangle$, $\llbracket \text{This cake is tasty} \rrbracket^{c; w'; j'} = 1$ iff

For all $\langle w', j' \rangle$ such that it is compatible with what John thinks at $\langle w, j \rangle$ for John to be j' at $\langle w', j' \rangle$, $\forall_{x, w''} [\text{Acc}(w', w'') \ \& \ C_3(\text{this cake}, x, w'') \ \& \ I(j', x)] [\text{tasty}(\text{this cake}, x, w'')]$

This cake is tasty would have essentially the same meaning as before, assuming that in matrix sentences the value of the judge parameter is the speaker. *John thinks that this cake is tasty* is predicted to be true just in case John's belief state entails that for all worlds w and all individuals x such that w is an accessible world inhabited by this cake and x , this cake and x are relevant in w , and John identifies with x , this cake is tasty to x in w . In other words, John thinks that the cake is tasty to him and to those he identifies with, which is what we want.

This shows that our semantics can be implemented within a framework where the evaluation index incorporates a judge parameter. However, it would be rather disappointing if this were the only way of implementing our proposal: nothing we have said before now requires this complication of the semantics, and so we might ask whether there is an alternative way of ensuring that the first person oriented interpretation of PPTs involves the speaker in matrix sentences, and the attitude holder in case of embedding. We think that there is, and that there are independent arguments in favor of the new assumptions that this route requires.

Suppose, with Lewis (1979), that attitude predicates are quantifiers over centred worlds: the set of world-individual pairs $\langle w, x \rangle$ such that it is compatible with what the attitude holder believes (wants, expects, etc.) that she is x in w . The clausal complement of an attitude predicate would be of type $\langle s, \langle e, t \rangle \rangle$. However, since we are proposing to dispense with the judge parameter, we revert to the traditional view that $\langle s, \langle e, t \rangle \rangle$ is the type of properties, not of propositions. That is, we side with Lewis in supposing that the complements of attitude predicates express properties. The centred world semantics for attitude predicates is intended primarily to provide an account of attitudes de se, conceived of as self-ascriptions of properties on the part of the attitude holder. In Chierchia's (1989) analysis of controlled PRO – the de se expression *par excellence* – he implements this view by postulating an abstraction operator in the left periphery of the control complement, which binds PRO.

In analogous fashion, instead of saying that the subject of the *identify with* relation is the judge, we let it be a variable bound by the most local abstraction operator. (The locality condition ensures that Stephenson's immediacy requirement is accounted for.) In unembedded cases, we could let the subject of the *identify with* relation be the speaker, so that the interpretation of the predicate introduced by the covert internal argument of a PPT would depend on whether it occurs below an attitude predicate. But this disjunctive formulation is somewhat inelegant. A more pleasing solution is to assume that the subject of the *identify with* relation is determined in the very same way in a root sentence as under embedding – on the view that we are entertaining, this would mean that it is bound by an abstraction operator. This obliges us to take the somewhat unorthodox view that in the left periphery of matrix sentences there is an abstraction operator that can bind individual variables, and hence that unembedded sentences express properties, not propositions. On the other hand, we would be in a position to retain the spirit of the Lewis/Chierchia view of attitudes de se, but to improve on it by assigning all CPs the same semantic type. On this view, regardless of its syntactic position, a CP is of type $\langle s, \langle e, t \rangle \rangle$. We illustrate this schematically below, providing a clausal structure annotated with semantic types.²⁶

²⁶ For completeness we add a second abstraction operator Op_2 that binds world variables. This is not crucial to the analysis, however, and for simplicity we omit this operator in the remainder of the paper.

73. $[_{CP} \langle s, \langle e, t \rangle \rangle \text{Op}_2 \text{Op}_1 [_{IP} \langle t \rangle \dots (pr_1) \dots w_2 \dots]]$

As shown, a sentence may include an element of type e that gets bound by the abstraction operator, provided the two are coindexed. If the sentence contains a PPT with no overt Experiencer argument, the covert internal argument will contribute the predicate that we characterised as being true of individuals whom the speaker identifies with. On the new view, the subject of the *identify with* relation is just a variable bound by the most local abstraction operator. We are almost ready to give a complete semantics for taste statements, but first let me say a little more about what I take sentence meanings to be.

On the orthodox view, a (declarative) sentence has truth conditions, and its meaning can be stated with reference to these. A linguistic expression that expresses a property cannot, of course, be associated with truth conditions. Following a proposal of Heim's (1991) I assume instead that to speak truly is to express a property that one has. I speak truly by uttering in w a sentence S that expresses a property P just in case P is true of me in w .

Our semantics for *This cake is tasty* will be very similar to that which we have already given, except for the abstraction operator in COMP, which binds a variable introduced by the internal argument of *tasty* as we have just described.

74a. This cake is tasty.

74b. LF: $[_{CP} \text{Op}_1 [_{IP} \text{This cake}_i [_{\text{GEN}} [t_i \text{ is tasty } \lambda x. I(y_1, x)]]]]$

74c. $[[74b]]^c = \lambda w \lambda y. \forall_{x, w'} [\text{Acc}(w, w') \ \& \ C_3(\text{this cake}, x, w') \ \& \ I(y, x)] [\text{tasty}(\text{this cake}, x, w')]$

(74c) expresses the property of being a y such that for all accessible worlds w and all x such that this cake and x are relevant in w and y identifies with x , this cake is tasty to x in w . I speak truly by saying *This cake is tasty* just in case this property is true of me. An occurrence of *This cake is tasty* embedded below an attitude predicate has the same semantics. We continue to assume a centred worlds semantics for attitude predicates. Here is an example.²⁷

75a. John thinks that this cake is tasty.

75b. $[_{\text{Op}_2} [_{\text{John thinks}} [_{\text{CP}} \text{Op}_1 [_{\text{This cake}_i} [_{\text{GEN}} [t_i \text{ is tasty } \lambda x. I(y_1, x)]]]]]]]$

75c. $[[75b]] =$

$\lambda w \lambda z. \forall \langle w', u \rangle$: it is compatible with what John thinks in w for John to be u in w' ,

$\lambda w \lambda y. \forall_{x, w''} [\text{Acc}(w, w'') \ \& \ C(\text{this cake}, x, w'') \ \& \ I(y, x)] [\text{tasty}(\text{this cake}, x, w'')](w')(u) = 1 =$

$\lambda w \lambda z. \forall \langle w', u \rangle$: it is compatible with what John thinks in w for John to be u in w' ,

$\forall_{x, w''} [\text{Acc}(w', w'') \ \& \ C(\text{this cake}, x, w'') \ \& \ I(u, x)] [\text{tasty}(\text{this cake}, x, w'')]$

We have given a semantics for PPTs that is quite simple: it requires very little new machinery given that there is independent evidence for the presence of the generic operator, and the

²⁷ The careful reader will notice that there is no individual variable in (75b) for Op_2 to bind. We treat such cases as involving self-locating content in Lewis' sense. I speak truly by asserting that John thinks that this cake is tasty just in case the property of inhabiting a world in which John thinks that this cake is tasty is true of me.

notion of first person oriented genericity is already needed in order to analyze generic *one*. We have also identified some conceptual advantages of the move to treat sentence meanings as properties rather than propositions, which was necessary in order to give a unified account of how the subject of the *identify with* relation is provided for both embedded and unembedded PPTs. We ought to ask, however, whether there is any independent empirical motivation for the presence of an abstraction operator in the left periphery of clauses.

What would constitute evidence for the presence of such an operator? An abstraction operator in matrix COMP would in principle be able to bind variables in its scope in much the same way that Chierchia (1989) proposes for PRO in control complements. As is well known, sloppy identity is diagnostic of bound variable readings. Stojanovic (2011) notes that first person pronouns can receive sloppy readings in matrix sentences.

76a. Alma (to Chris): I am a fool.

76b. Bruce (to Chris): I am a fool.

76c. Chris (to Bruce): That's what Alma said, too.

Reading (i) Alma said that she is a fool. ('sloppy' reading)

Reading (ii) Alma said that Bruce is a fool. ('strict' reading)

[Stojanovic 2011: 7, exs 11-13]

Like Stojanovic, we take this as evidence that sentence meanings have the type of a property, rather than a proposition. A potential objection to this view is that the sloppy reading might arise as a consequence of *that* in the examples above denoting the utterance, *I am a fool*, rather than the meaning expressed by that utterance. That is, reading (i) simply means, 'Alma said, "I am a fool"'. This would predict that we should find the same array of meanings with second and third person pronouns, contrary to fact.

77a. Professor Cheng (speaking of Alma): She is a fool.

77b. Chris (pointing at Daisy): She is a fool.

77c. In reply to Chris: (?) That's what Professor Cheng said, too.

[Stojanovic 2011: 8, exs 17-19]

Moreover, we can illustrate the availability of sloppy readings for first person pronouns by appealing to intuitions of same-saying. In (78), both conclusion (i) and conclusion (ii) are warranted, whereas with *she* the availability of these conclusions depends upon who the intended referent is in each occurrence.

78a. Mary: I am a fool.

78b. John: I am a fool.

Conclusion (i) Mary and John said the same thing.

Conclusion (ii) Mary and John said different things.

79a. Mary: She is a fool.

79b. John: She is a fool.

Conclusion (i) is only warranted if the intended referent of *she* is the same across both occurrences.

Conclusion (ii) is only warranted if the intended referent of *she* is different across both occurrences.

We take it then, following Stojanovic’s insights and a suggestion of Heim’s (1991), that ‘I’ is an obligatorily bound variable, and the abstraction operator in COMP serves as its binder.

80a. I am a fool.

80b. [_{CP} Op_I [_{IP} I₁ am a fool]]

80c. [[(80b)]]^c = $\lambda w \lambda x. x$ is a fool in w

This is a radical departure from the Kaplanian view that first person pronouns are directly referential expressions whose content in a context c is the author in c . By making this move, we treat I as having an almost identical semantics to obligatorily controlled PRO: the difference is that while the latter must be bound by an embedded abstraction operator, the former is (or at least can be) bound at the root level. That I and controlled PRO should have similar semantics is unsurprising: it captures the fact that to report that Pavarotti claimed to be a genius is to report that Pavarotti made the claim, “I am a genius”. On the traditional semantics of first person, by contrast, the intuitive parallel between I and PRO is opaque.

Thus far we have considered only declarative sentences. An anonymous reviewer points out that there is a speaker to hearer shift when PPTs occur in questions: *Is the cake tasty?* seems to ask for the addressee’s opinion of the cake. I agree broadly with Lasersohn and Stephenson that the reason for this is that to ask whether the cake is tasty is to invite the addressee to assert *The cake is tasty* (or its negation), which on my view would amount to self-ascribing the corresponding property. A possible implementation of this would involve associating individual abstraction operators with person features; individual variables are bound by these operators under condition of non-distinctness of person features. In declarative sentences, the occurrence of such an operator in the left periphery of the root clause bears [1st], and therefore binds first person pronouns; in questions it bears [2nd], and therefore binds second person pronouns. In either circumstance the operator binds the subject of the *identify with* relation associated with PPTs. According to this view, the person feature on the abstraction operator encodes information about the individual to whom the sentence meaning is to be applied – the speaker in the case of declarative sentences, or the addressee in questions. In the latter case, the presence of [2nd] indicates that of the set of properties that are possible answers to a yes-no question, the correct one is that which yields truth when applied to the addressee. I leave it to future research to investigate the many questions raised by this suggestion.

Let me close this section with a comparison of the view of conversational dynamics proposed here with that of Stephenson’s and other relativist proposals. The set of propositions that form the common ground is as usual modelled as a set of worlds, rather than as a set of world-individual pairs as in Stephenson’s system: propositions for us are sets of worlds, not sets of world-individual pairs. Furthermore, a common ground is updated in more or less the usual way, by eliminating worlds. The effect of update of a common ground with the content of S uttered by an individual x is to eliminate worlds in which the proposition obtained by applying the property expressed by S to x is false.

According to relativist views such as Stephenson’s, there is a class of expressions that give rise to sentences whose content can only be added to the common ground if the sentence is true relative to the addressee: *The cake is tasty* cannot be added to the common ground if there is some conversational participant relative to whom the sentence is false. Sentences

with first person pronouns, on the other hand, lack this property: if I say that I am a fool, the addition of what I have said to the common ground is not contingent upon my interlocutors being prepared to accept that *they* are fools. By contrast, our account unifies the two cases: both taste statements and sentences with first person pronouns are evaluated with respect to the speaker. The content of *The cake is tasty* may still fail to be added to the common ground because an addressee dislikes the cake, but this is not because the sentence is false relative to her, but rather because the sentence carries a false entailment – namely, that she finds the cake tasty. According to this view, whether I speak truly by uttering a taste statement or a sentence with *I* is determined in the same way: by application of the relevant property to me. This unification seems to us to be an appealing aspect of the analysis.

To conclude this section, we assigned a semantics to the predicate associated with the covert internal argument of PPTs without appeal to the judge parameter. This took us into new territory: the semantics of first person indexicals and the general issue of sentence meanings. In doing so, we encountered an empirical argument in favor of treating sentences as expressing properties, via abstraction by an operator in the left periphery of the clause. Now let's consider some consequences of our proposal for the semantics of taste predicates.

4 Consequences

4.1 Explaining subjective disagreement

Here is a reminder of the meaning assigned to *tasty* (81a) and to *This cake is tasty* (81b).

81a. $[[\text{tasty}]^c]^w =$

$\lambda x \lambda y: x$ has direct perceptual experience of the relevant kind of y in w . y is tasty to x in w

81b. $[[[\text{Op}_1 [\text{This cake}_i [\text{GEN} [t_i \text{ is tasty } \lambda x. I(y_i, x)]]]]]]^c =$

$\lambda w \lambda y. \forall_{x, w'} [\text{Acc}(w, w') \ \& \ C_3(\text{this cake}, x, w') \ \& \ I(y, x)] [\text{tasty}(\text{this cake}, x, w')]$

In this section, we spell out in more detail how our proposal accounts for subjective disagreements. Let's return to John and Mary disagreeing over whether the cake is tasty.

82a. John: This cake is tasty.

82b. Mary: No, this cake is not tasty.

Properties (I) and (II), described in the introduction, can both be explained on our account. We can explain Property (I), the intuition that the parties to a subjective disagreement genuinely disagree with each other, as follows. Suppose that each speaker selects an *identify with* relation that has the consequence that the set of individuals ranged over by GEN is the same. Suppose, as seems plausible, that John identifies with himself and with Mary, and that Mary identifies with herself and with John. In such a case, the truth value returned by applying the meaning of *This cake is tasty* to the speaker is not expected to vary depending on whether it is uttered by John or by Mary. What we have then, is a genuine disagreement.

To understand Property (II), whereby intuitively each of John and Mary speaks truly so long as they report their tastes sincerely, note that this property holds not only of *This cake is tasty*, but also of *This cake is tasty to me*. In the latter case, it is not difficult to see why this is: when I

say whether something tastes good to me, I do so from a position of epistemic privilege, since I am reporting a sensory experience which under typical circumstances I and I alone have access to. As long as I report this experience sincerely, then, I cannot fail to speak truly.

Note, however, that whereas taste statements with a (singular) first person internal argument always have Property (II), their counterparts with a covert internal argument do not always have it. Compare the following:

83a. Soapy dishwater is tasty.

83b. Soapy dishwater is tasty to me.

Intuitively, (83a) is just plain false: asked whether it is true or not I am unlikely to respond that it depends on who is doing the tasting, as I might do when asked whether spinach, say, is tasty. On the other hand, there might be people of whom (83b) is true. If John utters this sentence and I can be persuaded that he is sincere, then I cannot fail to accept what he has said as true. Nonetheless, if John asserts that soapy dishwater is tasty and I reply that it is not, the intuition that neither party to the disagreement is at fault evaporates: I have said something true, and John has said something false. What has John done wrong? Why did he speak truly in uttering (83b) but falsely in uttering (83a)?^{28 29}

It seems that John's mistake was to generalize on the basis of his personal experience - the experience that entitles him to utter (82b) - to the experience of others. He has failed to recognise that his tastes are exceptional, which is to say that others are not like him in tastes. The assumption that other people (one's interlocutors, say) resemble the speaker in tastes has been referred to as the 'Presupposition of Commonality'; Lopez de Sa (2008) suggests that the presupposition is taken for granted in ordinary discourse involving taste predicates.³⁰

I leave it to future work to investigate the place of the Presupposition of Commonality within the proposed account. Indeed, 'presupposition' may be a misnomer here: Commonality seems to fail the 'Hey, Wait a Minute!' test, for example. My own inclination is to say that Commonality is a pragmatic imperative to adopt the working assumption that the conversational participants are alike in tastes. Doing so enables us to engage in sensible exchanges about taste: I am justified in concluding from *The cake is tasty to me* that *The cake is tasty* only in so far as I assume that I resemble those with whom I identify in tastes. This assumption may involve a degree of pretence, however: when we persist in arguing about whether the cake is tasty we cannot help but accumulate evidence that we are *not* alike in

²⁸ See Egan (2010) for other examples of false taste statements.

²⁹ An anonymous reviewer raises the possibility that, contrary to what we claimed in Section 3.1 about the usage conditions associated with predicates of personal taste, it may be possible to describe something as disgusting without having tasted it. We think that discussions of soapy dishwater may constitute such a case: *Soapy dishwater is disgusting* seems to be more felicitous than *Carrots are disgusting* in a context in which I have no taste experience upon which to base my judgments. We suspect that this may be related to the facts in (82): just as *Soapy dishwater is tasty* is manifestly false, *Soapy dishwater is disgusting* is manifestly true. It is as though extreme examples such as this produce an instinctive response of disgust strong enough to overcome the patterns we find with more run of the mill cases such as *The cake is tasty*. Just why this should be remains an open question.

³⁰ See also Egan (2010) for discussion of the role of shared tastes in taste disputes.

tastes. By maintaining the convenient fiction that Commonality holds, Mary is permitted to respond to John that the cake is not tasty – that is, to produce a sentence that entails that the cake is not tasty to John – even though she has ample reason to believe that John does find the cake tasty given that he has just said *This cake is tasty*. Just how long the interlocutors are prepared to engage in the pretence that Commonality holds in spite of the evidence to the contrary will determine how long a disagreement will continue before somebody partially backs down and says, ‘Well, it is/is not tasty to *me*, at least’.

What we have learned is that Property (II) holds insofar as Commonality holds, at least as a reasonable working assumption. When I say that this cake is tasty I say that it tastes good to me and to others with whom I identify, and I take it for granted that the members of this group are alike in taste. If I am sincere in my report, and if the assumption of Commonality is justified (either because there is no evidence to the contrary, or because the conversational participants are willing to pretend that Commonality holds), then I am justified in generalizing to others’ experience on the basis of my own, and I cannot fail to say something true. When John utters (83a), he says that soapy dishwater tastes good to him and to others with whom he identifies, but he is *manifestly* wrong to assume that his tastes are like others’: since almost nobody would find soapy dishwater tasty, his tastes must be highly unusual, perhaps even defective. Consequently, what he says is false.

4.2 Explaining how ‘exocentric’ interpretations arise

A pleasing property of our analysis is that by appealing to genericity it ensures that taste statements always include a contextual restriction C_3 in their semantic representation. The contextual restriction is responsible for determining how and when ‘exocentric’ interpretations arise. Suppose for example that the contextual restriction excludes the interlocutors in a conversation, say because they are discussing cat food which they are unlikely to have tried. Together, they have a cat, Snowball, and they watch her to see how she responds to the new cat food.

- 84a. John: The cat food must be tasty. Snowball’s eaten half of it.
84b. Mary: It can’t be, she normally finishes the whole bowl.

In this context, it seems that John empathizes with Snowball and other cats, while Mary does the same. They presumably also empathize with themselves and with each other, but this is of no consequence: since John does not eat cat food, his tastes are irrelevant to whether the cat food is tasty. That is, there is no accessible world w such that John is a relevant individual in w . John, and likewise Mary, will therefore be excluded by the contextual restriction from the set of individuals in the restrictor of GEN. *The cat food is tasty* says that for every cat x and every accessible world w such that x inhabits w and x is relevant in w , the cat food is tasty to x in w . On this view, there is nothing special about ‘exocentric’ readings: they merely arise in contexts in which the speaker is not among the ‘target audience’, and is therefore excluded from the contextual restriction. Because the *identify with* relation relates the agent to more individuals than just herself, it is possible under these circumstances for a taste statement to report the tastes of individuals other than the speaker.

Since on our view subjective disagreement is accounted for by having the speaker generalize from direct experience, no subjective disagreement is predicted when the speakers’ direct

experience is irrelevant, as in this instance. We would contend that this is indeed the case: there is a fact of the matter about whether the cat food is tasty to Snowball and to other cats, and one can imagine circumstances that prove one of John and Mary right and the other wrong. Such scenarios would most likely involve further observation of Snowball's behavior, but if she suddenly acquired the ability to talk and told John and Mary just what she thought of the cat food, the argument would be settled then and there.

This much also follows from Stephenson's argument, since she denies that examples like (84) involve any judge-dependency, instead treating the internal argument of *tasty* as the context-dependent element *pro*. By her lights, (84) is no more a subjective disagreement than is a dialogue of form, *The cat food must be tasty to Snowball./No, it can't be!* We think our explanation of why (84) is not a subjective disagreement is to be preferred to Stephenson's, however, for the following reason.

- 85a. John: Snowball hasn't touched her cat food. It must not be tasty.
85b. Mary: *Observes a stray cat come into the kitchen and tuck in to the food.*
No, the cat food must be tasty, look at the way this cat is enjoying it. Snowball's just picky.

The felicity of Mary's reply indicates that the internal argument of *tasty* in John's utterance cannot be *pro*: since neither John nor Mary have set eyes on the stray cat before, it cannot be included in the referent determined by *pro*, which presumably must be familiar in the discourse. Moreover Mary's comment that Snowball is 'just picky' shows that she does not deny that Snowball dislikes the cat food. An account that appeals to genericity fares better.

Given that *pro* is eliminated from our system, we should expect that the problem of overgeneration due to unattested interpretations involving *pro* should also be absent. In the next section, we check whether this is so.

4.3 Explaining when 'exocentric' interpretations arise

Since our analysis depends on the contextual restriction to produce 'exocentric' readings, we predict that these interpretations will arise in just those situations in which the agent's tastes are irrelevant to the truth of the taste statement.

Happily, this is the very generalization that the data led us towards in section 3.1. We said that a sentence of form *x is tasty* will carry entailments about the agent's tastes unless (i) the agent has not tasted *x* or (ii) there is something about the agent that makes her an unsuitable candidate for being an individual to whom *x* is tasty. We have just discussed an example in the second category: talking about food that is not for human consumption. The second category can be thought of in terms of the presupposition of direct experience stated in the lexical entry of *tasty*. We assume that as with ordinary universally quantified sentences, presuppositions in the scope of the generic operator project universally. One way to avoid presupposition failure in a universally quantified sentence is by selecting a domain that is such that every element in it has the required property (von Stechow 2008). In this case, this means excluding world-individual pairs *w* and *x* such that *x* has not tasted the food item in *w*. If the speaker has not tasted it in the actual world, then the pair consisting of the speaker and the actual world will be excluded. However, pairs consisting of the speaker and

counterfactual worlds in which she has tasted the food item will be included in the contextual restriction, so that *The cake must be tasty* will entail that the cake is tasty to the speaker in those worlds in which she eats it. This modal entailment explains the infelicity of the continuation, ...*but I wouldn't like it*.

4.4 Explaining usage conditions

We also discovered a condition on when a PPT may be used. We noticed that *Shortbread is tasty* is infelicitous unless the speaker has tasted shortbread, even if she has ample justification for supposing it to be tasty without tasting it. Given the lexical entry for *tasty*, a presupposition failure will result if the generic operator ranges over individuals that lack direct perceptual experience of the relevant kind – that is, individuals who have not tasted the subject of *tasty*. Yet we cannot simply assume that this set always includes the speaker, since we would then expect that the presupposition will survive in the scope of any operator that is a hole for presupposition projection: the felicity of *Shortbread must be tasty* in a context in which the speaker has not tasted shortbread would be unexpected. This can be accounted for by the assumption that the speaker is automatically included in the set of individuals over which GEN ranges unless there is some reason to exclude her, such as the signal of indirect evidence associated with *must*. In this section, we explain this automatic inclusion of the speaker in terms of properties of generic sentences, and particularly the nature of the contextual restriction.

We propose that the contextual restriction does not rule out particular world-individual pairs unless the (linguistic or extralinguistic) context provides an explicit basis on which to do so. In fact, we have already seen some evidence in favor of this view – examples (40-42), repeated below.

86a. (?) Linguists write their papers using LaTeX, but I use Microsoft Word.

86b. (?) Linguists write their papers using LaTeX, but Mary uses Microsoft Word.

87a. Most linguists write their papers using LaTeX, but I use Microsoft Word.

87b. Most linguists write their papers using LaTeX, but Mary uses Microsoft Word.

88a. Linguists usually/typically/generally write their papers using LaTeX, but I use Microsoft Word.

88b. Linguists usually/typically/generally write their papers using LaTeX, but Mary uses Microsoft Word.

We saw that in a context which entails that Mary and I are linguists, the sentences in (86) are degraded with respect to those in (87) and (88). This is somewhat surprising given that generic statements are known to tolerate exceptions: if the contextual restriction were at liberty to freely exclude irrelevant individuals, (86) should be perfectly acceptable. (86) is improved by explaining what property of me makes me an atypical linguist:

89. Linguists write their papers using LaTeX, but I use Microsoft Word because I don't mind numbering examples by hand.

The same point can be made by probing our intuitions of what constitutes a valid argument when one of the premises is a generic sentence. (90) and (91) are both valid; the rule of Universal Instantiation is apparently applicable here despite the fact that generic statements tolerate exceptions. In (92) and (93) the premises are amended to make explicit why John should be excluded from the contextual restriction, and the arguments are no longer valid.

- 90a. Linguists write their papers using LaTeX.
- 90b. John is a linguist.
- 90c. *Conclusion*: John writes his papers using LaTeX.

- 91a. Linguists write their papers using LaTeX.
- 91b. John doesn't write his papers using LaTeX.
- 91c. *Conclusion*: John is not a linguist.

- 92a. Linguists write their papers using LaTeX.
- 92b. John is a linguist, but he doesn't mind numbering examples by hand.
- 92c. *Conclusion*: # John writes his papers using LaTeX.

- 93a. Linguists write their papers using LaTeX.
- 93b. John doesn't write his papers using LaTeX because he doesn't mind numbering examples by hand.
- 93c. *Conclusion*: # John is not a linguist.

In the unmarked case, then, suitable individuals are included in the restrictor of a generic quantification. Since the *identify with* relation is reflexive, in a taste statement the speaker is included in the restrictor of GEN unless it is made explicit that she should be excluded from the contextual restriction – say by use of epistemic *must*. In uttering *Shortbread is tasty*, the speaker makes no such information explicit, and the sentence will therefore entail that shortbread is tasty to the speaker. Since *Shortbread is tasty to me* presupposes that I have tasted shortbread, I can felicitously utter a sentence that entails that shortbread is tasty to me only if I have tasted shortbread.

On the other hand, if I add epistemic *must*, I overtly signal that I have not tasted shortbread, thereby ensuring that the pair consisting of me and the actual world is excluded from the contextual restriction. *Shortbread must be tasty* therefore does not entail that I find shortbread tasty in the actual world, and hence there is no presupposition that I have tasted shortbread. This is why the presupposition that I have tasted shortbread has the appearance of being filtered by epistemic *must*, even though *must* is a hole for presupposition projection.

This account provides an answer to the question of why *Shortbread must be tasty* is felicitous, but *Shortbread must be tasty to me* is not. Whereas the speaker can be excluded in the manner just described when there is no overt internal argument, with the PP *to me* this option is unavailable, leading to a clash between the signal of indirect evidence carried by *must* and the presupposition that the speaker has tasted the cake introduced by virtue of the semantics of *tasty*.

A prediction of our account is that a continuation in which the speaker acknowledges that she has not tried something will also serve to exclude the speaker and actual world pair from the contextual restriction, thereby resulting in a coherent discourse. An example raised by an anonymous reviewer suggests that this prediction is borne out.³¹

94. (*Uttered by a cook*) The cake is delicious – I haven’t tried it, but since I made it myself with the best ingredients, I can say with certainty that it is delicious.

The view that the contextual restriction excludes individuals only if there is explicit reason to do so helps to explain two further observations. Firstly, consider the first person orientation of PPTs, exemplified by the contrast between (95a) and (95b).

95a. # The cake that Mary and I ate was tasty, but I didn’t like it.

95b. The cake that Mary and I ate was tasty, but Mary didn’t like it.

On our view, the contextual restriction excludes neither Mary nor the speaker in these examples, since no explicit basis on which to exclude these individuals is provided. Since the *identify with* relation is reflexive, the speaker is included in the restrictor of GEN, explaining the infelicity of (95a). In (95b), there is no requirement that the speaker identify with Mary, and hence the first conjunct does not entail that Mary found the cake tasty. The first person orientation of PPTs can therefore be understood as arising from a conspiracy between the contextual restriction and the *identify with* relation.

Finally, we can now explain the contrast between (31a) and (32), repeated below.

96a. # The cake must be tasty, but I wouldn’t like it.

96b. The cake must be tasty, but I wouldn’t like it because I don’t like chocolate.

The information that the speaker doesn’t like chocolate constitutes evidence that she is not among the target audience of the cake, and hence should be excluded from the contextual restriction. The continuation in (96b) provides this information, yielding a felicitous sentence. Absent this information, the speaker is automatically included in the contextual restriction, so that (96a) is degraded. Note that the effect of the explicit denial that the speaker is in the target audience mirrors that of the explicit denial that the speaker has tried the cake in (94). That these two phenomena can both be blamed on the role of explicit information in excluding elements from the contextual restriction is a pleasing aspect of our proposal.

Having offered explanations for the puzzles encountered in the first part of this paper, we next identify two additional predictions of our account.

³¹ The reviewer suggests that in this example, the speaker’s providing the evidence upon which she bases her statement ameliorates it. We agree that this helps, however we believe that the variant of (94) in which *since I made it myself with the best ingredients* is omitted is also acceptable, suggesting that the crucial point is the explicit acknowledgment that the speaker has not tried the cake.

4.5 A prediction about arbitrary PRO

One of the initial motivations for appealing to genericity came from cases involving arbitrary PRO such as (37), which we repeat here.

97. It is fun [PRO_{arb} to ride the rollercoaster].

We said that the internal argument of *fun* covaries with the interpretation of PRO_{arb}, so that (97) means something like, ‘whoever rides the rollercoaster has fun’. Our analysis accounts for this, but it also predicts something a little stronger, which we shall argue is born out.

98. # It is fun to ride the rollercoaster, but I have never done it.

(98) shows that (97) is only felicitous if the speaker has ridden the rollercoaster. This cannot be explained by treating the entailment as a property of PRO_{arb}, since (97) is felicitous.

99. It is selfish [PRO_{arb} to ride the rollercoaster], but I have never done it.

This requirement seems to be peculiar to occurrences of PRO_{arb} bound by the internal argument of a PPT. Since we have seen that there is a requirement that unless signalled otherwise, the speaker must have direct experience of whatever is described by the subject of the PPT, the facts reported in this section are expected.

4.6 A prediction about partial retraction

A difference between our proposal and Lasnik's is that we predict that *The cake is tasty* asymmetrically entails *The cake is tasty to me*, whereas for him, the two sentences are true under exactly the same circumstances (assuming an autocentric context). In 4.1 we saw that at some point during a taste disagreement somebody may partially retract their original assertion, as John does in the following exchange.

100a. John: This cake is tasty.

100b. Mary: No, it's not tasty!

100c. John: Well, it's tasty to *me*, at least.

The possibility of partial retraction follows from the asymmetric entailment relation between *The cake is tasty* and *The cake is tasty to me* which is in turn predicted by our account.

An anonymous reviewer points out that Stephenson's proposal can also accommodate dialogues such as (100), given her view of the pragmatics of assertion: according to her, it is acceptable to update the common ground with the content of *This cake is tasty* just in case all conversational participants find the cake tasty. The possibility of partial retraction in (100c) may therefore be a consequence of the impossibility of adding John's original assertion to the common ground given Mary's response. The condition that renders this impossible is irrelevant to the proposal to add the content of (100c) to the common ground, and therefore this is an acceptable conversational move whose effect on the common ground is weaker, giving rise to the impression of partial retraction. As discussed in section 3.5, this view depends on an assumption about the pragmatics of conversation some version of which is

required for a relativist semantics, but which can be dispensed with on our account: namely, that there is a class of expressions that give rise to statements that can only be added to the common ground if all conversational participants agree that the content is true relative to them.

5. Conclusion

We have given a new semantics for predicates of personal taste that makes no appeal to a judge parameter as a component of the evaluation index. The account also differs from its predecessors in that it treats the covert internal argument of a PPT as a variable restricted to individuals identified with by the agent, which gets bound by the generic operator. In other words, the interpretation of taste predicates involves a peculiar type of genericity - the first person oriented type identified by Moltmann for generic *one*.³² The account is to be preferred to Stephenson's not only because it unifies the cases that she had treated by appeal to either a judge-denoting or garden variety pronominal, but also because this unification across what had appeared to be different classes of case is in fact a better fit with the data.

Let me close by identifying virtues of Stephenson's framework that are retained in our system. These concern the proper analysis of attitudes *de se*. Chierchia handled the obligatorily *de se* interpretation of controlled PRO by applying a Lewis-style analysis to control complements, whereby they are treated as expressing properties. Clausal complements of attitude predicates that do not contain a *de se* pronoun were assigned propositional type, as were matrix clauses. This amounts to a bifurcation of semantic content, whereby CPs have type $\langle s,t \rangle$ or $\langle s,\langle e,t \rangle \rangle$, depending on whether or not an abstraction operator is present. One consequence is that it is necessary to posit two lexical entries for attitude predicates such as *expect* that can take both finite clausal complements and control complements: one that composes with a constituent of type $\langle s,t \rangle$, and another that composes with a constituent of type $\langle s,\langle e,t \rangle \rangle$. An advantage of Stephenson's framework is that by ensuring that propositions are of type $\langle s,\langle e,t \rangle \rangle$, this bifurcation of content is rendered unnecessary: that attitude predicates are quantifiers over centred worlds fits quite naturally with the move to treat sentences as expressing sets of centred worlds.

Our proposal retains this virtue: all sentences are treated as expressing meanings of the same semantic type, $\langle s,\langle e,t \rangle \rangle$. For us, though, this is the type of properties, not propositions. That is, within Stephenson's framework the assignment of type $\langle s,\langle e,t \rangle \rangle$ to the CP is a consequence of the addition of the judge parameter, whose effects are felt throughout the inventory of semantic types. By inserting an abstraction operator in the left periphery of the

³² Perhaps unsurprisingly, it seems that Moltmann also noticed the applicability of her analysis of generic *one* to PPTs, and independently proposed an account building on this insight (Moltmann 2010b). Our analysis differs from hers in several respects, most significantly in how we each account for subjective disagreement. According to Moltmann, *This cake is tasty* has an objective, propositional content, but the sentence is accepted or rejected on the basis of a first personal application of this content by the addressee. As we understand it, this treatment of acceptance as potentially varying from individual to individual (which is not applied to other sentences such as those involving the first person pronoun) is a property that Moltmann's view shares with that of relativists. Our view is that one accepts *This cake is tasty* under just the same circumstances in which one accepts *I am a linguist*: namely, just in case the property expressed by those sentences is true of the *speaker* rather than of the *addressee*.

clause, further complication of the semantics is avoided. The trade off, of course, is that our analysis requires a complication of the syntax in the form of added abstraction operators.

Like Stephenson, we also predict that PPTs are interpreted *de se*. On our account, this follows from the fact that the subject of the *identify with* relation is bound by an abstractor, yielding a property that is self-ascribed by the agent. (101) suggests that this prediction is borne out. John's assertion in the scenario described is somewhat infelicitous; John would instead be expected to say, *That cake must be tasty*. Someone who heard John's assertion and later learned about the circumstances in which it was made would be entitled to feel misled.³³ On our view, this is because *That cake is tasty* typically entails that the cake is tasty to the speaker, due in part to the reflexivity of the *identify with* relation. If the subject of this relation could be construed *de re*, this would be sufficient for John's utterance to be felicitous. That it is not suggests that the subject of the relation is construed *de se*, as predicted.

101. *Context: John is watching someone eating a cake. He notices that he seems to be enjoying it, and concludes that cake is tasty to the guy he's looking at. He is unaware that what he is looking at is a reflection of himself in the mirror eating a cake.*

John: ?? That cake is tasty.

Within Stephenson's framework, the obligatory *de se* construal would be accounted for by letting the internal argument of the PPT denote the judge, and having the judge parameter be bound by the attitude predicate. She also proposes that controlled PRO is lexically specified as denoting the judge, so that the *de se* nature of PPTs and control receives a unified analysis. For us, *de se* expressions – not only PRO and the subject of the *identify with* relation associated with PPTs, but also the first person pronoun – are treated as bound variables. We thereby gain a unified treatment of all three types of *de se* element, an interesting result that we think warrants further investigation in the future.

Acknowledgments

I have profited from discussions arising from presentations of this work at Harvard, at the Cerisy conference on Meaning, Context and Implicit Content, and at Sinn und Bedeutung 16, from exchanges with Rajesh Bhatt, Amy Rose Deal, Irene Heim, Laurence Horn, Dirk Kindermann, Maria Polinsky, Craige Roberts, Jacopo Romoli, Susan Rothstein, Yael Sharvit, Tamina Stephenson and Igor Yanovich, and especially from many hours of conversations about this and related topics with Gennaro Chierchia. Special thanks to two anonymous *Journal of Semantics* reviewers. Responsibility for the remaining problems rests with me.

³³ A scenario where one is eating a cake and (unwittingly) looking in a mirror and does not realize that what one sees is the reflection of oneself eating the same cake is of course no more plausible than, say, the possibility that one might be unaware that one's pants are on fire.

References

- Carlson, G.N. (1977) *Reference to Kinds in English*. Doctoral dissertation, UMass, Amherst.
- Carlson, G.N. (1982) Generic Terms and Generic Sentences. *Journal of Philosophical Logic* 11(2):145-181.
- Chierchia, G. (1989). Anaphora and Attitudes De Se. In R. Bartsch, J. van Benthem, and P. van Emde Boas (eds.), *Language in Context*, Dordrecht: Foris, 1–31.
- Chierchia, G. (1995) Individual-level predicates as inherent generics. In G.N. Carlson & F.J. Pelletier (eds.), *The Generic Book*: 125-175. Chicago: University of Chicago Press.
- Egan, A. (2010). Disputing about Taste. Unpublished Ms., Rutgers University.
- von Fintel, K. (2008) What is Presupposition Accommodation, Again? *Philosophical Perspectives* 22(1): 137-170.
- von Fintel, K. & A. Gillies (2010). *Must...stay...strong!* *Natural Language Semantics* 18: 351-383.
- Heim, I. (1982). *The semantics of definite and indefinite noun phrases*. Doctoral dissertation, UMass Amherst.
- Heim, I. (1991) Lecture notes on ‘The first person’.
- Kamp, H. (1981) A Theory of Truth and Semantic Representation. Reprinted in P. Portner & B.H. Partee (eds.), *Formal Semantics: the Essential Readings* (2008). Oxford: Blackwell. 189-222.
- Keshet, E. (2005) A Matter of Taste. Unpublished manuscript, MIT.
- Koelbel, M. (2002). *Truth without Objectivity*. London: Routledge.
- Koelbel, M. (2003). Faultless disagreement. *Aristotelian Society*, 104, 53-73.
- Kratzer, A. (1995). Stage-level and individual-level predicates. In G.N. Carlson & F.J. Pelletier (eds.), *The Generic Book*: 125-175. Chicago: The University of Chicago Press.
- Lasnik, P. (2005). Context dependence, disagreement, and predicates of personal taste. *Linguistics and Philosophy* 28, 643–686.
- Lewis, D. (1975) Adverbs of Quantification. Reprinted in P. Portner & B.H. Partee (eds.), *Formal Semantics: the Essential Readings* (2008). Oxford: Blackwell. 178-188.
- Lewis, D. (1979). Attitudes *De Dicto* and *De Se*. *Philosophical Review* 88: 513–543.
- Lopez de Sa, D. (2008). Presuppositions of Commonality: An Indexical Relativist Account

of Disagreement. In M. Garcia-Carpintero and Max Koelbel (eds.), *Relative Truth*. Oxford: Oxford University Press. 297-310.

Milsark, G. (1974). *Existential sentences in English*. Doctoral dissertation, MIT.

Moltmann, F. (2006). Generic *one*, arbitrary PRO, and the first person. *Natural Language Semantics* 14: 257-281.

Moltmann, F. (2010a). Generalizing Detached Self-Reference and the Semantics of Generic *One*. *Mind and Language* 25(4): 440-473.

Moltmann, F. (2010b). Relative Truth and the First Person. *Philosophical Studies* 150: 187-220.

Saebo, K. J. (2009). Judgment ascriptions. *Linguistics and Philosophy* 32: 327-352.

Stephenson, T. (2007a). *Towards a Theory of Subjective Meaning*. Doctoral dissertation, MIT.

Stephenson, T. (2007b). Judge dependence, epistemic modals, and predicates of personal taste. *Linguistics and Philosophy* 30, 487-525.

Stephenson, T. (2010). Control in centred worlds. *Journal of Semantics* Vol 27, Issue 4, 409-436.

Stojanovic, I. (2011). De Se Assertion. Unpublished manuscript, Institut Jean Nicod.

Stump, G. (1985). *The Semantic Variability of Absolute Constructions*. Dordrecht: Reidel.