

Relative Truth, Speaker Commitment, and Control of Implicit Arguments*

Peter Lasersohn

University of Illinois

1. Introduction: Relativism and Faultless Disagreement

A fundamental problem in the theory of truth conditions is how to deal with sentences which intuitively are not about matters of fact, but seem to involve some ineliminable element of subjective judgment or opinion. Perhaps the clearest such examples are sentences dealing with matters of “personal taste,” such as (1)a. and b.:

- (1) a. Licorice is tasty.
b. Roller coasters are fun.

Recently, I have been developing and defending a semantic theory which deals with such examples by relativizing truth values to individuals, so that sentences like these may be true relative to you, for example, but false relative to me (Lasersohn 2005).

Perhaps the most obvious way to develop such a theory would be analyze these sentences as containing some sort of hidden indexical, so that “true relative to you” means “true when the value of the hidden indexical is set to you”, and “false relative to me” means “false when the value of the hidden indexical is set to me”. If we say this, then sentences like these have different truth values relative to different individuals because they express different contents relative to different individuals.

For several reasons, I think this approach is wrong, and that sentences like these do not express different contents relative to different individuals. Yet, I would suggest, they do vary in truth value from person to person. This means that the relativization is not like what we get with ordinary indexicals, and calls for a different kind of semantic technique.

* A version of this paper was presented at the Rutgers University Semantics Workshop, 9/30/2006, where John MacFarlane provided commentary. Thanks to John, to the audience at that talk, and to the audience at NELS for many useful comments and suggestions.

Peter Lasersohn

What I have suggested is that we recognize a role for context not just in fixing the contents of sentences on particular occasions of use, but also in determining the truth values of those contents once they are fixed. In this way, the same content can be evaluated as true relative to one context but false relative to another. If we regard each context as providing a relevant individual, then the same content can be true relative to one individual but false relative to another.

The idea that context plays a role not just in determining content, but also in assigning truth values *to* contents, can be found already in Kaplan (1989) — though of course he did not apply this idea in the way I am suggesting. And although I wasn't aware of it when I first started advocating this sort of approach to "personal taste" sentences, a whole series of very similar proposals to mine were already being made not just for personal taste examples, but in the analysis of aesthetic, moral, and probability statements more generally (Kölbel 2002), for future contingents (MacFarlane 2003), for vagueness (Richard 2004), and for epistemic modals (Egan, et al. 2005).

There are some differences among these proposals, both in motivation and in implementation. But I think it is fair to say that they all take essentially the same kind of strategy, which is a form of relativism; and in all of them, a central motivation for that strategy is to provide an account of what has come to be called "faultless disagreement" (Kölbel 2003) — that is, situations where one person asserts or believes some sentence content, and another person believes or asserts its negation, but where neither person has made a mistake.

For example, if John says "Licorice is tasty" and Mary says "No, licorice is not tasty," we intuitively regard John and Mary as disagreeing, even contradicting each other; yet we may also feel that each of them is entitled to his or her own views on this matter, and that therefore in some sense neither one is "making a mistake." Accounting for such cases in a conventional semantic theory is a challenge, because in such theories, if Mary is negating the very same content that John asserts, then one of them must be saying something false; and if she is not negating the same content as he asserts, they would not seem to be contradicting one another. On the face of it, then, conventional semantic theories allow us to account either for the intuition of disagreement, or for the intuition of faultlessness, but not for both simultaneously.

Because the analysis of faultless disagreement has played such a central role in arguments for relativist semantic analyses, I think the impression may be strong that this is practically the only motivation for this sort of semantic theory — and that therefore, if some other, less radical explanation for our intuitions of faultless disagreement could be found, the motivation for relativism would be removed, and we could retain a more conventional style of semantic theory.

What I hope to show in this paper is that there are arguments for a relativist semantic theory which do not depend in any direct way on the analysis of disagreement,

Relative Truth and Speaker Commitment

and that therefore a defense of a more conventional, non-relativist semantic theory must do more than simply provide an explanation for our intuitions of faultless disagreement.

The arguments I will give depend on the interpretation of predicates of personal taste in certain attitude contexts and presuppositional constructions, and on the “control” of putative implicit arguments. An examination of these topics will serve primarily to decide between a relativist analysis and what I take to be the leading alternative, namely an analysis which instead makes use of implicit indexical arguments, so that sentences like those in (1) express different contents relative to different people. It will not close off all other lines of analysis, so I will not be claiming that these arguments firmly establish the need for a relativist approach. But I think they will at least provide some reason for preferring such an approach to one of the more obvious and popular alternatives.

2. An Elementary Relativist Semantics

Before proceeding to the main arguments of the paper, it will be useful to review the major features of the kind of relativist semantic theory I have in mind, and present some formalism for making it precise. In doing this, I will naturally follow my own earlier presentation in most of the details, though most of what I have to say could be recast to conform to related proposals by others without doing too much violence to my main points.

To make things concrete, I give a little “toy” language to illustrate the technique. This is essentially the same language I gave in Lasersohn (2005), and is based directly on the system of Kaplan (1989). Some minor adjustments and simplifications have been made to my earlier system, and some aspects of the formalism are left implicit; readers desiring more detail are referred to the earlier article.

The basic expressions of the language include the constants *I*, *John*, *Mary*, *the-licorice*, and *the-roller-coaster*; and the predicates *fun*, *tasty*, *contain-sugar*, and *be-six-feet-tall*.

The formation rules of the language are the following:

- (2) a. If α is a predicate and β is a constant, then $\alpha(\beta)$ is a sentence.
- b. If α is a predicate and β is a constant, then $[\alpha \text{ for } \beta]$ is a predicate.
- c. If ϕ is a sentence, then $\neg\phi$ is a sentence.

We assign denotations to expressions relative to contexts, worlds, and individuals. We write $\llbracket \alpha \rrbracket^{c,w,i}$ for the denotation of α relative to context c , world w and individual i .

Each context c must specify at least an “agent,” or speaker/author, notated $A(c)$; a world $W(c)$; and a “judge,” $J(c)$, which will be the individual on whom the truth value of sentences containing predicates of personal taste depends.

The constant I is the first person singular pronoun, and always denotes the agent: for all c, w, i : $\llbracket I \rrbracket^{c, w, i} = A(c)$.

All other basic expressions have the same denotation relative to every context: $\llbracket \alpha \rrbracket^{c, w, i} = \llbracket \alpha \rrbracket^{c', w, i}$ for all c, c' .

Predicates will denote sets of individuals. We distinguish between “subjective” predicates like *fun* and *tasty*, whose denotation is a matter of personal taste and therefore varies from individual to individual, and “objective” predicates like *contain-sugar* and *be-six-feet-tall*, whose denotation does not vary in this way. That is, if α is an objective predicate, we require for all c, w, i, j : $\llbracket \alpha \rrbracket^{c, w, i} = \llbracket \alpha \rrbracket^{c, w, j}$.

Constants will denote individuals. All our constants are objective, and should not vary in denotation according to i ; we impose the same requirement as for objective predicates, that for all c, w, i, j : $\llbracket \alpha \rrbracket^{c, w, i} = \llbracket \alpha \rrbracket^{c, w, j}$. Proper names should also denote rigidly across worlds, but I will not bother to formulate that requirement here.

Sentences will denote truth values, with predicate-argument relations and negation working in the expected way:

- (3) a. $\llbracket \alpha(\beta) \rrbracket^{c, w, i} = 1$ if $\llbracket \beta \rrbracket^{c, w, i} \in \llbracket \alpha \rrbracket^{c, w, i}$
 $\llbracket \alpha(\beta) \rrbracket^{c, w, i} = 0$ if $\llbracket \beta \rrbracket^{c, w, i} \notin \llbracket \alpha \rrbracket^{c, w, i}$
- b. $\llbracket \neg \varphi \rrbracket^{c, w, i} = 1$ if $\llbracket \varphi \rrbracket^{c, w, i} = 0$
 $\llbracket \neg \varphi \rrbracket^{c, w, i} = 0$ if $\llbracket \varphi \rrbracket^{c, w, i} = 1$.

Modification of a predicate with a *for*-phrase makes it objective; *tasty for Mary* will denote, relative to any individual, those things which *tasty* denotes relative to Mary:

- (4) $\llbracket \alpha \text{ for } \beta \rrbracket^{c, w, i} = \llbracket \alpha \rrbracket^{c, w, b}$, where $b = \llbracket \beta \rrbracket^{c, w, i}$

The *content* of an expression α in a given context c , notated $\llbracket \alpha \rrbracket^c$, will be that function which maps any world w and individual i onto the denotation of α relative to c, w, i ; that is, $\llbracket \alpha \rrbracket^c(w, i) = \llbracket \alpha \rrbracket^{c, w, i}$.

We say that a sentence φ is *true in context* c iff its content in c maps the world and judge of c onto 1, that is iff $\llbracket \varphi \rrbracket^c(W(c), J(c)) = 1$.

Relative Truth and Speaker Commitment

If p and q are two sentence contents (that is, two functions from world-individual pairs to truth values), we say that p and q *contradict each other* iff there is no world w and no individual i such that $p(w, i) = 1$ and $q(w, i) = 1$.

It is worth noting a few interesting features of this system. Predicates, including subjective predicates, do not vary in content from context to context; $\llbracket \textit{tasty} \rrbracket^c = \llbracket \textit{tasty} \rrbracket^{c'}$ for all c, c' . Likewise, sentences consisting just of a constant other than I and a predicate receive the same content in all contexts: $\llbracket \textit{tasty}(\textit{the-licorice}) \rrbracket^c = \llbracket \textit{tasty}(\textit{the-licorice}) \rrbracket^{c'}$ for all c, c' .

Nonetheless, sentences containing a subjective predicate may vary in truth value from context to context. If two contexts have different judges, i.e. if $J(c) \neq J(c')$, it might happen that $\llbracket \textit{tasty}(\textit{the-licorice}) \rrbracket^c(w, J(c)) = 1$ and $\llbracket \textit{tasty}(\textit{the-licorice}) \rrbracket^{c'}(w, J(c')) = 0$. Of course then it would also be the case that $\llbracket \textit{tasty}(\textit{the-licorice}) \rrbracket^c(w, J(c)) = 1$ and $\llbracket \neg \textit{tasty}(\textit{the-licorice}) \rrbracket^{c'}(w, J(c')) = 1$. That is, the sentence and its negation could both be true (relative to their separate contexts).

Despite this fact, $\llbracket \textit{tasty}(\textit{the-licorice}) \rrbracket^c$ and $\llbracket \neg \textit{tasty}(\textit{the-licorice}) \rrbracket^{c'}$ will always contradict each other, even when c and c' are distinct.

This, in a nutshell, is the explanation for “faultless disagreement.” In a system where context plays a role not just in assigning contents to sentences, but also in assigning truth values to contents, and where contexts are assumed to provide an individual on whose tastes the truth values of sentences may depend, it is possible for two sentences to express contradictory contents, yet each be true in its context.

As a final comment, we should acknowledge an obvious simplification in the system, namely that it makes no provision for the fact that people’s tastes can change over time. This could be addressed by adding a time index alongside the world and individual indices, as in Lasersohn (2005); by interpreting the individual index as fixed to a “time slice” of an individual rather than a temporally persistent individual; or by interpreting it as fixed to a system of criteria for judging tastiness, fun, etc., rather than an individual at all, and adding formalism to match such systems with the individuals who employ them, allowing for the possibility that a given individual might employ different systems at different times. The choice among these strategies is an interesting issue, but it is not our main concern in this paper, so let us set it aside for the present.

3. Adopting a Stance

The formal rules just presented leave a great deal unsaid about how truth evaluation works in real, practical situations. The truth values of sentences may vary according the “judge” of the context, but what determines who the judge is?

Peter Lasersohn

I think that normally there is *nothing* in the real, concrete situation in which a sentence is used that determines who the judge is; it is indeterminate who the judge is. It is precisely because of this that the truth of sentences like *Licorice is tasty* or *Roller coasters are fun* is a matter of subjective judgment, and not objective fact. If some objective matter of fact privileged some particular person as “the” judge of the concrete situation of some assertion of *Licorice is tasty*, then we would have to consider that assertion to be objectively true or false, according to whether it was true or false relative to this privileged person. Because assertions like this are not intuitively objective, our theory should not portray the concrete situation of use as determining a judge.

Our formalism, however, requires that each context specify a judge. Therefore, we cannot identify contexts in the formal sense with concrete situations of use. Instead, each such situation corresponds to any number of distinct formal contexts — one for each individual (or at any rate, one for each sentient individual) — and no one of these can be regarded as “the” context in which a sentence is used.¹

All this having been said, it must be recognized that when we do engage in the practical act of assessing a sentence for truth or falsity, the only way to obtain a determinate truth value is to assess the sentence relative to a specific individual. Normally, we all use ourselves for this purpose, or as I will put it, we adopt an *autocentric stance*. Whether we are evaluating the claims of others or considering the truth of what we say ourselves, whether we are the addressee of an utterance or an accidental overhearer, the normal pattern is to assess relative to our own perspective.

An autocentric stance is not the only one we can adopt, however. Realizing that in matters of taste, our own perspective is no more valid than anyone else’s, we may also adopt a “bird’s eye view,” or *acentric stance*, in which we do not select any particular individual as judge. Because a truth assessment can only be done relative to a particular choice of judge, adopting an acentric stance precludes a truth assessment. If we adopt an acentric stance, sentences like *Licorice is tasty* seem not to have a determinate truth value.

It is perhaps worth noting that it is only when we adopt an acentric stance that “faultless disagreement” really seems faultless. And indeed, as long as we maintain an acentric stance, we are hardly likely to engage in disputes over matters of taste. But if you take an autocentric stance and assess *Licorice is tasty* as true, and if someone else asserts its negation, then (barring a shift of stance) you must assess that assertion as false — hardly “faultless.”

In addition to autocentric and acentric stances, it is also possible under limited circumstances to adopt an *exocentric stance*, assuming someone other than ourselves as

¹Alternatively, we might regard the judge as provided by a separate “context of assessment” rather than the context of use, using the technique of MacFarlane (2003). So far as I can tell, this gives an equivalent effect to the system presented here.

Relative Truth and Speaker Commitment

the relevant judge. In this case, we “place ourselves in someone else’s shoes” for the purposes of assessing truth values. This is especially common in free indirect discourse, and also occurs regularly in questions, and when one person speaks for another. Occasionally, it also occurs in ordinary assertions, especially when a predicate like *fun* is ascribed to a particular event, as in *Mary rode the roller coaster yesterday; it was fun*. An event of Mary riding the roller coaster on one specific occasion would not normally be fun for anyone but Mary; nonetheless someone else might assess a claim that this event had been fun. It would be pragmatically pointless for a person who was uninvolved in the event to assess such a claim relative to himself or herself; instead, the normal practice in such cases is to assess it relative to a pragmatically salient person who was involved.

The possibility of taking an exocentric stance complicates somewhat the distinction between an analysis that makes use of a hidden indexical to relativize truth to individuals, and an analysis like that advocated here, which instead uses an individual index alongside the possible world index. If a speaker clearly adopts an exocentric stance in making an assertion, there may be a tendency for listeners to coordinate their stances to that of the speaker, so that all participants assess the assertion relative to the same individual x . In this case, the assertion will seem intuitively true iff it is true relative to x , and false iff it is false relative to x — just as though it involved indexical reference to x . For this sort of example, the difference between the two theories becomes rather subtle, to say the least.

But such coordination does not always occur. Particularly when the speaker adopts an autocentric stance, there is no general expectation, I think, for listeners to coordinate their stances with that of the speaker; and it may easily happen that all participants in the discourse assess autocentrically, according to their own tastes. Such lack of coordination does not indicate misunderstanding or misinterpretation, but simply reflects our natural tendency to look at things from our own perspective.

4. Attitudes and Stances

Because belief and many other attitudes involve truth assessment, issues of stance will naturally arise in the semantic analysis of attitude verbs. Suppose John rode the merry-go-round with his young son Bill. John was bored the whole time, but observed that Bill was enjoying himself. In this circumstance, John would probably assess *The ride was fun* as false from an autocentric stance, but as true from an exocentric stance with Bill as the judge. Does John believe the ride was fun? The answer depends on his stance.

To capture the idea that one may believe different things while adopting different stances, in we may treat *believe* and similar verbs as 3-place predicates, denoting relations between an individual, a sentence content, and a context which that individual assumes in assessing the sentence content. Where c_1 is a context whose judge is Bill, we

may now write *believe*(*John*, c_1 , \wedge *fun*(*the-roller-coaster*)) to mean that John assumes context c_1 in believing that the roller coaster is fun, hence that he adopts an exocentric stance with Bill as judge. Here, the \wedge -notation forms a term denoting the content of the sentence it combines with; that is:

$$(5) \quad \text{For all } c, w, i: \llbracket \wedge \varphi \rrbracket^{c, w, i} = \llbracket \varphi \rrbracket^c$$

The context argument of *believe* is not overtly expressed in English syntax, and its value is fixed pragmatically; this is a real implicit indexical, unlike the judge parameter in a simple sentence like *Roller coasters are fun*. But notice that the implicit argument is in the higher clause — it is not in the subordinate clause expressing the object of belief. In fact, the object of belief is the same regardless of stance, and in no way encodes which individual is the one on whose tastes the truth of the belief depends.

Some additional constraints on the denotation of *believe* are probably necessary, but need not concern us here; see Lasersohn (2005) for some discussion.

It is interesting to compare *believe* with the verb *consider* (and related verbs like *deem*, *regard*, and *judge*).² *Consider* is much more limited than *believe* in types of complement clause it may combine with. It combines quite naturally with clauses expressing personal taste, but normally does not combine with clauses expressing completely objective matters of fact:

- (6) a. John believes the licorice to be tasty.
b. John considers the licorice to be tasty.
- (7) a. John believes the licorice to contain sugar.
b. ? John considers the licorice to contain sugar.

Actually, *consider* combines with a much wider class of clauses than just those which express personal taste; but they all, it seems to me, involve some sort of evaluative judgment or decision on the part of anyone assessing them for truth:

- (8) a. John considers Bill to be a genius.
b. John considers Bill to fall under the scope of Regulation #215.
c. John considers Bill to be ready for the competition.

Clauses headed by a vague or scalar predicate are allowed, but not if a particular value on the scale is overtly fixed:

- (9) a. John considers Bill to be tall.
b. ? John considers Bill to be six feet, two inches tall.

²Thanks to Manfred Krifka for starting me thinking about *consider*; his student Julia Staffel discusses the German equivalent *finden* in some detail in work currently in progress.

Relative Truth and Speaker Commitment

I presume (despite Williamson 1994) that this is because an assessment of *Bill is tall* involves some sort of decision on the part of the assessor as to where to draw the line between the tall and the non-tall; it is not purely a matter of objective fact.

Another interesting difference between *believe* and *consider* is that what one believes may be false relative to oneself; but what one “considers” apparently cannot, at least in matters of taste.

For example, suppose John does not know the taste of licorice — he has never tried it, or perhaps has forgotten what it tastes like. Even in this situation, I would claim, if John is the sort of person who would thoroughly enjoy the taste of licorice if he did try it, the generic statement *Licorice is tasty* is true relative to John. But now suppose John has heard all sorts of terrible things about licorice. On that basis, he might not believe it is tasty, but just the opposite. Later, he might try licorice and discover that it is tasty after all; he had been mistaken.

But if John doesn't know the taste of licorice, then *John considers licorice to be tasty* is a distinctly odd thing to say, no matter how much he has heard about it (as would *John considers licorice not to be tasty*). Saying that *x* considers *p*, unlike saying that *x* believes *p*, presupposes that *x* has the necessary knowledge for the kind of judgment that determines, and doesn't merely describe, truth relative to *x* in matters of taste.³ For this reason, a sentence like *John considers licorice to be tasty* does not lend itself easily to an interpretation where *Licorice is tasty* might be false relative to John.

Of course it might happen that John considers licorice not to be tasty, then later changes his mind and considers licorice to be tasty, so that he regards his earlier opinion as false. But this is a case of John's tastes changing, not of John belatedly discovering that the flavor of licorice appeals to what his tastes had been all along.

Probably because verbs like *consider* represent this kind of “truth-determining” judgment on the part of the subject, rather than a mere assessment for truth, they do not normally allow readings in which the subject takes an exocentric stance. Returning to the case where John is bored on the merry-go-round, but observes that Bill is having fun, it does not seem true to say *John considers the merry-go-round to have been fun*, even

³I don't mean to imply that the act of judgment itself brings it about that *p* is true relative to *x*; rather, *p* is true relative to *x* iff *x* would make this sort of judgment, if *x* had the appropriate kind of knowledge. Exactly how to characterize this kind of knowledge is a difficult question; how we answer it will depend in part on exactly which issues we consider to be matters of taste. From examples like *John considers licorice to be tasty*, it appears that the knowledge must be based directly on personal experience, such as tasting licorice. But a sentence like *John considers George Bush to be a terrible president* might be true even if John's only knowledge of George Bush and his actions comes from newspapers.

though *John believes the merry-go-round to have been fun* does admit a reading in which it is true in this situation.⁴

Because verbs like *consider* obligatorily require the subject to take an autocentric stance, unlike verbs like *believe*, there is no motivation to treat them as 3-place predicate rather than 2-place. We now write *consider(John, ^tasty(the-licorice))* to mean that John considers the licorice to be tasty.

At this point, perhaps it should be pointed out that the analysis I am suggesting for verbs like *consider* is formally very similar to that suggested by Lewis (1979) for *de se* belief. Lewis suggested that the objects of belief were sets of “centered worlds,” where centered worlds were essentially just pairs of a world and an individual.⁵ Of course this is exactly what I am suggesting for the content of sentences in general, including complement clauses to attitude verbs. This similarity in formal techniques should not be too surprising, since in both cases the point is to represent the dependency of truth on perspective.⁶

But beyond this general similarity, it should be noted that autocentric stances are actually *de se*. For example, suppose some portion of the roller coaster is adjacent to a large mirror. While riding the roller coaster past the mirror, John mistakes his own reflection for someone else, realizes that “that person” must enjoy roller coasters, and assesses *Roller coasters are fun* as true relative to “that person.” Surely we should count this as taking an exocentric stance rather than an autocentric one.

The *de se* status of autocentric stances suggests that it might be worthwhile to rethink the analysis of *believe*. Although *believe* allows an exocentric stance, the most typical understanding of a sentence of the form *x believes p* is certainly still one in which *x* adopts an autocentric stance towards *p*. Should we then represent sentences used this way as *believe(x, p)*, on the model of *consider*, or as *believe(x, c, p)*, retaining the implicit context argument as above? The latter formula certainly does not capture the *de se* nature of an autocentric stance. However, it seems reasonable to assume that taking a *de se* attitude automatically involves taking a *de re* attitude toward oneself (though not vice versa), and this raises the possibility that sentences of this form truth-conditionally have only a *de re*, exocentric reading, with the *de se*, autocentric interpretation being purely a matter of pragmatic implicature. I regard this issue as open, and will not attempt to answer it here; either answer is compatible with the main claims of this paper.⁷

⁴Perhaps marginally, *John considers the merry-go-round to have been fun* could be understood as true in this situation if the issue he was considering was not one of taste in entertainment, but of what degree of fun is necessary for something to count as fun, i.e. of where to draw the border for *fun* construed as a scalar predicate.

⁵Or more accurately for Lewis, pairs of a world and a space-time location. The term *centered world* originates with Quine (1969).

⁶Moltmann (2006) also stresses the connection between predicates of personal taste and *de se* attitudes.

⁷Some preliminary evidence in favor of the existence of an authentically autocentric, *de se* reading

5. Attitude and Stance in a Non-Relativist Semantics

If we consider these same issues from the standpoint of a theory which does not use an individual index alongside the possible world index, but instead assumes that predicates like *fun* and *tasty* take a hidden indexical argument, we arrive at a somewhat different picture. In this sort of analysis, we need no special notion of “stance”; the stance of a truth assessor is just his or her identification of a referent for the hidden indexical, no different from the identification of referents for any other indexical expression.

Adding this hidden argument for *fun* and *tasty* means that we can eliminate the argument in *believe* which was used to encode the subject’s stance. Now in the situation where John is bored on the merry-go-round, but observes that Bill is enjoying himself, we can say simply that John stands in the belief relation to two separate sentence contents: *believe(John, \neg fun(the-merry-go-round, John))*, but *believe(John, \wedge fun(the-merry-go-round, Bill))*.⁸

If we hold strictly to the idea that predicates like *tasty* and *fun* have a hidden argument whose value is fixed by the context in the same way as an ordinary indexical pronoun, then we should also write *consider(John, \wedge tasty(the-licorice, John))* for *John considers the licorice to be tasty*. But we should note right away that this formula does not adequately capture the *de se* nature of this sentence’s interpretation. A *de se* reading for this sentence falls out naturally under the relativist analysis, which is committed for independent reasons to the same sort of formalism which Lewis developed for *de se* attitudes; but the non-relativist analysis must make some special provision here.

Moreover, to adequately capture the *de se* reading it will probably be necessary to make some concessions to the relativist analysis. Appealing to a theory of *de se* attitudes like that of Lewis (1979), means adopting essentially the formalism I have been advocating, at least for this construction. Other approaches to *de se* pronouns also share crucial features with the relativist analysis. For example, we could “diagonalize” on the hidden argument of *tasty* in the manner of Stalnaker (1999). This would mean treating each possible world as specifying a context, and each context as specifying a value for the hidden argument. Then we analyze the subordinate clause *licorice to be tasty* as denoting the set of possible worlds *w* in which licorice is tasty for whomever is specified as the value of the hidden argument in the context specified by *w*. Like the relativist analysis, this treats *licorice to be tasty* as varying in truth value according to whose tastes

of *believe*, with no implicit context argument, comes from the fact that the context assumed by the believer need not be held constant across verb phrase anaphora: If John and Bill both autocentrically assess *The roller coaster is fun* as true, one may say *John believes the roller coaster is fun, and Bill does too*, even though John and Bill assume separate contexts with different judges.

⁸Of course all our symbols will have to be reinterpreted in such a theory; for example the rule for the \wedge -notation should be: $\llbracket \wedge \phi \rrbracket^c = \llbracket \phi \rrbracket^c$.

are considered, but without a concomitant variation in content — two of the most central features of the relativist analysis.⁹

The non-relativist analysis also poses potential problems for syntactic theory. In this line of analysis, a parallel emerges between verbs like *consider*, and “control” verbs like *want*. It has long been observed¹⁰ that the implicit subject of an infinitival clause embedded under *want* or a similar verb must be interpreted like a *de se* pronoun dependent on the matrix subject; indeed this appears to be the only construction which uncontroversially and unambiguously receives a *de se* reading. But if predicates of personal taste contain a hidden indexical argument, this too must be interpreted like a *de se* pronoun dependent on the matrix subject, at least when embedded under verbs like *consider*. This dependency would appear to be a case of control, and verbs like *consider* would therefore appear to be control verbs. However, control of non-subject arguments is not expected under most syntactic theories, and to accommodate it would require substantial revision to these theories.

To summarize, two main differences should be noticed between the relativist analysis of Section 4 and the indexical analysis of Section 5. First, in the relativist analysis there is no hidden argument for predicates of personal taste, but there is a hidden argument for certain attitude verbs, such as *believe*; the hidden argument is in the higher clause. In the indexical analysis, there is a hidden argument for predicates of personal taste, but not in verbs like *believe*; the hidden argument is in the lower clause. Second, in the relativist analysis, selection of a relevant individual for assessing sentences about personal taste does not affect the content of such sentences, but is purely a matter of the “stance” a truth assessor or attitude-holder takes toward that content. In the indexical analysis, this selection does play a role in determining content, and is accomplished by the ordinary means of reference-fixing for indexicals. However, this will not produce a *de se* reading for clauses subordinated to verbs like *consider*. Accounting for this reading may require significant concessions to the relativist analysis, and may also pose problems for existing syntactic theories.

6. Factive Constructions, Content and Speaker Commitment

Additional evidence for the relativist approach, and for a tight connection between issues of taste and *de se* attitudes, may be found by examining sentences whose truth commits the subject to some subordinate clause, and whose felicitous assertion also commits the speaker to that same clause. Such examples are well-known and well-studied under the

⁹Adopting a Lewis- or Stalnaker-style analysis of *de se* attitudes does not commit us to a fully relativist semantics, even for the complement clause of verbs like *consider*, if the individual index or value of the hidden argument is objectively fixed by the context of use. (Instead, we have what MacFarlane (to appear) calls a “non-indexical contextualist” analysis.) But such an analysis does share crucial features with a relativist approach, and the fact remains that a *de se* reading is expected under the relativist analysis, but requires special provision in an analysis which appeals to a hidden indexical.

¹⁰This observation goes back at least to Morgan (1970), though not using this terminology.

Relative Truth and Speaker Commitment

heading of *factive* predicates (Kiparsky and Kiparsky 1971 and much subsequent literature).

A factive predicate is one which takes a clausal complement, and whose felicitous use presupposes that this clausal complement is true.¹¹ Standard examples include *regret*, *recognize*, *forget*, etc. A speaker who asserted (10), for example, would normally be understood to presuppose that John's theory had flaws.

(10) John recognizes that his theory has flaws.

That this information is presupposed is easily seen by its preservation under negation:

(11) John doesn't recognize that his theory has flaws.

It should be acknowledged that under the right discourse conditions, these predicates may be used *without* commitment on the part of the speaker to the truth of their complements. For example, a speaker may assert (12) without committing to the claim that John stole the money:

(12) John imagines that he stole the money, and regrets that he did it.

I assume that a theory of presupposition projection will account for the pattern of when factive predicates do and when they don't carry this presupposition. The details of such a theory need not concern us here; see Beaver (1997) for an overview of different approaches. The important point is that the existence of examples like (12) should not keep us from recognizing the fact that ordinarily, factive predicates do commit the speaker to the truth of their complement clauses.

Nor will it matter much for our purposes whether this commitment is part of the truth conditions of sentences containing factive predicates, or pragmatic implicature, or something else. I have referred (and will continue to refer) to this commitment as required for the "felicitous use" of factive predicates; but the crucial point for our purposes again is just that the commitment exists, not its precise analysis.

Many factive predicates, though not all, also semantically entail that their *subjects* are committed to the truth of their complement clauses. For example, (10) entails that John believes that his theory has flaws; likewise, if John regrets that he stole the money, then John must believe that he stole the money.

¹¹Factive predicates are also sometimes defined syntactically, in terms of their ability to combine with the noun *fact* and with gerundive complements, their inability to combine with infinitival clauses, etc. As Kiparsky and Kiparsky (1971) point out, a few verbs, including *know* and *realize*, are semantically factive, but do not conform to the normal factive syntactic pattern. For our purposes here, it is the semantic and pragmatic constraints on factive predicates which will be important, not their syntax.

The fact that some factive predicates commit both the subject and the speaker to the truth of the complement clause makes them an interesting test case for the semantic analysis of predicates of personal taste. In order for a sentence containing such a predicate to be true and felicitously used, both the subject and the speaker must assess the content of the subordinate clause as true. In a relativist analysis like the one sketched in Section 4, both the subject and the speaker must adopt a stance in doing so; and nothing prevents them from adopting stances with different judges — in fact, since an autocentric stance is the norm, we should expect that most typically, the subject and the speaker would each assess the subordinate clause as true using himself or herself as the relevant judge. Our expectation, then is that (13) will have a reading which suggests that *Licorice is tasty* as an expression both of John's taste and of the speaker's — and in fact, that this will be the default interpretation:

(13) John recognizes that licorice is tasty.

This expectation is borne out; a speaker who asserted this sentence would ordinarily be understood as liking the taste of licorice, and as claiming that John liked it too.

Now consider this example in an indexical analysis like that sketched in Section 5. Here, *tasty* has a hidden argument whose referent is fixed like an ordinary indexical, so that the choice of referent helps determine the content of the subordinate clause. If the referent of that argument is fixed to John, then the subordinate clause means that licorice is tasty for John, and (13) should be understood as implying that the speaker believes licorice is tasty for John (though not necessarily for the speaker), and that John also believes this. Or, if the referent is fixed to the speaker, the interpretation should be that the speaker believes licorice is tasty for himself or herself, and that John also believes it is tasty for the speaker. However, under this analysis, there seems to be no way to obtain a reading which implies that the speaker believes licorice is tasty for the speaker, and John believes that licorice is tasty for John.

Such a reading might be obtained by adopting a suitable theory of *de se* attitudes, and treating the hidden argument of *tasty* as a *de se* pronoun — but as already pointed out, this would require some special provision under a theory in which predicates like *tasty* have a hidden indexical argument, while it is expected under a relativist analysis; it would also mean adopting several major points of the relativist analysis, and would pose problems for the syntactic theory of control.

One might suspect that the indexical analysis could be defended by claiming that the hidden indexical is fixed to a *group* containing both the speaker and the subject, and adopting a principle that if something is tasty for a group, it is tasty for all members of the group. Then (13) would imply that both John and the speaker were committed to the claim that licorice is tasty for this group, hence both for John and for the speaker. That

Relative Truth and Speaker Commitment

this is not the right solution, however, is easily seen by considering the following scenario: John considers licorice to be tasty, and erroneously believes that he is the only person to do so. Unbeknownst to John, the speaker also considers licorice to be tasty, and knows that John does too. In this situation, surely the speaker could truthfully and felicitously say *John recognizes that licorice is tasty*, but it is not the case that both John and the speaker are committed to the claim that licorice is tasty for a group containing both John and the speaker.

The indexical analysis faces a serious problem, then. In this analysis, the content of a clause containing a predicate of personal taste varies according to the value of the hidden argument; what the speaker and subject are both committed to should be truth of this content, *as determined by fixing this argument to a particular value*. As long as this value is fixed like an ordinary indexical, there seems to be no way to capture the most typical understanding of this sentence, in which *Licorice is tasty* is an expression both of speaker's individual tastes and of John's.

It should be acknowledged that this reading may not be the only possible interpretation for this sort of factive sentence; and it may even be that (13) can be understood as meaning that John and the speaker are both committed to the claim that licorice is tasty for John, or as both committed to the claim that licorice is tasty for the speaker, or as both committed to the claim that licorice is tasty for some other individual or group, as expected under the indexical analysis. But the relativist analysis can easily allow for such interpretations, since either John or the speaker may adopt an exocentric stance, assessing *Licorice is tasty* as true using the other one or some third individual or group as the relevant judge. The problem is not that the indexical theory predicts bogus readings, but that it fails to predict a reading the sentence actually has — in fact, I think, its most prominent, typical reading.¹²

7. Truth Evaluative Adverbs

A similar argument may be constructed using what I will call “truth evaluative adverbs.” Here, I mean adverbs such as *correctly*, *incorrectly*, *accurately*, *inaccurately*, *mistakenly*, etc., used to modify speech act verbs or attitude verbs. Somewhat like factive predicates, these adverbs express a commitment on the part of the speaker to a particular truth value for the complement clause to the verb they modify. For example, (14)a. commits the

¹²In a recent paper, von Stechow and Gillies (2006) observe that one can say *If Blofeld realizes you might be in Zürich, you can breathe easy* even in a situation where it is known to the speaker that the addressee is not in Zürich. They cite this as a problem for relativist theories of epistemic modals, according to which, they claim, the possibility of the addressee being in Zürich must be compatible with both the speaker's and Blofeld's epistemic states. But this reading falls out naturally in a relativist theory with exocentric stances, while the sort of theory von Stechow and Gillies favor, in which epistemic modals make indexical reference to a group, would seem to face the same problems just pointed out for a group-indexical analysis of predicates of personal taste.

speaker to the claim that *Bill stole the money* is true; (14)b. commits the speaker to the claim that this sentence is false.

- (14) a. John correctly believes that Bill stole the money.
b. John incorrectly believes that Bill stole the money.

Of course these sentences also imply that the subject is committed to the truth of the subordinate clause.

Now let us turn to examples in which the subordinate clause concerns an issue of taste:

- (15) a. John correctly believes that licorice is tasty.
b. John incorrectly believes that licorice is tasty.

Here we find again that a possible interpretation — indeed, the preferred interpretation — is that John is committed to the truth of *Licorice is tasty*, judged according to his own tastes; and that the speaker is committed to the truth of *Licorice is tasty* in (15)a., and to the falsity of this same sentence in (15)b., judged according to the speaker's own tastes. This reading is expected under the relativist analysis, but unexpected in an analysis which treats *tasty* as having a hidden argument whose value is fixed in the same way as an ordinary indexical. On such a theory, the expected readings are ones in which John and the speaker both assume the *same* value for this hidden argument in committing to the truth or falsity of the subordinate clause.

8. Summary and Conclusion

We compared two ways of relativizing the truth values of sentences containing predicates of personal taste: The first, or “indexical,” approach analyzed such predicates as containing a hidden indexical argument, so that the same sentence may express different contents depending on how the value of this indexical is fixed; the other, or “relativist,” approach assumed an individual index, analogous to a possible-world index, so that truth values could vary without a corresponding variation in content. The relativist approach bears a formal similarity to Lewis' theory of *de se* attitudes. It also necessitates a theory of stances which an individual may adopt in assessing a sentence content for truth or falsity, including an autocentric, or *de se*, stance. Because attitude verbs often imply truth evaluation by their subjects, the stance of the subject must be taken into account in the semantic analysis of these verbs.

We observed that some verbs, such as *consider*, imply an autocentric stance on the part of the subject, the *de se* nature of which is expected under the relativist analysis, but not under the indexical analysis. Moreover, revising the indexical analysis to account for this *de se* interpretation would appear to concede a number of important points to the relativist analysis, and pose problems for syntactic theory.

Relative Truth and Speaker Commitment

We also saw that sentences containing certain factive predicates with a complement clause imply a commitment to the truth of the complement clause on the part of both the speaker and the subject. When the complement clause contains a predicate of personal taste, the speaker and subject may each judge it to be true according to his or her own separate tastes, as expected under the relativist analysis. But under the indexical analysis, we should expect that the speaker and subject both judge the clause to be true using the *same* value for the hidden indexical, an expectation which is not borne out.

A similar argument was made using truth evaluative adverbs. Like factive predicates, these imply a commitment on the part of the speaker, and they may combine with verbs which imply a commitment on the part of the subject for the same clause. The speaker and subject may each assess this clause according to his or her own tastes, as expected under the relativist analysis. The pattern expected under the indexical analysis, namely that the speaker and subject both assess the clause using the same value for the hidden indexical, is not the only or even the preferred interpretation.

None of these arguments appeals directly to the idea of “faultless disagreement,” which has served as the primary motivation for recent relativist analyses. They therefore constitute independent evidence for this sort of analysis; and a defense of a conventional, non-relativist semantic theory must do more than simply provide an alternative explanation for faultless disagreement.

In presenting these arguments, I compared the relativist analysis only to one other approach — the hidden indexical approach. Of course this is not the only alternative to relativism, so these arguments should not be seen as fully establishing the relativist position. But they do, I think, give some reason for preferring it to one of the more obvious alternatives, and present some additional challenges which any analysis should account for.

References

- Beaver, David. 1997. Presupposition. In *Handbook of logic and language*, ed. Johan van Benthem and Alice ter Meulen, 939–1008. Cambridge, Mass.: MIT Press.
- Egan, Andy, John Hawthorne and Brian Weatherson. 2005. Epistemic modals in context. In *Contextualism in philosophy: knowledge, meaning and truth*, ed. Gerhard Preyer and Georg Peter, 131–168. Oxford: Oxford University Press.
- von Fintel, Kai and Anthony S. Gillies. 2006. CIA leaks. Ms., MIT, Cambridge, Mass., and University of Michigan, Ann Arbor.
- Kaplan, David. 1989. Demonstratives: an essay on the semantics, logic, metaphysics, and epistemology of demonstratives and other indexicals. In *Themes from*

Peter Lasersohn

- Kaplan*, ed. J. Almog, J. Perry and H. Wettstein, 481–563. Oxford: Oxford University Press.
- Kipasky, Paul and Carol Kiparsky. 1971. Fact. In *Semantics: an interdisciplinary reader*, ed. Danny D. Steinberg and Leon A. Jakobovits, 345–369. Cambridge: Cambridge University Press.
- Kölbel, Max. 2002. *Truth without objectivity*. London: Routledge.
- Kölbel, Max. 2003. Faultless disagreement. *Proceedings of the Aristotelian Society* 104: 53–73.
- Lasersohn, Peter. 2005. Context dependence, disagreement, and predicates of personal taste. *Linguistics and Philosophy* 28: 643–686.
- Lewis, David. 1979. Attitudes *de dicto* and *de se*. *The Philosophical Review* 88: 513–543.
- MacFarlane, John. 2003. Future contingents and relative truth. *The Philosophical Quarterly* 53: 321–36.
- MacFarlane, J. (to appear). Semantic minimalism and nonindexical contextualism. In *Context-sensitivity and semantic minimalism: new essays on semantics and pragmatics*, ed. G. Preyer & G. Peter. Oxford: Oxford University Press.
- Morgan, Jerry. 1970. On the criterion of identity for noun phrase deletion. In *Papers from the Sixth Regional Meeting, Chicago Linguistic Society*, 380–389. Chicago Linguistic Society, University of Chicago, Chicago, Ill.
- Moltmann, Friederike. 2006. First-person oriented genericity and relative truth. Ms., IHPST, Paris.
- Quine, W.V. 1968. Propositional objects. *Critica* 2: 3–22.
- Richard, Mark. 2004. Contextualism and relativism. *Philosophical Studies* 119: 215–242.
- Stalnaker, Robert. 1999. *Context and content: essays in intentionality in speech and thought*. Oxford: Oxford University Press.
- Williamson, Timothy. 1994. *Vagueness*. London: Routledge.

University of Illinois
Department of Linguistics
707 South Mathews Avenue
Urbana, IL 61801

lasersoh@uiuc.edu