| Title | De qualitate generalized |
|---|---|
| Sub Title | |
| Author | Tancredi, Christopher |
| Publisher | 慶應義塾大学言語文化研究所 |
| Publication year | 2019 |
| Jtitle | 慶應義塾大学言語文化研究所紀要 (Reports of the Keio Institute of Cultural and Linguistic Studies). No.50 (2019. 3) ,p.91- 154 |
| JaLC DOI | 10.14991/005.00000050-0091 |
| Abstract | In this paper I argue for a new analysis of belief statements in which the complement clause denotes not the putative content of a person's beliefs but a proposition that can be inferred from that content. The notion of inference that makes this analysis possible operates within a world rather than across worlds, setting it apart from the possible worlds analysis of entailment. Our analysis unifies standard de dicto, de re and de qualitate attributions while also distinguishing among beliefs in different necessities or impossibilities and accounting for possible and impossible inferences that can be drawn from such beliefs. I accept that belief statements are systematically ambiguous, but the only ambiguity unique to attitude attributions I take to be the de translato/non-de translato ambiguity argued for by Tancredi and Sharvit (forthcoming) : the embedded clause can be interpreted as the attributor takes the subject to interpret it (de translato) or as the attributor herself interprets it (non-de translato). I show how our analysis leads to the dissolution of Kripke's puzzle about belief while making it possible to reject Putnam's and Burge's claims that beliefs are not in the head. |
| Notes | 論文 |
| Genre | Departmental Bulletin Paper |
| URL | https://koara.lib.keio.ac.jp/xoonips/modules/xoonips/detail.php?koara_id=AN00069467-00000050-0091 |

*Reports of the Keio Institute of Cultural and Linguistic Studies* 50 (2019), 91~154

# De Qualitate Generalized

Christopher Tancredi

**Abstract**

In this paper I argue for a new analysis of belief statements in which the complement clause denotes not the putative content of a person's beliefs but a proposition that can be inferred from that content. The notion of inference that makes this analysis possible operates within a world rather than across worlds, setting it apart from the possible worlds analysis of entailment. Our analysis unifies standard *de dicto*, *de re* and *de qualitate* attributions while also distinguishing among beliefs in different necessities or impossibilities and accounting for possible and impossible inferences that can be drawn from such beliefs. I accept that belief statements are systematically ambiguous, but the only ambiguity unique to attitude attributions I take to be the *de translato*/non-*de translato* ambiguity argued for by Tancredi and Sharvit (forthcoming): the embedded clause can be interpreted as the attributor takes the subject to interpret it (*de translato*) or as the attributor herself interprets it (non-*de translato*). I show how our analysis leads to the dissolution of Kripke's puzzle about belief while making it possible to reject Putnam's and Burge's claims that beliefs are not in the head.
**Key words**: belief, attitude, attribution, inference, *de dicto*, *de re*, *de qualitate*, *de translato*

## 1 Introduction

According to Burgess's (2013, p.79) pithy summarization of belief statements, "To believe that snow is white is to have an attitude of belief toward, or stand in the 'belief-relation' to, the proposition that snow is white." While there is disagreement on many details, the general view that belief statements relate a believer to a belief designated by the verb's complement clause is widely agreed on by semanticists and philosophers alike, and so I will refer to it as *the standard view*. In this paper I take issue with the standard view. In particular, though I accept that belief statements relate individuals to propositions, I reject the view that the relation in question is that of the belief holder standing in the belief relation to the proposition identified by the complement clause. I argue instead that the internal belief of the belief holder enters the semantics only as the source from which the proposition denoted by the complement clause can be inferred. This minor difference in the semantics of attitude attributions will be seen to have far-reaching consequences.

Modern implementations of the standard view are largely based on Hintikka (1962). Hintikka took propositions to be sets of worlds and showed how to construct a semantics of belief statements upon such an assumption. The core idea behind his analysis of belief statements is that they compare the set of worlds compatible with an individual's beliefs, i.e. that individual's belief worlds, to the proposition, or set of worlds, denoted by the complement clause. Formally, where $[\![\ ]\!]^w$ is a function that takes a syntactic expression and gives back its denotation at world w,

(1) $[\![$ a believes that S $]\!]^w = 1$ iff a's belief worlds in w are a subset of $\{w' : [\![ S ]\!]^{w'}=1\}$.

This approach to the semantics of belief statements adopts standard assumptions from possible worlds semantics. Of particular importance to our concerns is its adoption of the possible worlds characterization of entailment: *S entails S′* iff the set of worlds in which *S* is true is a subset of the set of worlds in which *S′* is true. Based on this characterization I will somewhat loosely say that for Hintikka, *a believes that S* is true iff the sum total of *a*'s beliefs entails that S.[1]

1 The looseness lies in at least two places: (i) entailment is a relation between propositions, but a set of beliefs is distinct from the proposition it determines, and (ii) the proposition denoted by *that S* for any sentence *S* generally depends on the context in which *S* occurs, if for nothing other than fixing the temporal reference of any contained tenses. More adequate would be to say that the proposition determined by the conjunction of the sum total of *a*'s

Applying this analysis to a simple case of belief gives the following:

(2) $[\![$ John believes that the world is flat $]\!]^{@} = 1$ iff
in every world in which all of John's (actual world) beliefs are true, the world is flat.
(Here and throughout, @ = the actual world, i.e. the world of utterance)

This will be true if the sum total of John's beliefs entails that the world is flat, and false otherwise.

While Hintikka's approach of analyzing belief statements in terms of possible worlds is widely adopted, there are several well-known problems that show his particular analysis to be unsustainable without modification. Here I highlight two such problems. The first problem is that the analysis only works for attributing beliefs to people whose beliefs are taken to be internally consistent. This is because a set of internally inconsistent beliefs is true at no possible worlds, so the belief worlds of someone whose beliefs are internally inconsistent is the empty set. Coupled with the possible worlds analysis of entailment, this leads to the prediction that where *a* is known to have internally inconsistent beliefs, *a believes that S* is true regardless of what sentence *S* is. This is so since the empty set, identified with *a*'s belief worlds in this case, is a subset of all sets, including $\{w':[\![S]\!]^{w'}=1\}$ for all sentences S. Observationally, however, a person's indicating that they have inconsistent beliefs is not grounds for attributing absolutely any belief to them whatsoever, even if we take their inner beliefs to be as they indicate.

Some caution is needed in drawing conclusions from the observation here. Suppose that Mary utters "This glass is both full to the brim of water and not full to the brim of

beliefs entails the proposition denoted by *that S* in the context in which it occurs. I will stick with the simpler version in the text for expository convenience.

water." I check on her usage of the words in her utterance and find no cause for doubting that she means what I would mean by the sentence she utters aside from the obvious fact that the sentence is internally contradictory. Can I take this as establishing that Mary has inconsistent beliefs? Not conclusively. There is always a possibility that she is making a distinction in her usage that I simply didn't think to check. Because of this, I cannot know with certainty what the exact beliefs are that underlie her utterance and hence cannot know with certainty that she holds inconsistent beliefs. However, I can still show that there is a problem with Hintikka's semantics by exploring the speaker's beliefs. Given Hintikka's analysis of belief statements, a speaker who attributes a belief to Mary will be committed to the attribution being true. This can be modeled by saying that the attribution is true in all of the belief worlds of the speaker. Now, if the speaker believes that Mary has inconsistent beliefs, perhaps on the basis of Mary's utterance and the speaker's subsequent investigation into her meaning, then in all worlds compatible with what the speaker believes, Mary has inconsistent beliefs. The speaker is then predicted by Hintikka's analysis to be in a position to attribute any belief whatsoever to Mary. No speaker, however, will ever conclude from the premise that Mary has inconsistent beliefs that attribution of any belief whatsoever to Mary is licensed.

The second problem with Hintikka's analysis is that it fails for attributions of belief in necessary truths and necessary falsities. Under a possible worlds analysis, a sentence that is necessarily true denotes the set of all worlds. Since every set of worlds is a subset of the set of all worlds, the analysis predicts that everybody, and in fact everything (including non-sentient objects[2]), believes every necessary truth. Given the standard assumption that mathematical truths are necessary truths, this leads to the

2 Presumably, a non-sentient object has no beliefs. The set of worlds consistent with this absence of beliefs is then the set of all worlds. The conclusion that non-sentient objects believe all necessary truths follows since the set of all worlds, like all sets, is a subset of itself.

prediction that everything and everyone believes all mathematical truths. For instance, given that the sentence *$i^i$ is a real number* denotes a necessary truth, where *i* is the square root of negative one, Hintikka's analysis entails that any sentence of the form *a believes that $i^i$ is a real number* is true, where *a* ranges over people, coffee cups, and all other entities of type *e*. A sentence that is necessarily false, in contrast, denotes the empty set under Hintikka's analysis. Since the only set that is a subset of the empty set is the empty set itself, the only way for a person to believe something that is necessarily false under this analysis is for that person to have inconsistent beliefs. For instance, for any individual *a*, the claim that *a believes that $i^i$ is an imaginary number* will be tantamount to a claim that *a* has inconsistent beliefs. But this then brings us back to the first problem: anyone who can truthfully utter such a belief attribution should also accept as true any belief attribution whatsoever made to *a*. These predictions about attitudes toward necessary truths and falsities once again go strongly against intuition.

The literature is littered with analyses that attempt to overcome these problems. The two most prominent approaches to dealing with them are (i) to take propositions to be structured and analyze belief in a proposition as sensitive to the structure of the proposition (Cresswell & von Stechow 1982, Cresswell 1985, King 2007, 2014, 2017), and (ii) to admit impossible worlds into one's ontology so that different necessary truths can be distinguished from one another by the impossible worlds in which they are true (Jago 2014, Yagisawa 2010, Berto 2013). In this paper I will argue for a third alternative, namely that the underlying assumption of the role *S* plays in the semantics of a sentence of the form *a believes that S* is mistaken. In particular, I follow Bach (1997) in arguing that S does not directly identify the content of a belief of *a*'s. Our analysis will share with approach (i) above the idea that the objects of belief are structured. However, unlike standard structured propositions approaches to belief, I will provide a principled analysis of how and when belief in one structured object can justify attribution of belief in another structured object. Among other things, this will allow us to account for connections among impossible beliefs, i.e. beliefs in impossibilities, in ways that previous analyses cannot.

An analysis of the semantics of *a believes that S* that takes the content of S to be entailed by the sum total of *a*'s beliefs faces a wide range of challenges in addition to, and largely stemming from, the problems facing Hintikka's original analysis. On the philosophical side, it leads to Kripke's (1979) puzzle about belief, to Putnam's (1975) and Burge's (1979) conclusion that beliefs are not in the head, and retroactively to problems highlighted by Quine (1956) with quantification into opaque contexts. On the semantic side, facing up to Quine's concerns has historically led to a proliferation of interpretations, including *de re* and *de qualitate* interpretations on top of the Hintikkan *de dicto*, but in a way that sheds no light on Kripke's, Putnam's and Burge's concerns. In this paper I aim to show that there is a better way. In particular, I show how changing the underlying assumption of the role that *S* plays in *a believes that S* makes it possible on the philosophical side to wipe away Kripke's, Putnam's, Burge's and Quine's concerns in one fell swoop, and on the semantic side to give a single unified analysis of belief statements that covers *de dicto*, *de re* and *de qualitate* attributions. While I accept that belief statements are systematically ambiguous in ways that are independent of scope, the only ambiguity unique to attitude attributions I take to be that between *de translato* interpretation, argued to be necessary independently of any other potential ambiguities by Tancredi and Sharvit (forthcoming), and non-*de translato* interpretation. There is no additional *de dicto* / *de re* / *de qualitate* ambiguity.

In order to establish the need for a new analysis of belief attribution, I will proceed by examining the cases used to argue for separate *de re* and *de qualitate* interpretations. I accept that such additional interpretations (or other interpretations that do the same work) are required if one accepts the assumption that *S* in *a believes that S* identifies the content of a's purported belief. However, I also show that these modes of interpretation as formulated in the literature are not sufficient to account for the full range of attested belief statements. At the very least, an additional interpretation is needed to handle cases of *de qualitate*-like interpretation where the relevant property-denoting expression occurs in a downward entailing environment. I then show how all of the examples used to motivate treating belief statements as ambiguous can be given a unified semantics by changing the assumed role that *S* above plays. In a nutshell, I propose that the semantics of *a believes that S* requires only that the proposition *p* denoted by *S* be **inferable** from (not entailed by) a belief taken to be held by *a*. The differences among *de re* and *de qualitate* attributions will be seen to reduce to differences in how the required inference is drawn. The case in which *p* is identical to a belief taken to be held by *a*, i.e. what is generally referred to as *de dicto* attribution, will turn out to be the limiting case in which the inference is trivial.

The need for a separate *de re* interpretation under the standard view was first made clear by Quine (1956). Consider the following situation, slightly adapted from Quine:

(3) Situation A:
Ralph sees a man in a dark alley involved in a suspicious-looking activity. He says: "The man in the alley is a spy." Separately, he says: "Ortcutt is not a spy." I know, but Ralph does not, that the man in the alley is Ortcutt.

In this situation, taking his utterance as indicating his underlying beliefs, I can report on Ralph's beliefs by uttering the following:

(4) *Ralph believes that Ortcutt is a spy.*

To be sure, choosing to report on the situation in this particular way could well be misleading, and some effort will have to be expended to avoid misleading one's audience. But for someone fully apprised of the facts the report will generally be accepted as true. This example poses *prima facie* difficulties for Hintikka's analysis. Under that analysis, (4) is analyzed as follows:

(5) $[\![$ Ralph believes that Ortcutt is a spy $]\!]$ = 1 iff
in all worlds in which everything Ralph believes is true, Ortcutt is a spy.

Since Ralph believes that Ortcutt is not a spy, there are no worlds in which everything Ralph believes is true and in which Ortcutt is a spy. If the speaker takes Ralph to also believe that Ortcutt is a spy then she will take Ralph's beliefs to be inconsistent, in which case Hintikka's analysis correctly predicts that she should hold that (4) is true, but also predicts that she should hold any attribution of belief to Ralph at all to be true. On the other hand, if she fails to take Ralph to also believe that Ortcutt is a spy, then Hintikka's analysis predicts the sentence to be false. It does not generate a true interpretation that does not rely on Ralph having internally inconsistent beliefs. Ralph's beliefs, however, cannot, at least based only on the situation as described above, be taken to be internally inconsistent. Indeed they would all be true in a situation in which the man in the alley is a spy, the man in the alley is not Ortcutt, and Ortcutt is not a spy. The most that can be said against Ralph is that his beliefs look to be inconsistent with the way the world actually happens to be.

Kaplan (1968) overcomes Quine's difficulty by giving *Ortcutt* in (4) a *de re* interpretation. The core idea behind *de re* interpretation is to relate a believer to an individual through what has come to be known as an acquaintance relation, and to then make use of that acquaintance relation to spell out the actual belief held by the believer. Roughly this comes to the following:[3]

---

3 A more faithful rendition of Kaplan's analysis is:

*a believes P(b)* is true if $\exists\alpha[R(\alpha,b',a')\ \&\ a'\ \mathbf{B}\ \ulcorner P(\alpha)\urcorner]$,

where for all expressions x, x' is the interpretation of x, **B** is a primitive relation between an individual and an expression (roughly, the expression believed by that individual), α ranges over expressions, $\ulcorner P(\alpha)\urcorner$ denotes the expression P(α), and R(α,b',a') holds iff α is a vivid

(6) a believes P(b) if for some property Q

a is acquainted with b as *the Q* and

a believes P(the Q)

Applying this analysis to the case at hand gives the following result:

(7) ⟦Ralph believes Ortcutt is a spy⟧ = 1 if for some Q

Ralph is acquainted with Ortcutt as *the Q*, and

Ralph believes that *the Q* is a spy

If I take *Ralph believes that the Q is a spy* in the last line to be analyzed as in Hintikka's analysis,[4] for Q = *man in the alley*, both requirements in (7) are satisfied, giving us the correct prediction that *Ralph believes Ortcutt is a spy* is true.

While adding Kaplan's *de re* analysis to Hintikka's original analysis clearly extends the empirical coverage of the latter, it is not by itself problem free. Kaplan takes the acquaintance relation in this example to relate Ralph to Ortcutt via a name that constitutes a vivid name of Ortcutt for Ralph. However, while giving hints at what it means for a name to be vivid, he does not give either necessary or sufficient conditions

name of b' for a'. In the rendition of *de re* interpretation given in the text, *the Q* plays the role of α, a vivid name for Ralph that names Ortcutt. In principle Kaplan allows vivid names to be proper names in addition to definite descriptions, though we won't consider examples where appeal to a proper name is needed. If desired it is possible to approximate a proper name *N* with the description *the* $\lambda z.z{=}[\![N]\!]$. Presumably, such a description will count as a vivid name for an individual iff *N* does.

4 Kaplan himself does not commit to Hintikka's analysis or to any other particular analysis of attitude attributions. His main concern is to show how to get a Fregean approach to attitude statements to work in the face of Quine's objections to quantifying into opaque contexts. For this purpose he takes the verb *believe* to denote an unanalyzed predicate **B** that takes an individual and an expression as arguments. Connecting Kaplan's analysis to Hintikka's can be seen as an attempt to give an analysis of **B**.

for qualifying as such. Additionally, by restricting attention to vivid *names*, Kaplan's analysis fails to account for attitudes that could be analyzed as being toward properties, quantifiers, conjunctions or other entities of types other than type e, examples of which will be considered in detail below. Since such cases equally fall outside the range of what can be accounted for by Hintikka's original analysis, they remain unaccounted for under Kaplan's analysis taken as a supplementation to Hintikka's. Finally, by analyzing *believe* as denoting a primitive, unanalyzed relation between an individual and an expression, Kaplan's analysis provides no account for what beliefs can be plausibly taken to entail what other beliefs. Giving a Hintikkan analysis of the relation would provide such an account, but as we have already seen, the account that Hintikka provides fails in cases of attribution of beliefs in necessary truths and falsities.

Schwager (2010) notes empirical limitations of a *de re* analysis as well as problems for extending the notion of vividness from individual-denoting expressions to property-denoting expressions. One example highlighting these problems is the following (slightly modified):

(8) Situation:
A murder has occurred on campus. Detective Foyle investigates and concludes, 'someone with an office in Building Z is guilty.' Unbeknownst to him, all offices in Building Z belong to English department professors. Foyle has no *de dicto* beliefs about English department professors.

In the above situation, it is perfectly acceptable for a speaker apprised of the facts to say the following:

(9) *Foyle believes an English professor is guilty*. (true)
a. Foyle believes: 'An English professor is guilty' (unspecific (*de dicto*), false)

b. There is an English professor that Foyle believes is guilty (specific (*de re*), false)

c. ??? (true)

However, neither a *de dicto*/*unspecific* interpretation à la Hintikka nor a *de re*/*specific* interpretation of *an English professor* à la Kaplan results in the sentence being true. Under a Hintikkan *de dicto* interpretation, in order for the sentence to be true it has to be the case that the sum total of Foyle's beliefs entail that an English professor is guilty. However, Foyle lacks any *de dicto* beliefs about English department professors, including the key belief that would make such an entailment hold, namely that all of the people with an office in Building Z are English professors. The fact that people with an office in Building Z *are* all English professors is a contingent fact, not an analytic fact, and so does not necessarily extend to other possible worlds. This means there will be worlds compatible with the sum total of Foyle's beliefs in which a non-English department professor is guilty, for example in some of those worlds in which the relevant non-English department professor has an office in Building Z. This possibility renders the sentence false under Hintikka's analysis: the sum total of Foyle's beliefs do not entail that an English professor is guilty.

Kaplan's *de re* analysis fairs no better. Since that analysis only applies to individual-denoting expressions, adopting it here requires taking *an English professor* to be such an expression. (9) should then be true if the following obtains:

(10) For some Q,

Foyle is acquainted with an English professor as *the Q* and

Foyle believes that *the Q* is guilty (analyzed, e.g., as by Hintikka)

For this to be the case, Foyle has to have a particular English professor in mind: merely narrowing down the list of suspects to a bunch of English professors does not suffice. In the situation depicted, however, Foyle does not have any such individual in mind as the guilty one. The interpretation that Kaplan generates is a plausible one that will be true in other situations, but it does not account for the truth of (9) in the situation depicted in (8).

Can the truth of (9) be explained by extending Kaplan's *de re* analysis to property expressions? The most natural such extension would look like the following:

(11) a believes …P… if for some Q

a is acquainted with P as Q and

a believes …Q…

It is difficult to see how to make sense of this, though. What, in particular, does it mean to be acquainted with a property as another property or predicate? Taking our cue from Kaplan's *de re* proposal, we could say that it consists of two requirements: (i) that the two properties P and Q have identical extensions in the world of utterance, and (ii) that Q be a vivid predicate of the extension of P for the believer. In the situation depicted, neither of these requirements is unproblematic. The situation is not specific regarding whether the property of having an office in Building Z is co-extensive with the property of being an English professor, but if I add to the situation the existence of English professors who do not have an office in Building Z, then (i) will clearly not be satisfied. Crucially, such a change makes no difference to our intuitions, though: (9) remains true in (8) so extended. Regarding (ii), it isn't clear what it means for a predicate to be a vivid predicate of a property extension. Until this can be clarified, it remains unclear what this requirement comes to and hence how to determine whether the requirement is met. The relevant notion of vividness will have to be such that *English professor* can count as a vivid predicate of the extension of *person with an office in Building Z* in a situation in which English professors form a superset of the people with an office in Building Z. While this could be stipulated to hold, it is again difficult to see how applicability of the concept *vivid* can be reduced to necessary and sufficient conditions.

Based on these and other concerns, Schwager proposes to account for the true interpretation of (9) in the situation depicted in (8) via an independent *de qualitate* analysis of *English professor*. The core idea underlying her formal analysis is what she calls The Replacement Principle.

(12) **Schwager's Replacement Principle**: For the sake of reporting an attitude, a property that is involved in the content of the attitude that is to be reported (the reported property [*Q* above]) can be replaced by a different property (the reporting property [*P* above]) as long as the reported property is a subset of the reporting property at all relevant worlds.

According to this principle, (9) comes out as true since the set of people with an office in Building Z (the reported property) is a subset of the set of English professors (the reporting property) at all relevant worlds, licensing substitution of "English professor" for "person with an office in Building Z" in reporting Foyle's belief. The Replacement Principle is formalized as *de qualitate* interpretation as follows:

(13) ***de qualitate*** (from Schwager)

$Attitude_w(x,\langle P,Q\rangle)$ iff there is a property Q' s.t. at the w-closest worlds w' where $Q(w') \neq \emptyset$:

a. $Q'(w') \neq \emptyset$

b. $Q'(w') \subseteq Q(w')$

c. $Attitude_w(x,\lambda w''.P_{w''}(Q'))$ is true.

Let us step through this formalization. In the first line of the analysis, $\langle P,Q\rangle$ is a structured proposition, with P a function that applies to property Q to generate a proposition. This structured proposition is true iff P(Q) is true. Q and Q' are both properties of type ⟨s,⟨e,t⟩⟩, i.e. functions from worlds to sets of individuals. $\lambda w''$. $P_{w''}$(Q') in (c) is a simple, non-structured proposition, the set of possible worlds w'' in which $P_{w''}$(Q') is true, where $P_{w''}$ is the extension of P in world *w''*. While Schwager is not explicit about how to interpret the attitude attribution in (c), for concreteness I will again assume a Hintikkan interpretation.

To get (9) to come out true in the situation depicted in (8) under the analysis in (13), Q can be identified with the property of being an English professor. With this assumption, and taking the indefinite description to be interpreted as an existential quantifier, (9) can generate the following structured proposition and analysis:

(14) $\text{believes}_{@}$ (Foyle, ⟨λPλw.∃x(P(w)(x) & guilty(w)(x)), λwλx.English professor(x)(w)⟩)

This is true iff there is a property Q' such that at the @-closest worlds w' where *λx.English professor(x)(w')* is not empty,

a. Q'(w') is not empty

b. Q'(w') ⊆ λx.English professor(x)(w')

c. $\text{believes}_{@}$ (Foyle,λw''λP.∃x(P(w'')(x) & guilty(w'')(x))(Q')) is true.

For Q' = ⟦ person with an office in Building Z ⟧, (a)-(c) hold. (a) is true since worlds most like the world of utterance that contain an English professor also contain a person with an office in Building Z. This is so since the world of utterance @ contains a person with an office in Building Z and this fact is compatible with there being English professors. (b) is true because in the world of utterance the extension of *person with an office in Building Z* is a subset of the extension of *English professor* given that only English professors are people who have an office in Building Z. Since this fact is compatible with there being English professors, it also holds at all worlds closest to the world of utterance that contain English professors. Finally, (c) is true because *Foyle believes a person with an office in Building Z is guilty* is true according to Hintikka's analysis.

Empirically, Schwager's account, like Kaplan's before it, constitutes a clear improvement on prior analyses. However, also like Kaplan's analysis it is not problem-free. The first problem lies in the status of the analysis. It roughly has the form of *When sentence S is true, it's possible to say S'*. It is left unclear whether the reason that it's possible to say S' is because S' is semantically true whenever S is true or because it is pragmatically allowed to utter S' even when it's semantically false in just this situation. Additionally, in either case, it is unclear just what justifies the substitution of S' for S beyond the stipulation that it can be done. In this regard, the analysis can best be seen as a more formal and precise re-statement of the problem rather than as a formal solution to it. Second, the analysis as it stands makes wrong predictions for properties in downward entailing environments. In the situation described, for example, suppose that Foyle claims that a certain Professor Smith does not have an office in Building Z. Schwager's analysis incorrectly predicts that it should be ok to report on Foyle's beliefs by saying *Foyle believes that Professor Smith is not an English professor*, making the exact substitution that was seen to be acceptable in getting (9) to be acceptable. Finally, once again the analysis does nothing to address the problems posed by people with inconsistent beliefs and does not provide a plausible, principled analysis of belief in necessary truths and falsities. This is not a place where the analysis makes an incorrect prediction but rather where it makes no predictions at all, leaving unaddressed the question of what the proper underlying analysis should be.

To overcome the empirical problems Schwager's analysis faces with substitution in downward entailing environments it would be enough to make substitution sensitive to the environment in which it occurs. Formally this could be accomplished by changing line b in (13) to the following:

(15) b'. $Q'(w') \subseteq Q(w')$ and Q occurs in an upward entailing environment; or
$Q'(w') \supseteq Q(w')$ and Q occurs in a downward entailing environment; or
$Q'(w') = Q(w')$ and Q occurs in a non-monotonic environment[5]

While potentially formally adequate, however, such an extension amounts to little more than a listing of separate cases, without providing any deeper explanation for why it is these cases that are found. More importantly, this approach still leaves open the question of why the substitutions it licenses are acceptable, and it also fails to capture a clear parallelism that exists between *de re* and *de qualitate* examples. As they are spelled out here, both *de re* and *de qualitate* interpretations justify substitutions of one term for another in a belief attribution where the underlying attribution purportedly captures the relevant belief content directly but the result of substitution fails to do so. Furthermore, in both cases the person making the attribution has knowledge that allows her to infer the substitution content from the underlying belief content. Simply listing the kinds of substitutions that are possible provides no reason for why such a parallelism should be found.

Let me illustrate the parallelism in more detail. In the Ralph/Ortcutt example, the important piece of information that makes substitution possible is the speaker's knowledge that the man in the alley is Ortcutt. Combined with the content of Ralph's belief that the man in the alley is a spy, this piece of information makes it possible for the speaker to infer that Ortcutt is a spy. In the Foyle's investigation example, the important piece of information is that all people with an office in Building Z are

5 The third line has not been argued for, but would be required to account for attributions involving non-monotonic operators such as "exactly n". In a situation like (I) in which Foyle utters "Exactly one person with an office in Building Z is guilty", the attribution "Foyle believes that exactly one English professor is guilty" will be true if and only if all and only English professors have an office in Building Z.

English professors. Combined with the content of Foyle's belief that a person with an office in Building Z is guilty, this piece of information makes it possible for the speaker to infer that an English professor is guilty. Rather than treat these cases as formally distinct, I propose to unify them by incorporating inference directly into the semantics of belief attribution. To get the analysis to work, I will also propose an alternative foundation for attitude attribution as a replacement for the possible worlds foundation of Hintikka.

## 2 Analysis

The formal analysis that I propose is the following:

(16) $[\![$ *a believes that S* $]\!]^{j} = \lambda w. [\![ a ]\!]$ has a token (underlying) belief *B* in w from which $[\![ S ]\!]^{w}$ can be inferred by (judge) j.

Previewing, this analysis will be seen to have the following advantages over the analyses of Hintikka, Kaplan and Schwager:

I: It accounts uniformly for all of the intuitions underlying the *de dicto*, *de re* and *de qualitate* interpretations examined, stipulating only the underlying semantics of attitude predicates like *believe*.

II: It accounts for substitution behavior in downward entailing environments.

III: It extends to substitutions of expressions of other types without over-generating.

IV: It provides a properly constrained analysis of beliefs in necessary truths and falsities that allows one to draw reasonable conclusions from necessary or impossible beliefs without generating absurd consequences.

The process of *inference* employed in (16) is a process of deducing a conclusion from premises. I tentatively propose that all inference can be reduced to the following form:

(17) $a \in P$

$P \subseteq Q$

$\therefore\ a \in Q$

This notion of inference differs crucially from the notion of entailment modeled in possible worlds semantics. While entailment is analyzed as a relation among propositions analyzed as sets of worlds, inference is a world-internal relation. Since the beliefs of the inferrer end up playing a vital role in accounting for the acceptability of the *de re* and *de qualitate* examples I have looked at so far, the world within which inference applies has to be one where the inferrer's beliefs are true.[6] This relativity to the inferrer's beliefs is introduced through the superscripted judge parameter *j* in (16), formally taken to introduce a judge together with the set of propositions the judge holds to be true at the world of evaluation. For matrix utterances, the judge is identified with the individual evaluating the sentence, the Agent of the context (speaker) for determining whether it can be uttered, and the Addressee of the context (hearer) for deciding whether the proposition uttered can be accepted into the common ground.[7]

The inference process outlined above is based on the beliefs of the judge. This makes it possible to draw inferences that are intuitively but not logically valid. For example, given a judge's belief that all basketball players are tall, from the premise that John is a basketball player she can infer that John is tall. The conclusion clearly does not follow logically from the premise alone, but is intuitively felt to be valid nonetheless.

6 Below we will see that this holds only for non-*de translato* interpretation.

7 In Tancredi and Sharvit (forthcoming) we argue that the perspective parameter plays a role in *de translato* interpretation. However, the parameter value for generating *de translato* interpretations in a simple belief attribution statement is that of the embedded clause, identified as the perspective of the subject of the matrix verb, not that of the attributor as here.

The belief that all basketball players are tall plays a crucial role in the inference. Generalizing, in implementing the inference process above, $P \subseteq Q$ may not hold in all possible worlds, but if it holds according to the beliefs of the judge, then it is a valid premise.[8]

The general form posited for inference assumes no restriction on the types of expressions that can enter into such inference relations except for the restrictions implicit in the logical symbols "$\in$" and "$\subseteq$". Formally, if a is of type σ, then P and Q must be of type ⟨σ,t⟩, but no other restrictions are imposed. This makes it possible for inferences to be drawn based on a relation between properties, quantifier meanings, determiner meanings or any other non-propositional meanings of conjoinable type. We will see below that this is a welcome consequence of the proposal.

As with structured meaning accounts of attitudes, I assume that propositions are structured entities. In particular, I take all propositions to have the structure imposed by compositional semantics operating on syntactic trees prior to lambda conversion. I also take the structure of a proposition to affect the attitudes one can hold toward it. Where our account differs from standard structured proposition accounts is in the role that inference plays in connecting beliefs and their attribution. In this regard, the most important role that structuring plays in the proposed analysis is as input to an abstraction operation that creates a function-argument pair. Below, this shows up implicitly at the beginning of the inference process where a proposition *p* is split into

8 This aspect of the inference process can be mimicked under a possible worlds semantics as entailment under judge's belief. This kind of example thus does not provide a reason for not adopting a possible worlds analysis of belief attribution. Here, though, we are only concerned with laying out and illustrating the proposal. Motivation for the proposal will ultimately come from its ability to handle the well-known problems facing Hintikka and also in its ability to predict connections among attributable beliefs that do not follow from any known extensions of Hintikka's analysis that incorporate either structured propositions or impossible worlds.

two pieces, a functional piece $P$ and its argument $a$. Properly formulated, this abstraction process will guarantee that $P(a) = p$.

Formally, I take a belief to be a pairing of a structure with a model theoretic interpretation thereof. The structure I take to be homomorphic with the compositional semantic structure of a sentence prior to lambda conversion. In cases where the difference is immaterial, I often pretend that the structured part of the belief simply *is* a linguistic structure, though the need to account for our practice of attributing beliefs to pets without our thereby presupposing them to have a human linguistic ability requires that the two be distinguished in principle. This division allows us to make a distinction between two types of inference: inference based on a full belief – structure plus presumed model theoretic interpretation of the believer – and inference based only on the structured part of a belief, assigning values according to our own model theoretic interpretation. Until the role of a model becomes important, however, I will ignore it below, equating a belief with the structured portion of a belief.

We are now ready to see how the analysis proposed applies to the cases considered so far. I start with the *de dicto* case from (2), where I spell out all of the necessary details of the analysis, including the full compositional structure prior to lambda conversion of the proposition being attributed, the abstraction over this structure, and its simplified, partly lambda-converted structure, implicitly taken to generate sets rather than functions in order to match the form proposed earlier as the universal form for all inference.

(18) *de dicto*

John: "The earth is flat"

Me: "John believes the earth is flat"

John's utterance is evidence that he has a belief $B$ whose content is that the earth is flat.

My utterance is true iff there is a belief *b* of John's from which I can infer the earth is flat.

For *b* = *B*, I can trivially infer that the earth is flat.

*B*: The earth is flat

[[λPλQ.[the x: P(x)](Q(x)] (λy. earth(y))] (λz. z is flat)

(compositional structure)

[[ **the'** ] ( **earth'** )] ( **is-flat'** )

λz. z is flat ∈ λP'.[ [λPλQ.[the x: P(x)](Q(x)] (λy. earth(y))](P')

(abstraction)

λz. z is flat ∈ λP'.[the x: earth(x)](P'(x)) (lambda conversion)

λP. [the x: earth(x)](P(x)) ⊆ λP. [the x: earth(x)](P(x)) (trivially)

∴λx. x is flat ∈ λP. [the x: earth(x)](P(x)) (i.e. the earth is flat)

∴ My utterance is true.

This is a case where the inference process is formally appealed to but where it applies trivially. The inference has the form: *P(a);P⊆P; ∴P(a)*. While trivial, the preservation of the form of inference is important. Of particular importance is that the inference not be of the form: *p; ∴p*, a fact whose significance will become clear when we look at mathematical inferences below.

Next consider the case of Ralph and Ortcutt adapted from Quine. This is an example that was given a *de re* interpretation by Kaplan, formally introducing ambiguity into attitude interpretation to set the proposed interpretation apart from *de dicto* interpretation. Here we see that introducing such a *de re* / *de dicto* ambiguity is unneeded. (Here and below I simplify and abbreviate the process of dividing a structured proposition into its functional and argument pieces.)

(19) Ralph: "The man in the alley is a spy. Ortcutt is not a spy."

Me: "Ralph believes that Ortcutt is a spy."

John's initial utterance is evidence that he has a belief *B* whose content is that the man in the alley is a spy.

My utterance is true iff there is a belief *b* of John's from which I can infer that Ortcutt is a spy.

For *b* = *B* I can do so:

B: The man in the alley is a spy

[the x: man in the alley(x)](is a spy(x)) (definite description as quantifier)

$\lambda$x. x is a spy $\in$ $\lambda$P. [the x: man in the alley(x)](P(x))

$\lambda$P. [the x: man in the alley(x)](P(x)) $\subseteq$ $\lambda$P. P(Ortcutt) (my knowledge)

$\therefore$ $\lambda$x. x is a spy $\in$ $\lambda$P. P(Ortcutt) (i.e., Ortcutt is a spy)

$\therefore$ My utterance is true.

Here the important piece of information that allows the inference to go through is the information that Ortcutt is the man in the alley. If this is the case, then it follows that whatever properties hold of the man in the alley hold of Ortcutt. This is not a logical necessity, of course, but it is an relation that will hold given the judge's belief that the man in the alley is Ortcutt, which is all that is required.

Next consider the cases of Foyle's investigation, which Schwager used to motivate introduction of *de qualitate* as an additional mode of interpretation. We see in the first example below that our analysis makes such an additional mode of interpretation otiose, and in the second and third examples that our proposal improves empirically on Schwager's *de qualitate* analysis in its handling of property expressions occurring in downward entailing environments.

(20) (Upward entailing example)

Foyle: "A person with an office in Building Z is guilty."

Me: "Foyle thinks an English professor is guilty."

Foyle's utterance is evidence that he has a belief *B* whose content is that a person with an office in Building Z is guilty.

My utterance is true iff there is a belief *b* of Foyle's from which I can infer that an English professor is guilty.

For $b = B$ I can do so:

*B*: A person with an office in Building Z is guilty

$\exists x(\text{Building-Z-office-person}(x)\ \&\ \text{guilty}(x))$

$\lambda x.\ \text{guilty}(x) \in \lambda P.\ \exists x\ (\text{Building-Z-office-person}(x)\ \&\ P(x))$

$\lambda P.\ \exists x\ (\text{Building-Z-office-person}(x)\ \&\ P(x))$

$\subseteq \lambda P.\ \exists x\ (\text{English professor}\ (x)\ \&\ P(x))$

$\therefore \lambda x.\ \text{guilty}(x) \in \lambda P.\ \exists x\ (\text{English professor}(x)\ \&\ P(x))$

(i.e. an English professor is guilty)

$\therefore$ My utterance is true.

(21) (Downward entailing, good example)

Foyle: "No English professor is guilty."

Me: "Foyle thinks that no one with an office in Building Z is guilty."

Foyle's utterance is evidence that he has a belief *B* whose content is that no English professor is guilty.

My utterance is true iff there is a belief *b* of Foyle's from which I can infer that no one with an office in Building Z is guilty.

For $b = B$ I can do so:

*B*: No English professor is guilty.

$\neg\exists x(\text{English professor}(x)\ \&\ \text{guilty}(x))$

$\lambda x.\ \text{guilty}(x) \in \lambda P.\ \neg\exists x\ (\text{English professor}(x)\ \&\ P(x))$

$\lambda P.\neg\exists x(\text{English professor}(x)\ \&\ P(x))$

$\subseteq \lambda P.\neg\exists x(\text{Building-Z-office-person}(x)\ \&\ P(x))$

$\therefore \lambda x.\ \text{guilty}(x) \in \lambda P.\neg\exists x\ (\text{Building-Z-office-person}(x)\ \&\ P(x))$

$\therefore$ My utterance is true

Note that Schwager's analysis does not as it stands account for the acceptability of this example. Her analysis only allows substitution of property-denoting expressions, but for her to adopt the explanation given for the acceptability of (21) would require allowing substitution of quantifier-denoting expressions. It is trivial to extend her analysis to cover such cases, however. All that is needed is to remove the requirement that Q be a property in the first line of (13) and to allow the interpretation of *guilty* to be type-raised. This modification to her analysis will not help with the following example, though.

(22) (Downward entailing bad example)

Foyle: "No one with an office in Building Z is guilty."

Me: #"Foyle thinks no English professor is guilty."

The only relevant belief Foyle has indicated having is a belief *B* whose content is that no person with an office in Building Z is guilty.

My utterance is true iff there is a belief *b* of Foyle's from which I can infer that no English professor is guilty.

I cannot do so for $b = B$, and I have no basis for assuming Foyle has some other belief $B'$ from which I can do so.

[One unsuccessful attempt:]

*B*: No one with an office in Building Z is guilty.

$\neg\exists x(\text{Building-Z-office-person}(x) \ \&\ \text{guilty}(x))$

$\lambda x.\ \text{guilty}(x) \in \lambda P.\ \neg\exists x\ (\text{Building-Z-office-person}(x) \ \&\ P(x))$

$\lambda P.\neg\exists x(\text{Building-Z-office-person}(x) \ \&\ P(x))$

$\not\subseteq \lambda P.\neg\exists x(\text{English professor}\ (x) \ \&\ P(x))$

[Fail]

[All other attempts fail]

**Fail**

∴ My utterance is not true.

The falsity of the speaker's utterance in this last example is of particular importance since this example sets the proposed analysis apart empirically not only from Schwager's analysis as implemented but also from any plausible extension of the analysis that allows substitution of quantifiers. Since λx. Building-Z-office-person(x) ⊆ λx.English professor(x), Schwager wrongly predicts that "English professor" should be substitutable for "person with an office in Building Z" based on the structured proposition: ⟨*λP.¬∃x(P(x) & guilty(x))*, *λx.English-professor(x)*⟩. Her analysis therefore wrongly predicts the speaker's utterance to be true in this example, contrary to observation. The same structuring of the proposition fails to generate a true interpretation under our proposal since it fails to give rise to a valid inference of the form in (17). This inference pattern requires a subset relation to hold between the expression functioning as a function in the underlying belief and some other related function in the attributed proposition. Splitting the proposition in the above way, however, results in a functional part that is identical for the underlying belief and the attributed proposition.

In addition to accounting for the *de dicto*, *de re* and *de qualitate* examples from the literature, the proposed analysis also directly accounts for examples not previously looked at in the literature to the best of my knowledge. Consider, for example, a variant on the Foyle's investigation examples. Add to the original situation the speaker's knowledge that two thirds of the English professors have an office in Building Z. In this revised situation, if Foyle asserts "Everyone with an office in Building Z is suspect," the speaker can report on his beliefs by claiming "Foyle believes that most English professors are suspect." The process of showing this sentence to be true follows the exact same pattern as the other cases.

(23) Foyle: "Everyone with an office in Building Z is suspect."

Me: "Foyle thinks most English professors are suspect."

Foyle's utterance is evidence that he has a belief *B* whose content is that every person with an office in Building Z is suspect.

My utterance is true iff there is a belief *b* of Foyle's from which I can infer that most English professors are suspect.

From *B* I can do so:

*B*: Everyone with an office in Building Z is suspect.

[∀x: Building-Z-office-person(x)] (suspect(x))

$\lambda$x.suspect(x) $\in$ $\lambda$P. [∀x: Building-Z-office-person(x)] (P(x))

$\lambda$P.[∀x: Building-Z-office-person(x)](P(x))

$\subseteq$ $\lambda$P.[Most x: English professor (x)](P(x))

$\therefore$ $\lambda$x.suspect(x) $\in$ $\lambda$P. [Most x: English professor(x)] (P(x))

$\therefore$ My utterance is true.

It should be noted that Schwager can account for these kinds of cases as well by making the same change needed to account for (21), i.e. by first allowing substitution of quantifiers in addition to properties and then type-raising *λx.suspect(x)* to *λQ.Q(λx.suspect(x))* so that the quantifier phrase ends up as the second, i.e. argument, expression in the structured proposition.

## 3 Facing the problems facing other analyses.

Having set forth our proposal and demonstrated how it accounts for all of the judgments that have been considered so far, it is now time to show how the analysis fares vis-à-vis the problems that plagued earlier analyses. I start with Hintikka's analysis.

### 3.1 Hintikka

The first major problem we saw with Hintikka's analysis was that it only works for attributing beliefs to people whose beliefs are taken to be consistent. This problem arises because belief statements for Hintikka relate an individual to a proposition via the set of worlds compatible with that individual's beliefs. The analysis proposed here instead relates a believer to a proposition via a single token belief of the believer and the full set of token beliefs of the judge. As long as the token belief of the believer is not itself internally inconsistent there are no problems with the believer having inconsistent beliefs elsewhere, since the inference does not depend on these other beliefs. As for the beliefs of the judge, there is no need under the proposed analysis to analyze these beliefs as mere sets of worlds, and there is every reason to take them to be compositionally structured. Indeed, the possibility of their playing the roles given them in inference depends on their having such an internal structure, since this structure plays a vital role in generating the subset relation among sets that makes inference possible. Since the analysis only requires manipulation of individual beliefs of the judge rather than the sum total of all their beliefs, whether the judge's beliefs are consistent or not is again irrelevant.

The second major problem we saw for Hintikka's analysis was that it makes absurd predictions about beliefs in necessary truths or falsities: everyone and everything is predicted to believe every necessary truth, and anyone who believes a necessary falsity is predicted to believe every proposition. This problem arose since belief statements were analyzed in terms of entailment (*a believes p* iff the sum total of a's beliefs entail *p*), itself analyzed as a property of sets of possible worlds (*p* entails *q* iff the set of worlds in which *p* is true is a subset of the set of worlds in which *q* is true). In the analysis proposed here, entailment does not directly enter into the semantics of belief attribution at all. It is replaced by inference, which is a process applying within an individual world. Within a world, it is very easy to find beliefs that do not allow one to infer, say, that 2+2=4 using only the form of inference given above. The belief that Mary laughed, for example, fails to give rise to such an inference. For inanimate objects that lack internal beliefs, again our analysis gives the right results. Absence of belief makes it impossible for the inference relation to get off the ground. The required inference starts from a token belief, and token beliefs are something that inanimate objects lack.

Regarding beliefs in impossibilities, the analysis proposed here makes it possible to distinguish acceptable attributions of impossible beliefs from unacceptable ones, something that no other analysis I know of is capable of doing. To see this, suppose that John utters the sentence "2+2=1+4". I can report this by saying (24a), but I cannot report it by saying (24b).

(24) John: "2+2=1+4"

a. Me: "John believes that 2+2=5"

b. Me: #"John believes that $i^i$ is an imaginary number"

Under Hintikka's analysis, (24a) is true iff (24b) is true. This is because Hintikka has no way to distinguish among impossible propositions. All impossible propositions have the same denotation, namely the empty set, and so a person can only be truthfully said to believe one if they have inconsistent beliefs, but then that person can be said to believe anything, including any other impossible proposition. The proposed analysis, however, accounts for these facts easily, since it is straightforward to infer that 2+2=5 starting from the premise that 2+2=1+4, but it is far from straightforward to infer that $i^i$ is an imaginary number starting from that same premise. Indeed, the reason your typical person cannot report John's beliefs using (24b) is that they have no idea how to infer that $i^i$ is an imaginary number from the premise that 2+2=1+4, or even whether such an inference is possible.

A consequence of the proposed analysis is that for a speaker who does know how to construct the required inference, (24b) should be an acceptable way to report on John's beliefs. I believe that this is a desirable result, though finding a natural situation in which such an attribution comes out as unobjectionable is admittedly difficult. Your typical hearer would, of course, reject (24b) without receiving proper justification. However, if you were to observe John's utterance that 2+2=1+4 and then proceeded to spell out the inference steps needed to go from *2+2=1+4* to $i^i$ *is imaginary*, a hearer who accepted your inference as valid would then also have to accept (24b) as true.[9]

In addition to the above empirical problems, we also saw that Hintikka's *de dicto* analysis requires supplementation to account for *de re* and *de qualitate* attributions, making attitude attribution statements systematically ambiguous. We observed that it is possible to tweak Schwager's analysis of *de qualitate* attributions to overcome the problems posed by property expressions occurring in downward entailing environments. Making such a change, however, would be to maintain an ambiguity analysis of attitude statements. While this does not constitute a strong argument against the approach, I propose that all else being equal, a non-ambiguity analysis is preferable to an ambiguity analysis. This consideration favors the analysis proposed here, since it accounts for (intuitive) *de dicto*, *de re* and *de qualitate* attributions without positing an ambiguity.[10] The differences among these flavors of attribution on the proposed analysis derive not from ambiguity but merely from differences in how inferences are drawn.

### 3.2 Kaplan

As mentioned earlier, Kaplan does not give a formal analysis of what it means for an individual to stand in the belief relation to an expression, leaving the belief relation, his **B**, an unanalyzed primitive. Since I chose to assume a Hintikkan analysis of **B**, I saddled Kaplan with all of the problems facing Hintikka. This is certainly unfair to

9 The attribution becomes much more natural if prefaced by *So essentially*.

10 This is not strictly speaking correct. As we will see, getting the analysis to account for certain denials of belief requires taking belief attributions to be ambiguous between *de translato* and non-*de translato* interpretations. However, this ambiguity was argued in Tancredi and Sharvit (forthcoming) to be required independently, even under an analysis that admits a *de dicto* / *de re* / *de qualitate* ambiguity. Consequently, admitting *a de translato* / non-*de translato* ambiguity does not argue in favor of an analysis that admits a *de dicto* / *de re* / *de qualitate* ambiguity as well.

Kaplan himself, but is justified by the fact that the main development of attitude semantics that employs Kaplan's *de re* analysis does so by incorporating it into a Hintikka-type analysis. However, in addition to the problems coming from the adoption of Hintikka's framework for analyzing attitude statements, Kaplan's *de re* analysis introduced difficulties and faced limitations of its own.

The first problem noted for Kaplan's analysis came from the vividness requirement imposed on *de re* interpretation. Built into the acquaintance relation, this is a requirement that the attitude holder have a vivid name for the individual picked out by the expression being given a *de re* analysis. The problem was that neither necessary nor sufficient conditions were given for determining when a name satisfies this vividness requirement, making it impossible to either support or falsify the analysis. Additionally, it was difficult to see how a *de re* analysis could be extended to account for the kinds of example that Schwager examines. Doing so would require extending the notion of vividness in such a way that, for example, "English professor" ends up as a vivid predicate of the set of Building Z office holders even when the set of Building Z office holders is a proper subset of the set of English professors. Under the analysis proposed here, vividness plays no role, and so avoids the need to formally analyze the notion.

One of Kaplan's concerns, and indeed the primary reason that he introduced the concept of a vivid name, was to account for the differing beliefs that can be reported using the sentence below:

(25) John believes that the shortest spy is a spy

Under one way of understanding the sentence, it is taken to attribute to John something that could be attributed to anybody who believes that spies exist and that among all spies one is the shortest. Under another way of understanding the sentence, it requires

John to have beliefs about a particular spy. The analysis proposed allows us to reduce this difference in understanding to a difference in scope for the definite description. On a narrow scope reading, the sentence requires John to have a belief from which it can be inferred that the shortest spy is a spy. His belief that spies exist and that one among them is the shortest makes such an inference possible. On a wide scope reading, it requires that the shortest spy be an x such that John has a belief from which it can be inferred that x is a spy. This requires John to have a belief about the individual who the judge takes to actually be the shortest spy. The analysis thus predicts the very distinction that Kaplan was centrally concerned to account for.

The second problem, or rather shortcoming, of Kaplan's analysis is that it makes no predictions about what beliefs should follow from what other beliefs. This came from the fact that the verb *believe* was analyzed as a primitive relation **B** between individuals and expressions. The fact that an individual *a* bears **B** to an expression *S* says nothing about whether *a* also bears **B** to any other expression *S′* regardless of the relation that holds between *S* and *S′*. No inferences from *a believes that S* to *a believes that S′* are thus licensed under Kaplan's analysis. As can be seen in the response to the problems facing Hintikka, the analysis proposed here does not share this shortcoming. All belief attribution involves inference on that proposal, with inference restricted to the universal inference relation proposed in (17). Whenever *S′* can be inferred from *S*, our analysis predicts that *a believes that S′* will be entailed by *a believes that S*.

### 3.3 Schwager

From one perspective, the analysis proposed here could be seen as a generalization of Schwager's *de qualitate* analysis. However, there are some important differences between the two. As mentioned earlier, Schwager's analysis is essentially a formalization of her replacement principle, but the status of that principle was unclear. It allowed one expression to be substituted for another in attitude attribution contexts, but it gave no justification for why such a substitution was acceptable. Under our analysis the same facts are accounted for without positing any substitution. Rather, the analysis given of the examples directly predicts their status based on their semantic interpretation. The reason why what Schwager labeled *de qualitate* attributions are acceptable in certain circumstances, then, is under the proposed analysis because they are semantically true in those circumstances.

We also saw that Schwager's analysis makes some incorrect predictions about substitution in downward entailing environments, stemming from the fact that her analysis looks only at the expressions being substituted for one another and not at the environment in which the substitution occurs. In particular, it allows substitution of a set-denoting property for a superset-denoting property regardless of environment, but downward entailing environments only intuitively allow the reverse substitution. Our analysis gets these examples to come out right because it is based on inference relations, not on substitutions, and what counts as an acceptable inference relation depends not only on the denotation of the relevant argument term but also on the monotonicity properties of the functional terms employed in the inference, encoded in the requirement that $F \subseteq G$ in the proposed universal inference process in (17).

Finally, Schwager's analysis reduces a *de qualitate* attribution to a *de dicto* interpretation, but does not specify how *de dicto* attitudes are interpreted. This makes her analysis incomplete. Taking her analysis to be a revision to a version of the standard view that adopts Hintikka's analysis for providing the needed *de dicto* interpretation brings with it the problems examined above facing that analysis. Not doing so, however, leaves it unclear what the basis of her analysis is and hence makes it impossible to test. Our proposal offers a concrete alternative to Hintikka's analysis that makes clear, testable predictions, and so improves on Schwager's analysis in this respect.

## 4 Limitations of Inference?

Above I proposed a very minimalist analysis of inference. All inference, I hypothesized, can be reduced to the following form applying within a given world:

(26) $a \in P$

$P \subseteq Q$

$\therefore a \in Q$

There are obvious cases of inference that can easily play the role needed for getting belief attributions to come out right, however, that do not on the surface look like instances of this pattern. In this section I spell out how to reduce a small number of these inferences to this pattern. I make no attempt to offer a proof of the hypothesis, leaving that problem for future work.

One common pattern of inference that falls under the category of non-obvious instances of the hypothesized pattern is universal instantiation:

(27) $[\forall x{:}P(x)](Q(x))$

$P(a)$

$\therefore Q(a)$

That universal instantiation can be used as a basis for a belief attribution can be seen in the following example, where the parenthesized sentence indicates speaker's knowledge.

(28) John: "Everyone in my class passed."

Me: (Mary is in John's class.) "John believes that Mary passed."

I give the bare bones of reducing this case of universal instantiation to the hypothesized inference pattern below:

(29) Everyone in John's class passed.

$[\forall x$: x is in John's class] (passed(x))

$\lambda x$. x passed $\in$ $\lambda P$. [$\forall x$: x is in John's class] (P(x))

$\lambda P$. [$\forall x$: x is in John's class] (P(x)) $\subseteq$ $\lambda P$. P(mary)

$\therefore$ $\lambda x$. x passed $\in$ $\lambda P$. P(mary) (i.e. Mary passed)

The inference goes through given the assumption that Mary is in John's class. This assumption licenses the crucial step in the second to last line: If Mary is in John's class, then the properties that hold of all members of John's class also hold of Mary and are hence a subset of the set of properties that hold of Mary.

Another common inference step is existential generalization:

(30) P(a)

$\therefore \exists x$ (P(x))

Again this is a step that appears to be usable in justifying a belief attribution, as seen in the following example:

(31) John: "London burned."

Me: (London exists.) "John believes that something burned."

The relevant inference following the hypothesized pattern is the following:

(32) London burned

$\lambda x$. x burned $\in$ $\lambda P$. P(London)

$\lambda P$. P(London) $\subseteq$ $\lambda P$. $\exists x$ (P(x))

$\therefore$ $\lambda x$. x burned $\in$ $\lambda P$. $\exists x$ (P(x)) (i.e. Something burned)

The crucial second to last line is the simple observation that all properties that are true of London are true of something, a formula true given the speaker's knowledge that London exists. The inference then goes through for any speaker who accepts that London exists.

I can also look at less common but still intuitively valid inferences that support belief attributions. Consider, for example, the following:

(33) John: "All but one person in my class passed."
Me: (Mary and Sue are in John's class.) "John believes that Mary or Sue passed."

This example is formally more complex than the others we have considered so far, but it too can be beaten into the hypothesized universal form. The structure of the relevant inference goes as follows:

(34) All but one person in John's class passed.
[All but one x: x is a person in John's class] (passed(x))
$\lambda x$. x passed $\in$ $\lambda P$. [All but one x: x is a person in John's class] (P(x))
$\lambda P$. [All but one x: x is a person in John's class] (P(x)) $\subseteq$ $\lambda P$.(P(mary) or P(sue))
$\lambda x$. x passed $\in$ $\lambda P$.(P(mary) or P(sue))

The truth of the second to last line can be seen intuitively: If all but one person in John's class has some property P, and if Mary and Sue are two people in John's class, then no more than one of them can fail to have P. That is, either Mary or Sue (or possibly both) will have P. Since this holds regardless of the value of P, the set of properties that hold of all but one person in John's class is a subset of the set of properties that are true of either Sue or Mary (or both). A detailed formal justification for the second to last line can also be given. I make no attempt to shoehorn every step of inference into the hypothesized universal pattern in giving this justification, leaving open the question of whether it is possible to do so. Obviously the hypothesized universal form of inference will be more strongly supported if it applies to every step in the inference, though it's possible to countenance a weaker form of the hypothesis that requires only that the main line of inference be of the hypothesized universal form. (Below, "A\B" denotes the intersection of set A with the complement of set B, and |A| denotes the cardinality of set A.)

(35) $\forall P,Q\ (\langle P,Q\rangle \in \{\langle P',Q'\rangle\text{: [all but one x: } P'(x)](Q'(x))\} \leftrightarrow |P\setminus Q| = 1)$

[semantics of all but one]

$\{\text{mary,sue}\} \subseteq \{x\text{: x is a person in John's class}\}$ [my knowledge]

$\forall P,P',Q\ (P \subseteq P' \rightarrow P\setminus Q \subseteq P'\setminus Q)$ [logic]

$\therefore\ \{\text{mary,sue}\}\setminus\{x\text{: x passed}\} \subseteq \{x\text{: x is a person in John's class}\}\setminus\{x\text{: x passed}\}$

[universal instantiation]

$\forall P,Q\ (P \subseteq Q \rightarrow |P| \leq |Q|)$ [logic]

$\{\text{mary,sue}\}\setminus\{x\text{: x passed}\} \subseteq \{x\text{: x is a person in John's class}\}\setminus\{x\text{: x passed}\}$

[above]

$|\{x\text{: x is a person in John's class}\} \setminus \{x\text{: x passed}\}| = 1$ [semantics of all but 1]

$\therefore\ |\{\text{mary,sue}\} \setminus \{x\text{: x passed}\}| \leq 1$

$\forall P,Q\ ((P \cap Q) \cup (P\setminus Q) = P)$

$(\{\text{mary,sue}\} \cap \{x\text{: x passed}\}) \cup (\{\text{mary,sue}\} \setminus \{x\text{: x passed}\}) = \{\text{mary,sue}\}$

$|(\{\text{mary,sue}\} \cap \{x\text{: x passed}\}) \cup (\{\text{mary,sue}\} \setminus \{x\text{: x passed}\})| = |\{\text{mary,sue}\}| = 2$

$|(\{\text{mary,sue}\} \cap \{x\text{: x passed}\}) \cup (\{\text{mary,sue}\} \setminus \{x\text{: x passed}\})|$

$= |(\{\text{mary,sue}\} \cap \{x\text{: x passed}\})| + |(\{\text{mary,sue}\} \setminus \{x\text{: x passed}\})|$

$|(\{\text{mary,sue}\} \setminus \{x\text{: x passed}\})| \in \{0,1\}$

$\therefore\ |(\{\text{mary,sue}\} \cap \{x\text{: x passed}\})| \in \{1,2\}$

$\therefore\ \{x\text{: x passed}\} \in \{P\text{: } |P \cap \{\text{mary,sue}\}| \geq 1\}$

$$\{P: |P \cap \{mary,sue\}| \geq 1\} = [\![\text{Mary or Sue}]\!] = \lambda P.(P(mary) \text{ or } P(sue))$$

$$\therefore \{x: x \text{ passed}\} \in \lambda P.(P(mary) \text{ or } P(sue)) \qquad \text{(i.e. Mary or Sue passed)}$$

Below I give other examples of possible belief attributions based on non-obvious examples of the hypothesized inference pattern. I leave it to the reader to fill in the relevant inferences.

(36) John: "Exactly one person came."
Me: (Mary and Sue are people.) "John believes that if Mary came Sue didn't and vice versa."

(37) John: "No one in my class failed."
Me: (Mary is in John's class.) "John believes Mary didn't fail."

Note that the inferences do not depend on the speaker accepting the attitude holder's belief as true. They simply spell out a conclusion that follows if that belief is true.

## 5 Inference-Based Attribution and Hyperintensionality

As mentioned above, the inability of Hintikka's analysis to account for attributions of beliefs in necessary truths or necessary falsities has been addressed along two general lines in the literature: (i) taking propositions to be structured and belief in a proposition to depend on the proposition's structure; and (ii) expanding the domain of worlds to include not only possible worlds but also impossible worlds. These approaches come under the heading of hyperintensionality since they require reference to more than just the intensional interpretation of the clausal argument of an attitude predicate in order to determine whether a belief attribution is true. Both approaches implicitly or explicitly assume that the complement clause in a belief attribution identifies the purported belief of the subject. I argue here that our inference-based proposal is superior to both of these approaches.

While the analysis I have proposed shares with (i) the idea that the propositions that serve as the contents of belief are structured, it differs from (i) in how these propositions are identified. Under standard structured proposition accounts, the proposition denoted by the complement clause identifies this proposition directly. Under our proposal, the proposition denoted by the complement clause only identifies something that follows from some underlying purported belief of the believer. The significance of incorporating an inference process into the analysis can be seen in the following examples:

(38) John: "Mary saw Sue"
 Me: "John believes that Sue was seen by Mary"

(39) John: "Mary sold Sue a car"
 Me: "John believes that Sue bought a car from Mary"

In (38), John's utterance is in the active voice, and implies that John has a belief whose content is that Mary saw Sue. If the structure of a proposition affected whether the proposition is believed, the only way the speaker above could switch from the active voice in John's utterance to the passive voice in the belief attribution is if the propositions derived compositionally from active and passive share the same structure. King (2018) makes just such a proposal, arguing that composition applies in essence to the arguments in their base generated positions, not in their derived surface positions. Whatever the merits of such an analysis of passives, however, it does nothing to help account for (39). In (39) the two relevant sentences describe a single real-world situation in two different but equivalent ways, but it is not plausible to analyze the propositions corresponding to the two sentences as identically structured.[11]

11 Similar problems arise with pairs of sentences only one of which entails the other. For p = "Mary loaded hay onto the truck" and q = "Mary loaded the truck with hay", q entails p

In order to account for the acceptability of (39) under an analysis of belief attribution which takes the structure of a proposition to affect whether it is believed, it is necessary to be able to connect the pairs of propositions in a way that won't result in a parallel connection holding between all pairs of necessary propositions. Standard structured proposition analyses have nothing that accomplishes this required task.

Under the analysis proposed here, these examples are straightforwardly predicted to be acceptable. In (38), John's utterance can be taken to indicate that he has a belief with the content that Mary saw Sue. We can analyze this content as having the structure [**saw**(**sue**)](**mary**) while analyzing the embedded clause in the speaker's utterance as having the structure [[**was-seen**](**by**(**mary**))](**sue**) and still account for the acceptability of the attribution. The relevant steps in the inference from the former to the latter are the following:

(40) [[λxλy.saw(y,x)](sue)](mary)
[λy.saw(y,sue)] (mary) (simplifying)
mary ∈ λy.saw(y,sue) (abstracting over mary)
λy.saw(y,sue) ⊆ λy.was-seen(sue,by(y))
∴ mary ∈ λy.was-seen(sue,by(y))

The validity of the second to last line comes from the knowledge that the set of things that saw Sue is identical to the set of things that Sue was seen by, a piece of knowledge shared by all competent speakers of English. A largely parallel inference pattern will obtain for (39), with the crucial step (λy.sold(y,sue,a car) ⊆ λy.bought(sue,a

but not vice versa, since for q to be true the truck has to end up fully loaded while this does not hold for p. If John utters q, I can truthfully claim that John believes that p, but his utterance of p would not be enough for me to claim that John believes that q. Since entailment of p by q cannot be taken to be a sufficient condition for justifying the belief attribution on pain of losing an explanation of attributions of beliefs in necessary and impossible propositions, King has no obvious way of explaining these facts.

car,from(y))) licensed by the knowledge – again shared by all competent speakers of English – that the set of entities that sold Sue a car is identical to the set of entities that Sue bought a car from.

It might be objected at this point that our analysis fails to explain how it is that competent English speakers know what I claim they do. In part this is a fair complaint if our ultimate goal is a complete understanding of both what I understand when I know a language and how I come to have that understanding. Our analysis is based on an observation of what I know and gives no insight into how it is that I know that. The objection, however, in the end comes down to the mere observation that our analysis fails to go deeper than it does, which is an objection that can be levied against every analysis of any phenomenon that ever has been or ever will be produced. While I agree that a deeper explanation of how we know what we know is ultimately wanted, I take the depth of analysis given to constitute an improvement over past analyses, one that moreover points in a new direction for how to deepen our understanding of attitude attribution further.

Turn now to the second approach taken to accounting for attributions of belief in necessary truths and necessary falsities, the approach that adds impossible worlds to the mix. Such an addition will make it possible to distinguish the belief that 2+2=5 from the belief that 2+2=1+4, that 2+2=7+12, that $i^i$ is an imaginary number or that there is something that both exists and fails to exist, since each of these propositions will be true in mutually non-embedding sets of impossible worlds. It does not, however, offer an account for the range of beliefs that can plausibly be attributed to someone who indicates they have one of these impossible beliefs. Someone who declares "2+2=1+4" can thereby be said to believe that 2+2=5 but cannot thereby be said to believe that there is something that both exists and fails to exist. The set of impossible worlds in which 2+2=1+4, however, is not identical to and is not a subset or superset of the set of impossible worlds in which 2+2=5. If we maintain the possible worlds definition of entailment as a subset relation among worlds, extended to include impossible worlds as well as possible worlds, this means that 2+2=1+4 does not entail 2+2=5, and nor does the reverse entailment hold. Maintaining a Hintikkan approach to belief then results in the conclusion that believing that 2+2=1+4 is entirely independent from believing that 2+2=5, and so getting evidence of someone having the former belief should not be sufficient cause for attributing the latter. The fact that we easily accept attributing the latter belief to John upon hearing him declare that 2+2=1+4 is thus left without an explanation in an impossible worlds approach to impossible beliefs.[12]

## 6 Consequences of the analysis:

One of the consequences of the analysis, alluded to above in discussion of the shortest spy example, is that attitude contexts are transparent.[13] The embedded clause under the present analysis identifies not a belief in the head of the attitude holder but rather a proposition as analyzed by the attributor which is taken to follow from some such belief of the attitude holder. This makes it straightforwardly possible to allow quantification into attitude complements. Ambiguities involving quantifiers can thus be analyzed as normal scope ambiguities without the need to invoke all the apparatus of *de re* interpretation. This can be seen in the informal analyses available for the following example.

(41) John believes few people are happy.

[Few x: person(x)] (John has a belief from which *x is happy* can be inferred.)

---

12 I do not know of any work adopting an impossible worlds account of beliefs that addresses the question of how to extend the notion of entailment to encompass impossible worlds as well as possible worlds. The extension entertained here clearly fails to accomplish what is needed, and taking entailment to only apply among sets of possible worlds will also clearly fail to do so, since it will lead to no predictions about entailments among attributions of impossible beliefs.

13 This consequence does not hold for *de translato* attributions, as we will see below.

John has a belief from which *[Few x: person(x)] (x is happy)* can be inferred.

A related consequence is that existential generalization can apply unhindered in attitude contexts. From the attribution in (42a), for example, (42b) can be readily concluded.

(42) a. Ralph believes that Ortcutt is a spy.
b. There is someone who Ralph believes is a spy.

A further consequence of the analysis is that it maintains Semantic Innocence: the interpretation of sentences in attitude contexts is identical to that of sentences outside of attitude contexts. All of these consequences go against received wisdom. If Alice thinks that there are witches, for example, she could believe that several witches are evil. Intuitively, it does not follow from this that several witches are such that Alice believes they are evil. If Alice thinks one of the witches that she's only seen in her dreams, Esmerelda, is good, it does not follow intuitively that there is someone who Alice believes is good. And if Alice thinks Mary is a witch because she thinks that what it means to be a witch is to be a public figure, plausibly "Mary is a witch" is not being interpreted in the embedded clause in the same way it would be interpreted as a matrix utterance. In short, if the analysis stops here it succeeds at unifying *de dicto*, *de re* and *de qualitate* interpretations, but at the high cost of empirical coverage elsewhere. It is here that appeal to *de translato* interpretation becomes essential.

### 6.1 *De Translato*

On the analysis developed above, no formal semantic distinction is made between *de dicto*, *de re* and *de translato* interpretation. Distinctions exist, but they are all relegated to distinctions in how to draw inferences. Abandonment of any systematic semantic ambiguity unique to attitude attributions, however, goes one step too far. In addition to accounting for intuitions about the examples considered in the previous section, the existence of some semantic ambiguity plays a key role in explaining how sentences like the following could possibly be true:

(43) Ralph believes that Ortcutt is a spy, but he doesn't believe that ORTCUTT is a spy.

Under the standard view that takes *de dicto* interpretation to be semantically distinct from *de re* interpretation, the truth of (43) can be explained by analyzing the first clause *de re* (roughly: Ralph believes of Ortcutt that he is a spy) and the second *de dicto* (roughly: Ralph does not have a belief of the form *Ortcutt is a spy*). Lacking a *de re*/*de dicto* ambiguity, however, this path is not open under the proposed analysis. Worse, if we take the negation in the second conjunct to operate over that entire conjunct, as we should, the analysis as it stands predicts the second conjunct to be straightforwardly false given the truth of the first conjunct. The first conjunct will be true iff Ralph has a belief from which it follows that Ortcutt is a spy. The second will be true iff Ralph fails to have such a belief. Since Ralph cannot simultaneously have and not have such a belief, our analysis predicts that (43) should be a contradiction. And yet we can understand the sentence as true.

To account for the possible truth of (43) I introduce a distinct ambiguity argued for independently in Tancredi and Sharvit (forthcoming). The first clause I analyze exactly as argued for above, i.e. as saying that *Ortcutt is a spy* can be inferred from a belief of Ralph's. The second clause, however, I take to additionally involve *de translato* interpretation of the embedded clause. The core idea underlying *de translato* interpretation is that we sometimes attribute beliefs to people using language the way they use it rather than the way we use it. If John claims "26, a number with 4 factors, is prime", it is highly doubtful that he is using the word *prime* the way you and I use that word. Based on his utterance, however, we can easily claim:

(44) John believes that 26 is prime.

If we use our own characterization of the word *prime* as denoting a number with exactly two factors, this attribution comes out clearly false: John does not believe that 26 has exactly two factors. The attribution will be true, however, if the word *prime* can be assigned whatever interpretation John assigns it. Imagine we find, for example, that John takes a number x to be prime iff for some natural number n, $x = n^3 - 1$. By using *prime* to mean equaling one less than a perfect cube, the sentence in (44) will come out true, as desired. Formally, under a *de translato* analysis, (44) will come out as true iff John has a belief from which *[26 is prime]*$^{John}$ can be inferred, where the superscripted *John* indicates that the sentence is interpreted as John would interpret it rather than as we would standardly interpret it.

The significance of incorporating *de translato* interpretation into our analysis is that doing so allows us to account for examples like (43). Before getting to that example, however, it will be helpful to first look at a simpler case. Imagine our mathematically confused John making an additional claim. He says: "13, which has exactly 2 factors, is not prime". Based on the way I use the word *prime*, the fact that John believes 13 to have exactly 2 factors is basis enough for me to claim (45a). Based on the way John uses the word *prime*, I am equally justified in claiming (45b). But more importantly, if I interpret the embedded clause *de translato*, i.e. roughly as John would, I am also justified in claiming (45c).

(45) a. John believes that 13 is prime.
     b. John believes that 13 is not prime.
     c. John does not believe that 13 is prime.

Given an understanding of what John means by *prime*, (45c) on this *de translato* interpretation says that John does not have a belief from which it can be inferred that

13 is 1 less than a perfect cube. Given the situation as I have described it, this is very likely to be true.

With this as background, turn now to (43). I analyze the true interpretation of this example as deriving from a normal interpretation of the first conjunct combined with a *de translato* interpretation of the second. Informally, this leads to the following interpretation for (43):

(46) Ralph has a belief from which it follows that [Ortcutt is a spy (as I understand this sentence)], but Ralph does not have a belief from which it follows that [Ortcutt is a spy (as Ralph understands the sentence)].

Intuitively (46) is true in the situation envisaged earlier, where Ralph asserts that the man in the alley is a spy while also asserting that Ortcutt is not a spy. Given Ralph's understanding, the man in the alley is not Ortcutt – the two are distinct individuals. One of these two, according to Ralph, is a spy, the one he identifies as the man in the alley. Let us call him **a**. The other is not, the one he identifies as Ortcutt, who I will call **o**. Starting only from these beliefs of Ralph's, it is impossible to infer that Ralph's Ortcutt, i.e. **o**, is a spy. To see this, let us try to draw the inference in the obvious way:

(47) (Invalid inference)

The man in the alley is a spy

$\lambda x.x$ is a spy $\in \lambda P.$[the y: man-in-the-alley(y)](P(y)) (abstraction)

$\lambda P.$[the y: man-in-the-alley(y)](P(y)) $\subseteq \lambda P.(P(\mathbf{o}))$ (speaker's knowledge)

$\lambda x.x$ is a spy $\in \lambda P.(P(\mathbf{o}))$ (Ortcutt is a spy)

Both the inference pattern and the conclusion drawn here have the right form. However, the premise introduced as speaker's knowledge is not valid. According to the speaker, the man in the alley is Ortcutt. However, the speaker's Ortcutt cannot simultaneously be identified with Ralph's Ortcutt and with the man in the alley within a single world, and inference is under our proposal always world bound.[14] This is easiest to see if we understand the definite description used to describe Ralph's beliefs as referential. On this understanding, Ralph's belief is equivalent to the belief that **a** is a spy. Given that Ralph also believes that **o** is not a spy, to avoid unfairly taking Ralph to have internally contradictory beliefs, I have to conclude that **a** is not identical with **o**. This means that there is no guarantee that the properties that hold of **a** will be a subset of the properties that hold of **o**, speaker's knowledge notwithstanding.

The case of Ralph is a case in which the differences in interpretation between Ralph and the speaker cannot be reduced to a simple difference in extension within a fixed ontology. This sets these double vision cases apart from cases like John's misunderstanding of the word *prime*. In John's case, we had no reason to suspect that John's ontology was any different from our own and so could readily assign to his predicate *prime* an extension within our ontology. This is not the case with Ralph. In order to assign Ralph's interpretation to an expression it is necessary to adopt Ralph's ontology instead of our own. The extensions of the relevant predicates will come out wrong otherwise. Ralph's interpretation of the word *spy*, for example, includes **a** but not **o** in its extension, while his interpretation of *man* includes both **a** and **o**. If I add that Ralph believes that Ortcutt is, but the man in the alley is not, an upstanding citizen, then his interpretation of *upstanding citizen* will include **o** but not **a**. It is impossible to duplicate these predicate interpretations using the speaker's ontology which contains only one individual in place of **a** and **o**. The decision to interpret an attribution of beliefs to Ralph *de translato* then brings with it a commitment to Ralph's ontology in the conclusion to be derived through inference: once Ralph's ontology is adopted for drawing an inference, there's no going back.

14 More specifically, it is always about the actual world. See section on inference below.

Given the distinctions observed between the speaker's ontology and Ralph's and the effect they have on inference, we need to now go back and reassess our non-*de translato* interpretation of the first conjunct of (43). I claimed earlier that that conjunct comes out true because Ralph has a belief from which it can be concluded that Ortcutt is a spy, where "Ortcutt" occurs here as a name in the speaker's language. Given the different ontological commitments of me and Ralph, how does this inference go through? Taking speaker's "Ortcutt" to designate **O**, informally I want to say that in some sense $\mathbf{O} = \mathbf{a}$ and that $\mathbf{O} = \mathbf{o}$. However, there is no way of formalizing this intuition given the premise that $\mathbf{o} \neq \mathbf{a}$. Imagine, as one attempt to do so, connecting **O**, **o** and **a** via the properties that hold of them. If I think of **O** as somehow the fusion of **a** and **o**, for example, I could entertain the following possible semantic connection between the two.

(48) $\lambda P.P(\mathbf{O}) = \lambda P.P(\mathbf{a}) \cup \lambda P.P(\mathbf{o})$

Based on this I could infer from Ralph's belief that the man in the alley is a spy, formalized now as *spy(**a**)*, that Ortcutt is a spy, formalized as *spy(**O**)*:

(49) The man in the alley is a spy.

$\text{spy}(\mathbf{a})$

$\lambda x.\text{spy}(x) \in \lambda P.P(\mathbf{a})$ (abstraction)

$\lambda P.P(\mathbf{a}) \subseteq \lambda P.P(\mathbf{O})$

$\therefore \lambda x.\text{spy}(x) \in \lambda P.P(\mathbf{O})$ (speaker's Ortcutt is a spy)

Unfortunately, this approach is untenable. It fails to get off the ground right from the beginning. The reason for this can be seen by looking at the particular property used in the starting assumption of the inference, namely $\lambda x.spy(x) \in \lambda P.P(\mathbf{a})$. In order for this to be true, **a** has to be in the extension of $\lambda x.\text{spy}(x)$. Since **a** cannot be identified with **O** but **O** has to exist for the inference as a whole to go through, this inference has to be seen as taking as its starting point that the world contains two distinct individuals, **a** and **O**. Combined with parallel inferences about being an upstanding citizen, acceptance of (48) is tantamount to acceptance that there are 3 Ortcutt-related individuals in the world: Ralph's **a** and **o** plus the speaker's **O**. But this is something that nobody accepts, not Ralph, not the speaker, and not anyone else who knows of Ralph's predicament.

I get around this problem by separating the structured part of a structured proposition from its model theoretic interpretation, and taking tentative acceptance of Ralph's beliefs to start off as tentative acceptance of the structured part of a structured proposition. Like Soames (2014), I take propositions to be internal to an individual, though transmittable through language. Given a particular propositional structure, the first step in the inference process is to determine extensions for each part of that structure. This can be done either by making use of Ralph's presumed interpretation of that structure or by making use of the speaker's interpretation. The first case involves tentatively accepting that the world contains two individuals, **a** and **o**, denoted respectively by "the man in the alley" and "Ortcutt", despite the speaker's beliefs to the contrary. The second case involves taking the world to contain only one individual, **O**, denoted by both of these expressions. Formally, this step is not an inference step. It is a decision about whose interpretations and hence whose ontology to use for the purposes of determining what is to count as the content of the belief associated with a given structure. Once the interpretation of the relevant propositional structure is decided on, and hence the ontological commitments are fixed, inference can proceed apace.

### 6.2 Pet beliefs

The above analysis is based on the assumption that beliefs pair structured semantic objects with model theoretic interpretations. But what do we then make of attributions of beliefs to our pets? When Bill says "My dog Fido believes that I will feed him", this analysis makes it look on the surface like it requires Bill to presuppose that his dog has linguistically structured beliefs. We can avoid this consequence if we make a distinction between underlying beliefs and their semantic representation, however. For people with a fully developed language, we make the transition from underlying belief to semantic representation without even blinking. For Fido, however, this transition is impossible, Fido lacking the requisite linguistic abilities as he does. We can understand attributions of beliefs to Fido nonetheless in one of two counterfactual ways: (i) were Fido to have the ability to turn his beliefs into linguistically structured propositions, they would translate into a propositional structure from which one could infer that Bill will feed him; or (ii) were I to have Fido's underlying belief, it would be translated into such a propositional structure. The presumed structured part of a belief thus needs to be formally distinguished from a linguistically structured proposition, even if they stand in a 1-to-1 relation for competent language speakers.

## 7 What is inference?

The analysis proposed places a large burden on the process of inference. It is thus worth examining this process in more detail. Inference is a process we go through regularly in our daily lives. It involves starting with one or more propositions as premises and deriving from them at least one, typically distinct, proposition as conclusion. If inference is to operate entirely world-internally as proposed here, the propositions it operates on cannot be conceived of as mere truth values. This is where appeal to structure comes in. While all propositions that are true in a given world by definition have the same extension at that world, intuitively I do not wish to say that they are all the same proposition, nor even that they somehow all count as the same proposition with respect to that world. Associating propositions with compositional semantic structures is one way of distinguishing propositions within a world, and since semantic composition is already a well-established part of our apparatus for analyzing language, it is a way of making the needed distinctions at minimal theoretical cost.

I have posited above that all inference involves separating a proposition p into a function f and an argument a such that $f(a) = p$, and deriving a new proposition $g(a)$ from the fact that everything $f$ is true of $g$ is also true of. This inference process applies to extensions, not intensions. *That everything f is true of g is also true of* is merely a colloquial way of saying that the extension of $f$ is a subset of the extension of $g$. With respect to inference, the structuring of the propositions only serves to facilitate division of a proposition into a suitable function and argument. It is because of the crucial role that extensions play in inference that I do not admit inference to proceed from one proposition directly to another. Given that inference has to proceed from the truth of the premises to the truth of the conclusion, basing the process on propositional extensions alone would allow one to infer any true proposition from any other. The vast majority of such putative inferences, however, are not intuitively valid.

One could imagine trying to block this unwanted consequence by requiring inference to always be from one structured proposition to another without appeal to extension. While this would indeed succeed in blocking inferences among intuitively unrelated propositions, however, it would also have the unwanted consequence of blocking inferences that are felt to be valid. For example, what makes it possible to infer that most English professors are suspect given the premises (i) that everyone with an office in Building Z is suspect and (ii) that the people with an office in Building Z comprise two thirds of the English professors is not some structural affinity among the premises and conclusion. It is the extensional fact that given the premises, the properties in the extension of *everyone with an office in Building Z* form a subset of the properties in the extension of *most English professors*. Eliminating the role of extensions in the inference process would render this inference formally invalid.

The easiest way to implement the notion of inference needed in the proposed analysis is as inference about the actual world. Beliefs on this view are properties posited to hold of the actual world, as in King (2007). While no person's beliefs are sufficient to fully characterize the world, correct beliefs fit the world in a way in which incorrect beliefs do not. The world thus acts as a testing ground for beliefs, allowing us to accept some beliefs and reject others depending on how we find the world to be. While the actual world cannot be manipulated, since we do not have god's knowledge about the world, we are all of us in the position of not knowing which of our beliefs actually hold of the world. Conversation can be seen as having as at least one of its typical goals reducing our uncertainty about how the world is, and judging other people's beliefs plays a crucial role in achieving this goal.

Depending on our own degree of certainty about a subject under discussion, we may be willing to adopt another person's beliefs as giving a more accurate or more plausible description of the world, or we may take certain of our own beliefs to be solid enough as to warrant disputing any incompatible beliefs we discover others to have. These two approaches to others' beliefs correspond to two different ways we could reason based on what we take another person to believe. On the one hand, we could take their beliefs to be correct as we identify them in order to derive consequences from them. Alternatively, we could mold the beliefs into a form that is compatible with those of our own beliefs that we are unwilling to compromise on and determine what follows from the resulting combination of their beliefs and our own. Under the analysis proposed above, the former approach corresponds to typical *de translato* interpretation and the latter to non-*de translato* interpretation. When the beliefs of another person that we are reporting on are compatible with our own, then the *de translato* / non-*de translato* distinction becomes semantically inert.

There are three important ways that another person's beliefs can differ from our own. Like in the case of Ralph, there can be a difference in assumed ontology: one person can see two objects where the other sees only one. Alternatively, there can be a difference in assumed meaning: one person can think the word *prime* means having exactly two factors while the other thinks it means being one less than a perfect cube. Finally, there can be a mere difference of extension: two people can both understand intuitively what it means to be a professor while one person takes *professor* to apply to a,b,c and another takes it to apply to b,c,d. All three types of difference can lead to inconsistencies when attempting to fit both people's beliefs to the real world.

Suppose, now, that I wish to draw inferences based on a purported belief of another person when I know that that belief is inconsistent with my own beliefs. Mary utters "every math professor is tall", for example, claiming that a, b and c are the relevant math professors, while I take the math professors to be b, c and d, with a being a mere math assistant. Fully apprised of all the facts, I can infer one of two things from Mary's belief: either that a, b and c are tall, or that b, c and d are tall. The first comes from accepting not only Mary's statement but also Mary's assumed extension for *math professor*, functioning here as the domain of quantification. The latter comes from accepting Mary's statement but evaluating it using my own extension for *math professor*. Only the first way can be seen as deriving consequences from Mary's statement that she herself will be likely to endorse. Still, the second way is used regularly, and is common in arguing against a claim one thinks to be false.

Turn now to Ralph. Ralph uttered "the man in the alley is a spy. Ortcutt is not a spy." I see the man in the alley and Ortcutt as a single person, though obviously Ralph does not. Here again I have two choices for how to draw inferences. Both start from tentatively accepting Ralph's statement to see what follows. The first way also accepts Ralph's ontology and hence the values assigned to the expressions he used. Doing so makes it impossible to conclude that Ortcutt is a spy since Ortcutt is under these assumptions a distinct individual from the man in the alley. The second way rejects Ralph's ontology, substituting in the speaker's ontology instead and assigning a single individual as the denotation of both "the man in the alley" and "Ortcutt". This assumption makes it possible to conclude from the statement "the man in the alley is a spy" that Ortcutt is a spy, and similarly allows one to conclude from "Ortcutt is not a spy" that the man in the alley is not a spy.

Finally consider John. John claimed "26, a number with 4 factors, is prime". I can here again draw inferences in two ways, both starting from the tentative acceptance of the sentence "26 is prime" as a linguistically structured but not yet model-theoretically interpreted proposition. The first mode of inference tracks John's understanding, using the property he associates with the word "prime". Given his understanding of that word, I can conclude that 26 is one less than a perfect cube. The second mode of inference makes use of our own understanding of the word "prime". Given this understanding, I can conclude that 26 has exactly 2 factors. The first mode of reasoning leads to a true conclusion and the second to a false conclusion, but both ways of reasoning are valid.

## 8 Philosophical questions:

I turn now to consideration of three philosophical claims revolving around belief: Kripke's puzzle about belief, Putnam's twin Earth based claim that belief ain't in the head, and Burge's arthritis in the thigh based claim that mental contents have an ineliminable social, i.e. non-internal, component. In all three cases I show that the conclusions drawn do not follow from the observations they are based on.

### 8.1 Kripke's puzzle about belief:

Kripke imagines an individual, Pierre, who grows up as a monolingual French speaker in France, during which time he comes to believe that, as he would put it, Londres est joli, translatable into English as "London is pretty". Later he moves to London, where he picks up English by the direct method, never translating between English and French. Being in a rather run down part of the city, he comes to believe that, as he would put it, London is not pretty. He hasn't given up his beliefs from his childhood in France. He has simply added to them. Kripke's puzzle comes in the form of a question: Does Pierre or does he not now believe that London is pretty? Kripke takes no answer to this question to be satisfactory. Starting from the idea that Pierre's beliefs are about London itself, to answer "yes" is to acknowledge his French beliefs at the expense of his English beliefs, to answer "no" is to do the opposite, and to answer both "yes" and "no" is to commit to a contradiction.

On the analysis argued for here, while Pierre's beliefs are in one sense about London, London itself is not a component of his beliefs. It is the unique real-world London that both his "London" thoughts and his "Londres" beliefs are in fact about, but there is no need for the content of his beliefs to reflect this fact. Indeed, it is the fact that Pierre fails to make this connection that gives rise to the seeming paradox in the first place. If London itself does not figure in his thoughts, however, and there is no implicit requirement that all of his thoughts about London be internally identified as being about the same object, then no puzzle arises regarding the content of his beliefs.

Given the analysis proposed for interpreting belief attributions, this problem can be reduced to the double vision problem of Ralph. Like Ralph, Pierre sees two things where the speaker sees one. This makes the properties Pierre takes to hold of the world incompatible with the properties the speaker takes to hold of the world. They cannot both be true of the same world, and so in drawing inferences from Pierre's beliefs I have to make a choice: do I respect his implicit belief that there are two distinct objects in the world identified by the names "London" and "Londres", or do I respect my own belief that these two names are names of the same one place? I cannot accept both ideas simultaneously on pain of contradiction, but neither choice on its own is contradictory. Pierre's belief, of course, turns out to get the facts wrong – it fails to fit the world in a particular way in which my beliefs succeed in fitting it – but that's the worst that could be said of them.

The analysis of belief attribution proposed above took it as a given that people can differ in the properties they take the world to have in ways that are incommensurable. Of particular relevance here is the fact that these properties can include ontology: two people can disagree on what objects the world contains in ways that cannot be reconciled. The fact that both people take their beliefs to be about a real world object does nothing to ensure commensurability in their ontologies. At most it commits them to accept that observations about the world can bear on the question of which one (if either) has their ontology right.

Kripke, no doubt, may well accept the description given above of the situation as in some sense complete, and yet it is doubtful he would take it to bear on the paradox. The reason for this, I suspect, is that Kripke is interested in modeling the beliefs of both Pierre and others with a single model in which there can be but a single object that is London. If this requirement is imposed on an analysis of belief, then Kripke is correct that the paradox is unresolvable. The rational reaction to this conclusion, however, is not to decide that the paradox is unresolvable but to challenge one of the underlying assumptions that leads to this conundrum. After all, in the situation described it is observationally acceptable to claim both that Pierre believes London is pretty and that he does not believe that London is pretty without contradicting oneself. If the purported analysis fails to account for this fact then the problem is not in the way people speak but in the analysis thereof. The obvious candidate assumption to challenge is the assumption that all speakers' beliefs have to be able to be modeled by a single model that contains only a single object that is London. While such a model might be sufficient for your beliefs and mine, such a model is patently ill equipped to model Pierre's beliefs.

The analysis presented above dissolves Pierre's puzzle by assigning distinct models to distinct individuals for giving interpretations to shared structured propositions. If we think of a model as a way of representing all the ways that individuals in its domain can be related, differences in domain will give rise to different interrelations among individuals, i.e. to different sets of possible worlds. On this view, the only thing that connects models based on different domains is the communication-based need for a model to fit the real world. It is only the connection through the real world that makes one person's model relevant to another person's.

This extended use of models[15], what could be called I-models, i.e. interpretational components of what Chomsky (1986) dubs I-languages, still makes it possible to adopt Kripke's analysis of modality. Metaphysical necessity, for example, can be identified as necessity with respect to a single I-model. Multiplication of models brings extra benefits, however. In particular, it allows for an analysis of epistemic modality that can make sense of statements like "John might be Bill, and he might not be", looked at in Tancredi (2009). Restricted to a single model with names interpreted as rigid designators, this sentence should be a contradiction. Within any single model, either John is Bill, in which case John is Bill in every world of the model in which John exists, or he is not, in which case the non-identity of the two holds in every world of the model in which they both exist. In the one case, "John might be Bill" will come out true and "he might not be" false. In the other case, the truth values switch. In no model, however, are both sentences true. If we analyze epistemic modality as quantifying over epistemically possible models in addition to over possible worlds within a model, however, then we can easily account for the simultaneous truth of both conjuncts. We can similarly account for similar problems related to logical modality. "For any distinct names a and b, it is logically possible that a=b and it is also logically possible that a≠b". By adopting an extended use of models, this can be analyzed as saying that there is a logically possible model in which in some world a=b, and that there is a logically possible model in which in some world a≠b. For each model, of course, if the identity/non-identity holds in one world it will hold in them

15 See Tancredi (2007, 2008, 2009) for a formal analysis of modality and propositional attitudes based on multiple models.

all. The models themselves, however, need not be mutually compatible.

### 8.2 Putnam's twin Earth:

Putnam provides us with another thought experiment. We are to imagine two individuals, Oscar on Earth, and twin Oscar on twin Earth, identical molecule for molecule, down to the last detail. Earth and twin Earth are also as identical as it is possible to be except for one pervasive difference: where earth has $H_2O$, twin Earth has XYZ and vice versa. Inhabitants of both worlds call the liquid that fills their lakes and streams and that falls as rain from the sky "water". Putnam notices that we can describe this situation by saying that *Oscar believes that the lakes around him are filled with water but twin Oscar does not*. From this he concludes that beliefs ain't in the head. However, while I accept his observation as correct, his conclusion does not follow. Indeed, the analysis proposed in this paper gives a perfectly straightforward way of accounting for the truth of the attributions while taking beliefs to always and only be in the head. According to that analysis, the attribution says that *Oscar has a belief in his head from which it follows for the speaker that the lakes are filled with water, and Twin Oscar does not*.

The difference in the belief attributions we are willing to make to Oscar and twin Oscar does not lie in an internal representation difference between Oscar and Twin Oscar. It lies in an extensional difference in what their beliefs are about. If this difference is known to the speaker, as it is taken to be as part of the thought experiment, then it cannot be ignored in drawing inferences from Oscar's and from Twin Oscar's beliefs. To draw inferences from a putative belief of Oscar's, I need to know what that belief is about. In this case, it's about what both Oscar and the speaker call "water", i.e. $H_2O$, so it is straightforward for the speaker to take Oscar's putative belief and infer from it that lakes are filled with water, i.e. speaker's water, i.e. $H_2O$. To draw parallel inferences from a putative belief of twin Oscar's, I start from the knowledge that twin Oscar's beliefs are about XYZ. From the proposition that the lakes are filled with XYZ, it is impossible for the speaker to infer that the lakes are filled with $H_2O$. He thus cannot use his own term “water” in describing the beliefs of twin Oscar, accounting for our understanding of the claim that twin Oscar does not believe that the lakes are filled with water. Significantly, coming to this conclusion does not require positing any notion of what constitutes a belief other than the naïve one that locates beliefs in the head.

Of course, on the analysis I proposed of belief attribution it also follows that it should be possible to say that twin Oscar does believe that the lakes around him are filled with water, and I believe that this is a correct prediction. The sentence comes out true when I use twin Oscar’s term “water”, complete with its associated extension of XYZ, i.e. when I interpret the embedded clause *de translato*. And given the freedom to switch between *de translato* and non-*de translato* interpretations freely, I correctly predict that the situation can be described as:

(50) Twin Oscar believes that lakes are filled with water, but he doesn’t believe that they are filled with WATER.

On its true understanding, the embedded clause is interpreted *de translato* in the first conjunct and non-*de translato* in the second. Our analysis leads to the further prediction that a speaker on Earth for whom “water” is known to refer to $H_2O$ will not be able to make a parallel pair of attributions to Oscar.

(51) #Oscar believes that lakes are filled with water, but he doesn’t believe that they are filled with WATER.

Since such a speaker shares with Oscar the extension assigned to the word “water”, there is no basis for contrast between the occurrence of “water” in the first clause and that of “WATER” in the second, and hence there is no combination of *de translato* and non-*de translato* interpretations that will render the sentence true.

From these examples and their analysis, we see that Putnam is only partially vindicated. He is vindicated in his discovery that belief *attributions* depend on more than the identity of internal beliefs, since in all cases, *de translato* included, inference based on attributor's knowledge can affect the truth of the attribution. His conclusion that belief ain't in the head, however, goes out the window, as it should. The inferences can only get off the ground with an initial premise, and internal beliefs supply that premise.

### 8.3 Burge's arthritis in the thigh:

Burge follows Putnam in arguing that the mental contents of beliefs include more than just what is in the head. He gives the following, slightly modified, example to support this view.

Alan has a large number of attitudes commonly attributed with content clauses containing 'arthritis' in oblique occurrence. For example, he thinks (correctly) that he has had arthritis for years; that his arthritis in his wrists and fingers is more painful than his arthritis in his ankles, that it is better to have arthritis than cancer of the liver, that stiffening joints is a symptom of arthritis, that certain sorts of aches are characteristic of arthritis, that there are various kinds of arthritis, and so forth. In short, he has a wide range of such attitudes. In addition to these unsurprising attitudes, he thinks falsely that he has developed arthritis in the thigh.

With this as the actual situation, we are asked to consider the following counterfactual situation: Alan might have had the same physical history and non-intentional mental phenomena while the word 'arthritis' was conventionally applied, and defined to apply, to various rheumatoid ailments, including the one in Alan's thigh, as well as to arthritis.

According to Burge, in the counterfactual situation, Alan lacks some – probably all – of the attitudes commonly attributed with content clauses containing 'arthritis' [the word as actually used by you and me, not as used in the counterfactual situation CT] in oblique occurrence. He lacks the occurrent thoughts or beliefs that he has arthritis in the thigh, that he has had arthritis for years, that stiffening joints and various sorts of aches are symptoms of arthritis, that his father had arthritis, and so on. He goes on to conclude that however we describe the patient's attitudes in the counterfactual situation, it will not be with a term or phrase extensionally equivalent with 'arthritis'. So the patient's counterfactual attitude contents differ from his actual ones…. [T]he patient's mental contents differ while his entire physical and non-intentional mental histories, considered in isolation from their social context, remain the same.

Once again I am in a position to be able to grant the observation and deny the conclusion. The description of the situation, however, is theory laden in a way that is at odds with our analysis. Burge talks of "attitudes commonly attributed with content clauses containing 'arthritis' in oblique occurrence". He includes in such attitudes Alan's believing that he has arthritis in the thigh. But what counts as an "oblique occurrence" for Burge? While he does not specify, I assume he means roughly "in a position that does not allow for substitution of actual identicals". Under the analysis I have given, substitution of identicals is always possible if the result is interpreted with respect to the model w.r.t. which the terms are identical. Thus, in a *de translato* interpretation of a belief attributed to John, if John takes *prime* and *composite* to be synonyms, the one can be substituted for the other in a belief attribution to John as long as the result is also interpreted de *translato*. The fact that I equate *prime* with *having exactly two factors*, on the other hand, will only license substitution in belief contexts when interpreted non-*de translato*. If we take actual identicals to be expressions that are synonymous for the speaker, then substitution of actual identicals is only guaranteed to be possible under our analysis for non-*de translato* interpreted expressions.

With this as background, what Burge appears to be saying when translated into the terms of our present analysis is that the speaker's word "arthritis" cannot be used *de translato* to describe the counterfactual situation. But this is hardly surprising: the speaker's word "arthritis" cannot be used *de translato* to describe the actual situation either. The actual situation requires use of Alan's word "arthritis", which is both extensionally and conceptually distinct from the speakers in the actual world. This is parallel to John's usage of the word "prime" above, where John associated a different concept with the word than the speaker. Since Alan's word "arthritis" includes diseases of the thigh in its extension, there is no problem using Alan's word to describe not only the actual situation but also the counterfactual situation from our perspective. The fact that the surrounding community agrees with Alan's usage in the counterfactual situation is of no import. The belief attribution depends only on the meaning Alan assigns to the word.

Given this state of affairs, the proposal of this paper predicts that the following sentence can be used by us in the actual world to describe Alan in the actual world, and it can be used by us in the actual world to describe Alan in the counterfactual situation, but it cannot be used by us in the counterfactual situation (where I take "arthritis" to apply to diseases of the thigh too) to describe Alan in the counterfactual situation.

(52) Alan believes that he has arthritis in the thigh, but he doesn't believe that he has ARTHRITIS in the thigh.

In the actual-actual and actual-counterfactual cases, "arthritis" in the first conjunct is interpreted *de translato* while that in the second conjunct is not. Since the two interpretations differ, a contrast can be generated. In the counterfactual-counterfactual situation, on the other hand the *de translato* interpretation ends up identical with the non-*de translato* interpretation, leaving no room for contrast between the two conjuncts and therefore only giving rise to a contradictory interpretation. Under Burge's assumptions, on the other hand, this sentence should only give rise to a contradictory interpretation since in all cases the two occurrences of "arthritis" are required to be occurrences of the same word. Here, intuition clearly supports the proposed analysis.

## 9 Sentential propositions and *the proposition that S*

On our analysis above, "a believes that S" relates a person to an implicit belief *b* in addition to relating that person to the interpretation of S (as inferable from *b*). The implicit belief on this analysis essentially serves as an implicit argument of the verb. Support for this view comes from differences between such statements and statements of the form "a believes the proposition that S". While statements of the former form allow for an inferential connection between the proposition denoted by S and the presumed internal belief, statements of the latter form do not. To see this, consider a situation in which Building Z is burning. I know, but John does not, that two thirds of the English professors are in Building Z. In this situation, when John says "Everyone in Building Z is in danger", I can report on the situation as in (a), but not as in (b).

(53) a. "John believes that two thirds of the English professors are in danger."
b. "#John believes the proposition that two thirds of the English professors are in danger."

The report in (a) is predicted to be acceptable given the analysis of "a believes that S" provided above. Extending that analysis to (b) would lead to the incorrect prediction that (b) should be as acceptable as (a) in the situation given. However, there is no reason to have to extend the analysis to (b). If we take the NP headed by "proposition" to directly designate a structured proposition that is part of John's purported beliefs, then we expect (b) to come out as false in the situation described. Separating internal beliefs from their inferable consequences in the semantics makes such an analysis possible.

## 10 On the Judge role

I take the Agent parameter and the Judge parameter to play formally distinct roles, despite the fact that with respect to an utterance the two roles will be filled by the same individual. The difference needs to come out as something like the strict / sloppy identity difference for pronouns. Reference to an Agent acts like strict reference. If a speaker A makes reference to herself as Agent, then hearers judge the truth of what was said as depending on how things stand with A. Reference to a Judge, on the other hand, acts like sloppy reference. If a speaker A makes reference to herself as Judge, then hearer B will judge the truth of what was said as depending on how things stand with B, not A.

There is no Judge-based pronoun in English, but I take the role of Judge to play a role in matrix utterances nonetheless. In particular I take the Judge role to be implicated in the inference encoded in belief attribution. To illustrate, if John utters "Mary believes it is raining", I can conclude that John believes that *it's raining* can be inferred by him from a belief of Mary's. For me to accept the sentence as true, however, it is not enough that I accept that John has this belief. I have to accept that *it's raining* can be inferred *by me* from a belief of Mary's as well.[16] This is a "sloppy" interpretation, where the acceptability of the sentence depends on whoever is judging the sentence.

**Bibliography**

Bach, Emmon 1997 "Do Belief Reports Report Beliefs?" *Pacific Philosophical Quarterly* 78, 215–241.

16 Perhaps more cautiously we should say that *it's raining* can be inferred on the basis of my beliefs, something I might be willing to grant even in cases where I know that I myself do not have the ability to construct the necessary inference from my own beliefs.

Berto, Francesco 2013, “On Conceiving the Inconsistent,” *Proceedings of the aristotelian society 135th session*, issue no. 1 volume CXIV.

Burge, Tyler 1979 “Individualism and the Mental” In P. French, T. Uehling, and H. Wettstein, eds. *Midwest Studies in Philosophy: Studies in Metaphysics*. University of Minnesota Press, Minneapolis.

Burgess, John 2013 *Kripke*, Polity Press, Cambridge, UK.

Chomsky, Noam 1986 *Knowledge of Language: Its Nature, Origin and Use*. Greenwood Publishing Group.

Cresswell, Maxwell 1995 *Structured Meanings*, MIT Press, Cambridge.

Cresswell, Maxwell, & Arnim von Stechow 1982 *“De Re Generalized” Linguistics and Philosophy* vol. 5 no. 4, 503-535.

Hintikka, Jaakko 1962 *Knowledge and Belief: an Introduction to the Logic of Two Notions*, Cornell University Press, Ithica.

Jago, Mark 2014 *The Impossible: An Essay on Hyperintensionality*, Oxford University Press, Oxford.

Kaplan, David 1968 “Quantifying In”, *Synthese* 19, 178-214.

King, Jeffrey 2007 *The Nature and Structure of Content*, Oxford University Press, Oxford.

King, Jeffrey 2017 “W(h)ither Semantics!(?)”, *Nous* 52: 4, 772-795.

King, Jeffrey 2018 “On Propositions and Fineness of Grain (Again!)”, ms. Rutgers.

Kripke, Saul 1979 “A Puzzle about Belief” in A. Margalit (ed.) *Meaning in Use*, 239-283, Dordrecht, D. Reidel.

Putnam, Hillary 1975 “The Meaning of Meaning” in K. Gunderson (ed.) *Language, Mind and Knowledge*, Minnesota Studies in the Philosophy of Science, VII, University of Minnesota Press, Minneapolis.

Quine, Willard van Orman 1956 “Quantifiers and Propositional Attitudes” *The Journal of Philosophy*, Vol. 53, No. 5., pp. 177-187.

Schwager, Magdalena 2010 “Speaking of qualities”. In Ed Cormany, Satoshi Ito & David Lutz (eds.), *Proceedings of Semantics and Linguistic Theory* (SALT) 19, 395–412. eLanguage.

Tancredi, Christopher 2007 “A Multi-Model Modal Analysis of I-Semantics: Part 2, Identity and Attitudes,” in Yukio Otsu (ed.) *Proceedings of the Tokyo Conference on Psycholinguistics 8*, pp. 21-55.

Tancredi, Christopher 2008 “Multiple Models,” report for The Center for Advanced Research on Logic and Sensibility, pp. 1-32.

Tancredi, Christopher 2009 “Domains of Quantification, Rigid Designation and Modality: The Case for Multiple Models”, talk handout, NYU.

Tancredi, Christopher, and Yael Sharvit forthcoming “Qualities and Translations” *Linguistics and Philosophy*.

Yagisawa, Takashi 2010 *Worlds and Individuals, Possible and Otherwise*, Oxford University Press, Oxford.