

DRAFT: please do not quote

De Jure Anti-Coreference and Mental Files

N.A. Pinillos

Arizona State University

Abstract:

I introduce the phenomenon of de jure anti-coreference and argue that in contrast to its opposing notion, de jure coreference, it is rarely found in human representational systems. I explain how the Fregean can hope to explain this asymmetry by appealing to senses or mental files. I argue, however, that such approaches, in order to account for dynamic coordination, must ultimately appeal to semantic relationism.

1. Coordination

We possess the indispensable ability to track objects in our environment. We can do so not only visually and with our other senses, but also in thought and communication.

With natural languages, this can be achieved in many ways, including repetition of the same word and by the use of anaphoric devices. These methods, when applied correctly, are epistemically rewarding. They typically place the agent in a position to know that it is the same object that is being represented.

In contrast, there are ways of tracking the same object that are much less epistemically rewarding. If I use distinct names for the same object as in ‘hesperus’ and ‘phosphorus’, then I manage to represent the same thing twice over, but I do so in an epistemically opaque manner. If participants manage to know that the same thing is being represented, it will be known in a less direct way. Following the literature, we call the first way of coreferring ‘de jure’ and the second, ‘de facto’. I will use the term ‘coordination’ as a broad term to include not only de jure coreference but also cases that are just like those except where reference fails (as when the non-referring term ‘Vulcan’ is repeated).

We may hope for an explanation of coordination, or even better, a reductive account. But consider the possibility that from the perspective of semantic theorizing, coordination is just a primitive relation. The relation can be implemented in different ways. In the familiar formal languages, repeating the same symbol gives rise to the phenomenon. But we can imagine a formal language where coordination happens between different symbols (say all the ones whose subscript numerals represent prime numbers). This suggests coordination is multiply realizable.¹

The Fregean, however, disagrees. She offers a semantic reduction of coordination to senses or their modern naturalized incarnation, mental files. No special provisions beyond them are needed.² The simplest explanation says that coordination

¹ It’s not mysterious to speak of semantically irreducible but multiply realizable notions. Arguably, denotation (or some such representation-world relation) is such a notion. There may be different ways for representations to pick out things in the world, but a compositional theory of meaning may be agnostic between these.

² The idea has its roots in Frege (1956) and has been carried out by a number of neo-Fregeans including Evans (1985b), Campbell (1987) and Recanati (2012). Kit Fine (2007), has argued that such a reduction

between A and B happens for human thought and language whenever A and B are associated with the same mental file.

After some stage-setting (section 2), I develop an argument (sections 3, 4 and 5) that actually favors the Fregean approach over recent relationist accounts of coordination. The Fregean, unlike the relationist, has an elegant explanation for why a certain asymmetry is found in human representational systems (why *de jure* coreference but not *de jure* “anti-coreference” is pervasive). I argue, however, that despite initial appearances, Fregeans actually lack the resources to account for coordination that is dynamic (sections 6 and 7). I recommend that the Fregean adopt a relationist strategy.

2. Two Tests.

I begin by focusing on *de jure* coreference proper (cases of coordination where reference is successful). Kit Fine (2007) has provided a test for this notion.³ Here’s the explanation in his own words:

A good test for when an object is represented as the same [*de jure* coreference] is in terms of whether one might sensibly raise the question of whether it *is* the

fails though it is less clear to me about what he thinks about the implementation projects. In general, there has not been a clear distinction in the literature between implementation and reduction.

³ Fine holds that the referential cases of coordination are primary and the failed reference cases are explained in the terms of the former. In my view, the general notion (coordination) is primary and the successful cases (*de jure* coreference) are explained in terms of the general notion plus referential success. I hold this because, contra Fine, I do not think the failed reference cases can be explained in terms of the successful cases (Pinillos 2015). In this section, I will be discussing Fine, so I begin with the notion that requires successful reference.

same. An object is represented as the same in a piece of discourse only if no one who understands the discourse can sensibly raise the question of whether it is the same (page 40) [his italics].

Fine immediately goes on to give an example of a specific application

Suppose that you say 'Cicero is an orator' and later say 'Cicero was honest', intending to make the very same use of the name 'Cicero'. Then anyone who raises the question of whether the reference was the same would thereby betray his lack of understanding of what you meant (page 40).

It is not totally clear from the explanatory part quoted (first part of the quoted material), what it means to raise the question of whether *it* (the object represented) is the same. How, exactly, can such a question be raised? In the example (the second part of the quoted material) Fine talks about raising the question of 'whether the reference was the same'. But there is still some unclarity. The expression 'the reference' could be interpreted either de re or de dicto, and a de dicto reading could be filled in various ways. A de dicto meta-linguistic reading of 'the reference' makes for a plausible test:

- (1) S raises the question of whether the referent of the first occurrence of 'Cicero' is the same as the referent of the second occurrence of 'Cicero'.

This gloss gives us what we want, since it does seem that in the discourse Fine discusses, if someone were to raise the question in (1) then we would think that they don't understand the utterance.

Another issue concerning Fine's test is that in the explanatory section (first part of the quoted material) he talks about understanding the discourse, but in the example he talks about understanding what the speaker meant. Since a speaker might be mistaken about the semantics facts concerning the discourse, the notions can come apart. I propose we focus on the discourse.

Putting these ideas together suggest the following:

Fine's test: If representation occurrences *a* and *b* are de jure coreferential in a discourse, then if a rational agent understands the discourse, they cannot sensibly raise the question of whether *a* and *b* refer to the same thing.^{4 5}

⁴ It is good to remember that this is only a test. So it need not even aim to give us strict necessary conditions for de jure coreference. Imagine a sophisticated philosopher who is wondering whether error theory about intentionality is true (because of worries about normativity). They will therefore wonder whether any word will refer to anything at all. A fortiori, they will then wonder, for any discourse, including those involving de jure coreference, whether the occurrences at issue refer to the same thing. This would be a counter-example if the test aimed to also give us necessary conditions for de jure coreference. However, we can treat this sort of case as "fringe" and so save the test. As an aside, I add that the alleged counter-example is interesting for a different reason. It constitutes a challenge to the claim that competent speakers of a language must know the various semantic facts concerning their language. For an explicit statement of this thesis see Larson and Segal (1995).

⁵ A reasonable way of characterizing the phenomenon assumes that de jure coreference happens between occurrences or tokens and not expression *types*. This is because we want to say that in the trivial 'Cicero is Cicero', there is a sense in which de jure coreference holds between two things and not just between something and itself (the expression type 'Cicero'). It may be that the correct theory will explain the phenomenon in terms of a relation holding between an expression type and itself, as Fiengo and May (1994) argue. But this should be seen as a theoretical explanation, and not part of the "theory-neutral" description of the phenomenon.

The test, as stated, is approximately adequate though there is an important class of examples that it misdiagnoses (or at least, it is unclear about). Consider a third person belief report concerning Jonah the biblical character. Let us also suppose that ‘Jonah’ does refer:

(2) Steve thinks if *Jonah* got eaten by a big fish, *he* would have disliked it.

Pre-theoretically, the terms in italics are de jure coreferential. However, an adult can understand this third person attribution perfectly and still wonder whether ‘Jonah’ refers at all, and a fortiori, whether ‘Jonah’ and ‘he’ refer to the same thing. It would be preferable to have a test that covers these types of examples.

One way of resolving the problem just discussed is to switch from talk of coreference to talk of conditional coreference. Consider the following modification of the test:

(Fine’s Test Modified) If representation occurrences *a* and *b* are de jure coreferential in a discourse, then if an agent understand the discourse, they cannot sensibly raise the question of **whether *a* and *b* refer to the same thing if they refer at all.**

This particular construal avoids the previous problem. The rational person who understands the belief attribution about Jonah (‘Steve thinks if *Jonah* got eaten by a big

fish, *he* would have disliked it') and is wondering whether 'Jonah' refers at all, cannot also wonder whether 'Jonah' and 'he' refer to the same thing if they refer at all. That is, the person who understands the attribution won't doubt the proposition expressed by the following: 'if 'Jonah' and 'he' both refer, then they refer to the same thing'.

Here's a different worry about Fine's test. Fine appeals to the notion of whether it is sensible to raise a question about coreference.⁶ But consider a situation in which one is forced at gunpoint to raise a question about coreference. Isn't this a situation in which it is sensible to raise a question about coreference? Or consider a situation in which it is simply impolite or too time consuming to raise a question about coreference, isn't this a situation in which it is not sensible to raise a question? Of course, these are not senses of "sensible" that Fine had in mind. We need an epistemic not a practical interpretation. But it is very unclear to me how to give the locution such a reading.

In previous work I have emphasized that the correct epistemic notion associated with de jure coreference is knowledge. There is good reason to think the concept plays a central role in predicting and explaining behavior.⁷ So it is particularly apt for our purposes:

(Knowledge of Conditional Coreference—KCC) If representation occurrences *a* and *b* are de jure coreferential in a discourse, then if an agent understand the

⁶ Other theorists have considered related notions. Schroeter (2007) and Campbell (1987) discuss Cartesian certainty before rejecting it.

⁷ Williamson (2000).

discourse, they know that *a* and *b* refer to the same thing if they refer at all (Pinillos 2006, 2011, Recanati 2012).

Let's revisit one of our earlier cases, 'If *hesperus* is a planet, then it revolves around the sun'. Here, someone who understands this utterance (in a natural discourse) would know that the occurrences of '*hesperus*' and '*it*' refer to the same thing if they refer at all. It's not just that they believe it or have some justification for believing it. Something stronger follows from their understanding of the discourse: knowledge.⁸

I prefer KCC over Fine's Modified Test. Although the differences are subtle, they are important. I illustrate this by giving a specific scenario which KCC and Fine's test give different results.

For a variety of purposes, we sometimes willingly presuppose falsehoods. Although there are subtle differences between compounds A and B, a chemist may treat them as the same when solving a problem. So in the context of solving the problem, it is presupposed that 'A' and 'B' corefer. It is plausible that when something is presupposed, conversational participants cannot sensibly raise a question about whether it is true. If they do raise such a question, then in virtue of raising such a

⁸ One worry is that agents who understand the discourse might not have grasped the concept of 'reference'. If they haven't, then they can't know that some terms refer to the same thing if they refer at all. Of course, I agree that competent agents need not possess lexical items 'reference', but I doubt that a competent speaker could lack the concept of reference. Understanding discourse often involves solving the problem of figuring out who or what is being talked about, and who or what is being talked about with this name or that pronoun. Arguably, solving this problem requires having a concept of reference. But even if we concede the point, we can get around this worry in a familiar way. Instead of talking about what an agent who understands the discourse knows, we can talk about what she is in a position to know if she possessed the relevant concepts. This should assuage the worries that the condition over-intellectualizes the phenomenon.

question, the proposition no longer counts as presupposed in the discourse (it is no longer assumed or taken for granted). If this is right, then in a discourse in which it is presupposed that 'A' and 'B' corefer, it is not sensible to raise a question about whether it is true. Hence, Fine's modified test cannot give us the right result that in fact 'A' and 'B' are not de jure coreferential. However, KCC will give us the right result. This is because our agent will nonetheless not know that the terms corefer if they refer at all.

3. De jure anti-coreference.

In this section, I introduce the notion of de jure anti-coreference. I argue that there is an interesting asymmetry between it and de jure coreference. De jure coreference is commonplace, but de jure anti-coreference is rare.

We can use the template for KCC to characterize other de jure semantic notions beyond coreference. For example, we could provide a condition for de jure anti-coreference:

(Knowledge of Conditional Anti-Coreference-KCAC) If representation occurrences *a* and *b* are de jure anti-coreferential in a discourse, then if an agent understands the discourse, they know that *a* and *b* refer to distinct objects if they refer at all.

Consider the following utterance

(3) Mercury is closer to the sun than Venus

Aren't 'Mercury' and 'Venus' here de jure anti-coreferential? After all, many of us know that 'Mercury' and 'Venus' refer to distinct objects. But must someone who understands the utterance know that the names refer to different objects if they both refer? I do not think so. Suppose you just have minimal evidence that the names refer to distinct planets (imagine you are learning about the planets for the first time). The evidence is not strong enough for knowledge. This does not bar you from understanding the utterance.

Here's another way of putting the point. Suppose someone wanted to know whether or not Mercury and Venus were the same planet. Suppose all they knew was that (i) 'Mercury' refers to Mercury, (ii) 'Venus' refers to Venus and (iii) 'Mercury' and 'Venus' both refer to existing planets. If 'Mercury' and 'Venus' in (3) really were de jure anti-coreferential, then (according to KCAC) all the agent would need to do to learn the astronomical fact is understand (3) and deploy a disquotational scheme.⁹ But this seems like too easy a way to know.

Perhaps understanding (3) requires knowing other things. Perhaps a hearer who understands the utterance has to know that the speaker intends to refer to different

⁹ Since they start knowing that 'Mercury' refers to Mercury and 'Venus' refers to Venus, they can combine this knowledge with KCAC to come to know that 'Mercury' and 'Venus' refers to different things if they refer at all. Since the agent knows that the terms refer, he can deduce that that they refer to different things.

things, or that it is presupposed that there are different things. But this is still not the same as knowing that terms actually refer to different things if they refer at all.

3.1 Lasnik and Evans

It may be thought that restrictions from binding theory provide examples of de jure anti-coreference. The structure of the following utterances pushes an anti-coreferential reading for the italicized occurrences:

(4) *He* is happy when *Oscar* is in love

(5) *Oscar* is happy when *Oscar* is in love

Howard Lasnik (1976) proposed a rule of anti-coreference to account for these and related cases:

(Non CR) If NP1 precedes and c-commands NP2 and NP2 is not a pronoun, then NP1 and NP2 are disjoint in reference.

According to Lasnik's proposal, certain grammatical structures require anti coreference. Assuming that agents know rules like (Non CR), then they can come to know that there is anti coreference in (4) and (5). If this line of reasoning is correct, we would have good

reason to think that de jure anti-coreference is a robust and interesting natural language phenomenon.

Gareth Evans (1985a) objected to Lasnik's proposal. He provided the following counter-examples to (Non CR):

(6) Everyone has finally realized that *Oscar* is incompetent. Even *he* has realized that *Oscar* is incompetent.

(7) Everyone has finally realized that *Oscar* is incompetent. Even *Oscar* has realized that *Oscar* is incompetent.

Evans notes that in these constructions, the second sentences have the same structure as (4) and (5) respectively. However, anti coreference is not required. Relatedly, true identity statements like 'hesperus is phosphorus' also provide counter-examples to (Non-CR).¹⁰ In my view, these examples are decisive.¹¹

3.2 Refinements

Consider a different case

¹⁰ Lasnik also considers weakening (non CR) so that instead of prohibiting coreference, it prohibits intended coreference. The examples (6) and (7) (and true identity statements) are also counter-examples to this proposal.

¹¹ Gareth Evans opts for a modified rule which prohibits referential dependence in certain constructions (as opposed to prohibiting coreference). Another option for formulating the rule prohibits de jure coreference (which is different from requiring anti coreference).

(8) *Oscar and his brother* went to the park.

Arguably, competent speakers who understand (8) will know that if ‘Oscar’ and ‘his brother’ refer at all, then they must refer to different things. (8) therefore passes the test. Does this then show that de jure anti-coreference is a genuine phenomenon? Not quite, KCAC only gives us (approximate) necessary conditions for de jure anti-coreference. But it does not give us sufficient conditions.

One way to resist the idea that (8) exhibits genuine de jure anti-coreference is by adding a further condition to de jure anti-coreference: that knowledge of conditional anti-coreference has to be genuine semantic knowledge. It cannot arise from “worldly” or extra-semantic knowledge. In (8), knowledge arises from our general understanding of family relationships, not semantics.

The same issue arises concerning de jure coreference. Consider the following:

(9) $1+1=2$

Suppose that everyone who grasps (9) is in a position to know that ‘ $1+1$ ’ and ‘2’ corefer. We should still not be convinced that the terms corefer de jure. This is a case where coreference is obvious, but still extra-semantic. Obvious coreference is not de jure coreference.

The reason we just rehearsed relies on making fine distinctions between semantic and extra-semantic knowledge. But the point can be made without attempting

to adjudicate on such cases. Notice that examples (8) and (9) involve complex terms.

We can restrict the phenomenon accordingly:

(Asymmetry) Regarding simple terms (like proper names), de jure coreference is a pervasive phenomenon. In sharp contrast, de jure anti-coreference is rare.¹²

How should we account for Asymmetry?

4. Fregean accounts of Asymmetry

According to Fregean semantics, sense and reference are the two main components of meaning. Senses play at least two roles for Frege. First, senses determine reference. A symbol refers to an object in virtue of it expressing a sense which refers to that object. Second, senses are the cognitive content of symbols. Cognitive contents are ways of thinking of referents. Accordingly, they are more fine-grained. For example, two symbols that refer to the same thing may express different senses (as with 'hesperus' and 'phosphorus').¹³

On this conception, coreference can happen in two ways. It can happen when the senses are the same, and it can happen when senses are different but turn out to refer to the same object. The Fregean can say that this distinction corresponds to the de

¹² I do not want to say that there is no de jure anti-coreference for simple terms at all. It may be that '7' and '6' exhibit de jure anti-coreference (and they may be simple terms). Nonetheless, I hold that these are relatively rare cases.

¹³ Frege (2000)

jure/de facto coreference distinction. In particular, restricting the claim to referring senses, they can say that representations A and B de jure corefer just in case they express the same sense.

The epistemic property of de jure coreference, KCC, can be explained by appealing to a general feature of senses: Transparency.

(Transparency) For an agent that grasps representations X and Y, if X and Y express the same sense, then the agent will know that X and Y express the same sense. And if X and Y express distinct senses, then the agent will know that X and Y express distinct senses.

Transparency can be seen as a general requirement of meaning (senses, for our purposes) which Michael Dummett (1978) claims is “undeniable”:

It is an undeniable feature of the notion of meaning—obscure as that notion is—that meaning is transparent in the sense that, if someone attaches a meaning to each of two words, he must know whether these meanings are the same (131)

It is crucial to note that Dummett appeals to knowledge in his definition of transparency and not something weaker like ‘justified belief’ or ‘warrant’. This is precisely what the Fregean needs to explain KCC.

There is a venerable tradition among philosophers of taking transparency to have somewhat of an axiomatic status. In a discussion of Dummett's passage, Paul Boghossian (2008) says that Dummett is following the footsteps of Russell and Frege 'both of whom gave epistemic transparency a pivotal, if unargued, role in their respective theories of linguistic and mental content.' (pg. 159). In addition, it is well known that the Fregean has made much of Transparency as a key element of rationality. The reason why it should be irrational for someone to hold that 'hesperus is a planet and hesperus is not a planet' as opposed to 'hesperus is a planet and phosphorus is not a planet' is that in the former, the senses of the 'hesperus' occurrences are the same and this is something that the speaker is able to recognize from the inside without further empirical investigation.

Assuming that transparency is a principled notion, the Fregean has a plausible explanation of the asymmetry between de jure coreference and anti-coreference. The pervasiveness of de jure coreference follows from the fact that when wishing to refer to the same thing, agents will often re-token a sense. When a sense is repeated, competent agents will know this (by transparency), and hence come to know that the same object is being represented (if anything is being represented at all).¹⁴

What about de jure anti-coreference? The paucity of de jure anti-coreference is explained by noting that there are no commonly instantiated relations between tokened senses that are both (i) transparent and (ii) guarantee distinctness of reference (when there is reference). For example, cases in which distinct senses are tokened will

¹⁴ We also need to assume that competent agents will be in a position to know senses (for singular terms) refer uniquely.

certainly be recognized as such (because of transparency), but will not normally put agents in a position to know whether or not the referents are the same.

These thoughts suggest that the asymmetry between de jure coreference and anti-coreference depend on a different asymmetry. This is an asymmetry between senses (of singular terms) and their designations: Each such sense can designate at most one object, but each object can be designated by more than one sense. Hence, a re-tokening of the same sense guarantees the same referent, but multiple tokens of distinct senses fail to guarantee distinct referents. Moreover, since these alleged facts about senses are part of our understanding, these facts about guarantees are epistemically accessible to competent agents.

So far, we have been trying to give a Fregean account of Asymmetry. We can resist it by resisting the Fregean explanation of de jure coreference. We can do this in at least three ways. First, we can question whether Fregean senses exist and give the semantic content of singular terms.¹⁵ Second, granting that senses do play that role, we can further question whether senses are transparent, or at least challenge the Fregean to give us a deeper explanation of why transparency applies to certain relations between senses (like identity and distinctness) but not others. Third, even under the assumption that senses are the contents of singular terms and are transparent, we can question whether de jure coreference can even be explained by appealing to senses. I will challenge the Fregean on this score. I have argued elsewhere (2011) that de jure coreference is not a transitive intra-discourse relation. If this is right, then since identity

¹⁵ Kripke (1980)

is transitive, de jure coreference cannot be explained by appealing to the identity of senses. Consider the following:

(10) We were debating whether to investigate both *Hesperus* and *Phosphorus*; but when we got evidence of their true identity, we immediately sent probes *there*.

Here, 'Hesperus' and 'Phosphorus' are de facto coreferential but each is de jure coreferential to 'there'. Hence the failure of transitivity.

In the next sections, we will see how other types of commonly found "dynamic" cases, put pressure on the Fregean. I will argue that the Fregean approach must be seriously modified and, ultimately, must be supplemented with a relational approach to meaning.

5. Semantic Relationism

We just saw how some Fregeans would wish to explain de jure coreference. She would say that A and B de jure corefer whenever A expresses a sense S, B expresses a sense S' and $S=S'$. Crucially, the semantic facts which explain de jure coreference simply concern the meaning of A on its own and the meaning of B on its own. The transparency of senses then guarantees the epistemic features of de jure coreference including KCC.

The semantic relationist denies this. Some further relational semantic fact is needed. There are at least two varieties of relationism. According to Pinillos (2006, 2011), this further semantic fact includes the holding of a semantically primitive relation, p-linking, between A and B. Kit Fine (2003, 2007) does not accept primitivism of this stripe. Instead, he holds that what it is for A and B to be de jure coreferential is that it is a *semantic requirement* that A and B refer to the same thing. The sense in which this is a relationist view is that the following inference is invalid:

(11) It is a semantic requirement that: A refers to *o*

(12) It is a semantic requirement that: B refers to *o*

(13) It is a semantic requirement that there is an *x* such that A and B refer to *x*.

That is, the notion of a semantic requirement is not closed under logical consequence.

The idea here is then that the non-relational semantic requirement about A and B in (11) and (12) are not sufficient to explain the relational requirement in (13).¹⁶ Hence, from the perspective of a semantic theory, there will be irreducible relational semantic requirements. In this sense, the theory will be relationist.¹⁷

¹⁶ I say the semantic fact that A refers to *o* is non-relationist because it doesn't (typically) relate two or more representations. Of course, reference is a relation nonetheless, but it is not a relation in the intended sense.

¹⁷ I am also a relationist so I agree that (13) does not follow from (11) and (12). Where I disagree with Fine is that I think that (13) is explained in part by a primitive semantic relation holding between A and B.

On Fine's account, some relational semantic requirements cannot be explained by appealing to non-relational semantic requirements. But the relational semantic requirements need not just concern coreference. For example, according to Fine, it is a semantic requirement that 'even prime' denotes the intersection of the denotation of 'even' and 'prime'. But this is not reducible to non-relational semantic requirements.

Elsewhere, I have argued against Fine's account in favor of primitivism (Pinillos 2015). I do not want to rehearse the arguments here. Rather, I want to notice an important shortcoming that both relationist theories face. They both have a very difficult time accounting for Asymmetry.

In my account, de jure coreference is explained by appealing to a primitive semantic relation. The absence of de jure anti coreference would presumably be accounted for by appealing to the absence of a primitive anti-coreference relation. But there is now a mystery about why there should be this absence. Why shouldn't there also be a primitive "anti link" ensuring non-coreference just as there is one guaranteed coreference? Similarly, on Kit Fine's approach, A and B are de jure coreferential whenever it is a semantic requirement that there is an x such that A and B refer to x. But it is now a mystery why there shouldn't be anti-coreference requirements? Why shouldn't there be semantic requirements (in natural languages, thoughts or familiar formal languages) that there are distinct x and y such that A refers to x and B refers to y? Again, this absence is mysterious.

The Fregean, however, seems to have an explanation for Asymmetry. This counts in favor of the Fregean and against the relationist. But the matter is actually

much more complicated. Dynamic coordination cannot be explained by the Fregean, at least not without importing relationist machinery.

6. Mental Files

Francois Recanati (2012) has recently argued that mental files can play the role of senses. Mental files are posited to be stable yet dynamic mental representations that bundle information about objects in the world. This brief characterization contains two important elements. First, mental files are dynamic in that information may be added and deleted as we learn more about the objects they are about. This encoding of information is what allows mental files to explain cognitive content.¹⁸ Second, mental files have semantic properties. They can refer to things in the world. Recanati, unlike Frege, eschews descriptivism. The referent of a mental file is not normally determined, for example, by selecting the object which satisfies most or all of the information included in the file. The referent is often determined in some other way. For example, it can be determined by some causal-historical relation connecting the mental file (or the information therein) and the object in question. In this way, Recanati is able to avoid familiar Kripkean objections to Fregean approaches to content.

Importantly, it is possible for an agent to have two distinct files for a single object. The most obvious ways this can happen is if the agent is in a “Frege” case and thinks of an object in two distinct ways.

¹⁸ To simplify matters, we can identify the information contained in a file with a predicate in system of mental representations. When a predicate that expresses a property P is contained in a file about an object, the agent believes of that object that it has property P.

Now what is the connection between mental files and de jure coreference? Some philosophers have thought there is an intimate connection between these two notions. Laura Schroeter (2007), for example, treats the connection as definitional.

For present purposes, we can simply define mental files as whatever internal cognitive mechanisms normally explain the subjective appearance of de jure coreference (pg. 600).

I do not think that mental files are definitionally connected with de jure coreference (or their subjective appearance) in this way. I think of mental files as theoretical mental posits which explain how we organize information and how we are able to think about things in the world. In my view, it is substantive empirical question, not a definitional one, whether the subjective appearance of de jure coreference can be explained by appealing to mental files.

Krista Lawlor and Kit Fine dismiss mental files as mere metaphors which lack explanatory power. Here's Lawlor (2010):

Nor will it do to say that for thoughts to represent their objects as the same [de jure coreference] is for the thoughts to somehow involve the same file of information; invoking a file isn't a way of giving an account as much as a metaphor that presumes the very notion to be explained (page 491).¹⁹

¹⁹ See also Fine (2007, 67-68)

There are two criticisms here that should be separated. One is that the notion of a file assumes what needs to be explained, and the other is the idea that a mental file is a metaphor and so, presumably, lacks explanatory power. Regarding the former, the argument I give below should be distinguished from this type of criticism. Although my argument aims to show that appealing to mental files is not sufficient to account for de jure coreference, I do not also say that mental files presuppose de jure coreference.

Regarding the second worry, I do not think that uses of 'mental files' need be metaphorical. When we say that computers 'have files' we are using the notion in its literal sense. The same applies to files in the mind. 'files', in this sense, is a functional term (not one that must denote a contiguous physical location).

6.1 Mental Files and Coordination.

Let us shift from de jure coreference to coordination. Recall that by coordination, I mean to include all the cases of de jure coreference plus ones that are just like them except where reference has failed. This change won't affect the test/necessary condition (KCC) we discussed above. This is because KCC is still an approximate necessary condition/test for coordination even for non-referring cases.

Following the Fregean account of de jure coreference discussed above, the natural explanation of coordination in terms of mental files says that two representations

are coordinated just in case they are associated with the same mental file.²⁰ Francois Recanati, in his *Mental Files* book agrees with one direction of this bi-conditional. Although he thinks that mental files are central to explaining coordination, and (as I read him) thinks sameness of mental files is sufficient for coordination, he does not think it's necessary. He thinks that coordination can also occur when files are merged. Recanati seems to accept the following pair of sufficient conditions for coordination:

Identity: If representations X and Y are associated with the same mental file, then they are coordinated.

Merge: If representation X is associated with file X*, Y with file Y* and Y* is the product of the merging of file X* with some file, then X and Y are coordinated.

The reason why Recanati adopts Merge is to account for one of the cases discussed above:²¹

²⁰ Below, we will speak of mental files and mental file stages themselves as being coordinated. Since files are not natural language representations, then it is unclear how our test KCC (which appeals to discourses) could have application in this domain. The test can be extended naturally in the following way: Files or file stages X and Y in a thought are coordinated whenever their natural language counterparts in a natural language expression of the thought are coordinated. This is not an endorsement of the priority of natural language over thought, however.

²¹ Recanati (2012, 111-112)

(10) We were debating whether to investigate both Hesperus and Phosphorus; but when we got evidence of their true identity, we immediately sent probes there.

Recall that 'there' is allegedly coordinated with 'Hesperus' and 'Phosphorus' ('Hesperus' and 'Phosphorus' are not themselves coordinated). But the file associated with 'there' is not the same file associated with just either 'Hesperus' or 'Phosphorus'. Rather, 'there' is associated with a file that is the product of the merger of the 'Hesperus' and 'Phosphorus' files.

Insofar as the Fregean finds cases like (10) convincing, then the account of coordination that appeals to sameness of files will have to be amended to include something like Merge. It is worth pausing to note that this is a significant move for the Fregean. Coordination is now no longer reduced to identity of sense. Moreover, Dummett's Transparency will have to be amended to cover more than just identity of meaning.

I would like to find a unified Fregean approach to coordination. Eventually, we will find a single condition to cover both Identity and Merge above. The first step, however, is to replace talk of files with file stages (which are time-slices of files). This will give us more theoretical leeway.

To get us started, here are two initial proposals for sufficiency conditions which replace the above conditions with conditions concerning stages. Note that the first, No

Change (provisional), is not a full replacement of Identity. Rather, it gives us that condition at the limit, where it is clear that the condition must hold.

No Change (provisional): If representations X and Y are associated with mental files stages that contain the exact same information, then X and Y are coordinated.

Merge for Stages (provisional): If representations X and Y are associated with mental file stages X* and Y* respectively, and Y* has all the information from file stage X* plus information from some other file stage (adjusted for consistency and coherence), then X and Y are coordinated.

Consider the following putative counter-example to No Change. Suppose there are two mental file stages with the same information, but these file stages are either in different persons or accidentally in the same person (perhaps separated by decades). Arguably, No Change mistakenly predicts coordination.

What is missing from the accounts is the idea that mental file stages can be *connected*. Here is where the analogy of physical files come to play. There could be two distinct files containing the exact same information but located in different offices. From the perspective of someone utilizing the files, there is an important difference between these files. For example, a person who adds information to one of them will search for it

later in that location and not the other. Phrasing the idea in terms of file stages, the person will attempt to retrieve the information from a file stage that is physically connected to the file stage where she placed the information in the first place.

Similarly, mental file theory must assume that there is a theoretically important difference between mental files stages that are connected and those that are not. In the case of physical files, connectedness boils down to space-time contiguity. In the cases of mental files, connectedness is presumably not also explained in terms of space time contiguity. Some other type of contiguity is in play. In what follows, I will take on board the notions of connectedness and contiguity that are important mental file theorizing. I will assume also that we have some intuitive grasp of these notion even if they are “theoretician’s intuitions”.

We can now isolate three important relations defined over stages: “Same File”, “Connectedness” and “Coordination”.

- (i) The ‘Same File’ relation: Stages X and Y are related by the Same File relation just case they belong to the same file.²²
- (ii) The ‘Connected’ relation: Stages X and Y are related by the Connected relation just in case there is a contiguous series of file stages connecting X and Y.

²² I am assuming that files can survive change (two file stages which contain different information can be related by “Same File”).

(iii) The 'Coordination' relation: Stages X and Y are related by the Coordination relation just in case the thought tokens associated with X and Y are coordinated.

There are various views one could have about how these relations themselves relate to one another other.²³ My own position, as we will see, is that these three relations are all distinct. An extreme view in the opposite direction identifies all three relations. This position is implausible once we accept merging. In merging, two distinct files convert into a single file which presumably is yet a third numerically distinct file (it cannot be identical to both initial files because identity is transitive and symmetric). However, all three files are connected. This shows that 'Connectedness' and 'Same File' are distinct relations.

Another reason to think that Connectedness is distinct from Same File is this. Suppose you visually track an object X starting at t_1 as it slowly morphs into a numerically distinct object Y. The tracking is accompanied by a running commentary about whatever object is in front of you at the time. At some point in this process, your words and accompanying thoughts will be about Y. Let this be time t_n . The series of thoughts involved in this process invoke connected mental files stages since at no point in the series do you intentionally open a new file. Rather at each point in the series, you are taking the previous file stage and simply updating it with new information. So

²³ In discussing how these relations relate to each other I suppose that the domain of relata are fixed to mental file stages to all three relations.

although the corresponding file stages at t_1 and t_n are connected, they must be associated with distinct files (because they refer to different objects).²⁴ Connectedness and Same File are therefore not the same relation.

We can now rewrite the conditions above taking into account Connectedness:

No Change: If representations X and Y are associated with mental file stages that contain the exact same information and these stages are Connected, then X and Y are coordinated.

Merge for Stages: If representations X and Y are associated with Connected mental file stages X^* and Y^* respectively, and Y^* has all the information from file stage X^* plus information from some other file stage (adjusted for consistency and coherence), then X and Y are coordinated.

6.2 Dynamic Coordination

We should add to our list of sufficient conditions for coordination. I think one could deploy a file, modify it, and redeploy it in such a way that the original and the redeployed file are coordinated. Suppose I see a cat and say 'this cat is grey'. I then

²⁴ Note that if we insist that the stages at t_1 and t_n are stages of the same file and yet refer to different things, then we give up the idea that sameness of file is even a sufficient condition for coordination. I assume this is non-negotiable.

touch the cat and judge 'she is soft'. Here, 'this cat' and 'she' are coordinated. Yet, the associated file has undergone change. The file for that cat has added information to the effect that she is soft. Consider a different case from Frege (1956) and Evans (1985b). Suppose I think of a particular time t as 'now'. Later on I think of that same time t as 'then'. These representations are coordinated yet the later one contains a new perspective, and is thus associated with a slightly modified file. No Change and Merge for Stages fail to capture this.

Let us call these types of cases, cases of dynamic coordination.²⁵ The following principle is an attempt to explain cases of dynamic coordination in terms of mental files:

Dynamic (provisional): If representations X and Y are associated with Connected mental file stages X^* and Y^* where the information in X^* and Y^* substantially overlap then X and Y are coordinated.

There are some points about Dynamic (provisional) that I want to bring out right away. First, No Change and Merge for Stages are arguably derivable from it, so it gives us a unified explanation of the phenomenon. Second, there is no mention of mental files, per se. To be specific, there is no mention of the Same File relation. Hence, insofar as this principle is plausible, the role of mental files in explaining coordination is minimized.

²⁵ Note that if one thinks that sameness of files is sufficient for coordination (and also think that files can survive some change of information contained), then one will also think that dynamic coordination is possible.

According to Dynamic (provisional), when two file stages are connected and exhibit substantial overlap in information (call this 'Overlap'), then coordination takes place. As I mentioned before, Connectedness on its own is not sufficient for coordination (two distinct files referring to different things could be connected).

Overlap on its own is also not sufficient. There could be files about distinct objects with overlapping information. This is because reference is not determined in a descriptivist way. Rather, reference is fixed by how the predicates that constitute the information got into the file. Recanati adopts Evans' conception of reference (for the types of files at issue) where the referent of a term is the dominant source of the information associated with the term. The dominant source of the information is not the object which descriptively satisfies the information. Rather, it is the object that causally brought about that information. With respect to files, this means that the referent of the file is the dominant source of the predicates, as they appear in the file. In principle, there could be files possessing the same information (same predicates) but referring to different things.

With this idea as background, we can develop a counter-example to Dynamic (provisional). Suppose an agent opens a file F at time t_1 with information encoded with predicates A , B and C . The source of this information, as they appear in the file stage, is object X . As time goes by, the agent modifies the file so that B and C are deleted. Hence, the file stage at t_2 only contains predicate A . Now at time t_3 , predicates B and C are reinstated, however, the source of these predicates, as they appear in the file stage,

is some distinct object Y. So the file stage at t_3 contains information A, B, and C but the dominant source of the information is object Y which is distinct from X.

We have then the following situation. The file stages at t_1 and t_3 are connected and contain overlapping information (in fact, they contain the same information—the same predicates) but they are about different objects. This means they are not coordinated. This is a counter-example to Dynamic (provisional).

It won't solve the problem to require that the predicates in the two file stages at t_1 and t_3 have the same object as the dominant source. This is because the source might be the same object, accidentally, as in a Frege case. What we need is some psychological mechanism where it is transparent that tokens of predicates have the same source. Let us call this relation between tokens of predicates 'transparency of source' or TS. Transparency of source is not the same relation as Coordination. If I open a file about Smith with the predicates 'is happy' and 'is Canadian', tokens of these predicates in the file are related by TS but they are not coordinated. They are not coordinated because the predicates do not even mean the same thing. Relatedly, two tokens of the predicate 'is happy' in separate mental files will have the same meaning but may have different sources.

The notion of transparency here is naturally seen as epistemic. It is tempting to explain it in terms of knowledge as Dummett does. We can be more cautious here and simply require the agent to have access to sameness of source and that this access play a role in explaining coordination and the accompanying knowledge of (conditional) coreference.

It is unclear whether TS is a genuine semantic relation. The notion of “source” is not the familiar reference relation. First, reference can compose, but “source” does not.²⁶ Second, as I mentioned before, the semantics of predicates is wholly distinct from the notion of the source of a predicate (in a file). It is preferable, therefore, to think of “source” as a proto-semantic notion. Let us now rewrite Dynamic.

Dynamic: If representations X and Y are associated with mental file stages X* and Y* respectively where most of the information in X* and most of it in Y* are TS related, then X and Y are coordinated.²⁷

Dynamic gives us a plausible unified implementation of coordination. Three points are in order. First, neither Same File nor Connectedness make an appearance. Thus, mental files do not play the role in coordination that the Fregean was supposing. Second, it can be a vague matter whether most of the information across two stages are related via TS (on account of the vagueness of ‘most’). This leads to vagueness and accompanying uncertainty about whether there is coordination in certain dynamic cases. Third, a

²⁶ The referent of ‘the father of bob’ is derived from the referent of ‘the father of’ (a function) and the referent of ‘bob’ in a familiar way. But the source of ‘the father of bob’ is not compositionally derived from the source of the component representations. It is not even clear what it means to ask for the source of a complex expression where the parts have different sources.

²⁷ To be clearer, most of the predicates in X* must be TS related to some predicates in Y* and most of the predicates in Y* must be TS related to some predicates in X*. When this happens, coordination is instantiated. This is just a simplified implementation of Evans’ conception because it assumes that what counts as the dominant source is just the source of most of the predicates in the file stage. As Evans (1973) himself pointed out, the picture is more complicated. The coordination account needs to be modified accordingly.

question can arise as to whether Dynamic gives a reduction of coordination. I now turn to this.

I have argued elsewhere that coordination is a primitive semantic relation (it cannot be reduced to other semantic notions). Because other implementations seem possible, Dynamic does not give us a reductive account of coordination (it just gives us an implementation). But let us suppose that I am mistaken about this, and that Dynamic does give us a reduction of coordination. A further question is whether Dynamic reduces coordination to only non-semantic notions. Suppose that it does. In particular suppose that TS in Dynamic is not a semantic relation. In this case, the result would be compatible with my earlier claim that coordination is irreducible to other semantic notions.

Suppose on the other hand that the key notion in Dynamic, TS, is a genuine semantic notion. In that case, coordination would be reduced to a distinct semantic relation TS. But it is open that that this relation is itself irreducible to other (non-relational) semantic notions. Hence relationism would not be refuted (even if coordination gets a semantic reduction as specified in Dynamic).

6.2.1 Why not just revert to The Same File relation?

In the previous section, I developed a counter-example to Dynamic (Provisional). One way to solve the problem is to choose simplicity and reinstate the Same File relation (instead of the TS relation):

Dynamic (Same File): If representations X and Y are associated with mental file stages X* and Y* related by Same File, where the information in X* and Y* substantially overlap, then X and Y are coordinated.

We should reject this alternative. There are two related problems. First, as we saw earlier, the merging phenomenon reveals coordination among file stages belonging to distinct files. The revised condition we are now considering cannot account for that phenomenon. We would need an additional principle, thereby thwarting our attempt at a unified explanation.

The second worry with this approach is instructive. It gets the order of explanation wrong. We saw above that the referent of a mental file stage is the dominant source of the information contained in the file at the time. In dynamic reasoning, information is added or deleted from a file. This means that the dominant source of information could suddenly change, thereby obliterating coordination. One way of ruling out this possibility is to add in Dynamic (Provisional) that the source of information must (mostly) remain the same. This is the approach I favored in Dynamic (invoking TS). But we are now considering a different strategy. According to this strategy, the possibility is ruled out by saying instead that the file must remain the same. Given the account of how reference is fixed for files, it seems like the ultimate explanation for why the referent has not shifted (in coordination) is that the source of information has remained the same. It is not that the file has been preserved. In fact, it

seems like part of the reason why a file is preserved is that the source of information has mostly remained the same, and not the other way around. Hence, a better of explanation for coordination appeals to Transparency of Source (TS) and not sameness of file. Therefore, we should prefer Dynamic over Dynamic (Same File). The approach is more complicated but philosophically more satisfying.

6.3 Resisting Dynamic Coordination

Philosophers have worried quite a bit about dynamic coordination.²⁸ The structure of the problem is this. Two token thoughts may be intuitively coordinated and yet these thought tokens constitute two different ways of thinking about the object. This creates a problem for the Fregean who holds that coordination just boils down to thinking about the object in the same way (deploying the same sense).

There are many attempted resolutions to this problem. My goal in this paper was not to solve this particular problem for the Fregean. But some remarks are in order. One solution is to say that the two ways of thinking about the object are really a single way of keeping track of the same object across time. This is how Evans (1985b), following Frege (1956) attempts to solve the problem. My view is compatible with this perspective. I acknowledge that many (but not all) cases of dynamic coordination will involve the same file. However, I have argued that a deeper explanation of coordination makes no appeal to sameness of file. One must adopt relationism.

²⁸ Burge (1988), Campbell (1987), Evans (1985b), Frege (1956), Papineau (2006), Schroeter (2007)

A second approach, which is compatible with the first option, is to weaken the notion of coordination so that there is less of a guarantee of coreference. This view corresponds to Schroeter and Campbell's position that senses are not transparent. My view on transparency is that knowledge of conditional coreference must be respected (or an analogous notion at the level of thought). Although this is an epistemically strong notion, I do not think it requires Cartesian certainty, incorrigibility or anything like that.

A third option puts pressure on the notion of a Fregean sense or a mode of presentation (Millikan, 1997). In this paper, I have assumed that the Fregean approach is on the right track. But this is not a requirement. In fact, Kit Fine holds that relationism can solve the puzzles that gripped Frege without appealing to senses.

A fourth option is to accept that there is only static coordination in reasoning. So when a sense/file changes, there cannot be coordination. Recanati has argued in favor of this approach.²⁹ Consider the case above where an agent first thinks 'This cat is grey', later thinks (upon touching her) 'she is soft' and then even later concludes 'there is a soft grey cat'. The intuitive thing to say is that the file has changed in the process. But if so, Recanati (2015) says there is no coordination. There is a related inference, however, in which there is coordination. With this inference, at the time our agent wants to draw the conclusion, he re-tokens the premises and conclusion all at once. And when the agents does this, he tokens the same unchanged file throughout. Recanati supports this stance by saying that 'I deny that legitimacy of that 'diachronic' understanding of

²⁹ Ninan (2014) and Onofri (2014) separately press Recanati on this issue. Onofri, in particular, claims that if file identity does not explain coordination then the role of mental files in explaining coordination is curtailed. Although I think that this criticism is in the right track, it is too quick. Other notions, central to files,

that...train of thought. For logic purposes, a train of thought has to be thought of as synchronic' (page 7).

Although this is not the place to discuss the foundations of logic, I believe that it is hasty to say logical notions such as validity fail to apply to dynamic real-time reasoning. I grant that the classical notion of validity does not explicitly concern itself with diachronic dynamic reasoning, but it is hasty to conclude that it concerns itself only with synchronic reasoning. Another option is that validity is neither synchronic nor diachronic. It is an *a temporal* notion.

What is central to logical validity and related notions is the notion of logical form. Dynamic reasoning can exhibit form. But this form is not a pattern formed by the fluid sense/file structure invoked to dynamic reasoning. Rather, the form is realized by coordination itself (which the relationist thinks is a further aspect of meaning).

Dynamic thought is puzzling for the Fregean because she expects a single theoretical entity, a Fregean sense, to account for two phenomena: modes of presentation and coordination. I have urged that the Fregean should reject this assumption.

7. Revisiting Asymmetry

We saw above that the Fregean has an elegant account of Asymmetry (the prevalence of de jure coreference and the scarcity of de jure anti-coreference). This explanation depended on the analysis of coordination in terms of file identity. But we rejected this analysis. Our new approach is expressed in Dynamic. Can this new approach explain Asymmetry?

The new approach incorporates the TS relation (Transparency of Source) and thereby inherits the trouble relationism faced in explaining Asymmetry. The problem there was to explain why semantically required coreference but not anti-coreference should be prevalent. The problem has now been converted to explain why TS but not “transparency of distinct sources” should be prevalent. Why shouldn’t it also be commonly transparent that two predicate tokens have distinct sources?

Here’s an explanation sketch for the asymmetry. One way that two mental representation tokens A and B in an agent can come from the same source is if they are part of the same causal chain originating from that source. The fact that B is part of the same causal chain as A (A partially causes B) is a fact that is grounded on the local features of the agent. Now if this causal fact is what guarantees sameness of source (as the Evansian picture suggests) then the fact’s locality can explain why the agent is at least in a position to apprehend it. Hence we have a clue about how transparency of source is possible.

In contrast, whatever grounds the fact that two tokens C and D come from distinct sources is something not local to the agent. The ground is not that C and D belong to distinct chains. That fact does not guarantee distinctness of source (it could be a Hesperus/Phosphorus type of case). Rather, the responsible ground presumably just concerns the origins of the chains of C and D. This is something that perhaps happened long ago and in a distant place. Hence the agent will have a different type of access to that fact.

In sum, the causal structure of information we get from Evans suggests that some facts which guarantee sameness of source (for mental tokens instantiated in a single agent) are local. In contrast, facts that guarantee distinctness of source are not. I claimed that this asymmetry can help explain a different asymmetry, one we noted between de jure coreference and anti-coreference.

Bibliography

Boghossian, Paul (2008) *Content and Justification: Philosophical Papers*, Oxford University Press.

Burge, Tyler (1988). Individualism and self-knowledge. *Journal of Philosophy* 85 (November):649-63.

Campbell, John (1987). 'Is sense transparent?' *Proceedings of the Aristotelian Society* 88:273-292.

Dummett, Michael (1978) 'Truth and Other Enigmas' Harvard University Press
Evans, Gareth 1973: 'The Causal Theory of Names'. *Proceedings of the Aristotelian Society, Supplementary Volume 47*: 187 – 208. Reprinted in *Collected Papers* (Oxford University Press: 1985), page references to reprinted version.

Evans, Gareth (1985a). 'Pronouns' in *Collected Papers*. Oxford University

Evans, Gareth (1985b). 'Understanding Demonstratives' in *Collected Papers*. Oxford University Press

Fiengo, Robert. & May, Robert. (1994). *Indices and Identity*. The MIT Press

Fine, Kit (2003). The Role of Variables. *Journal of Philosophy* 100 (12):605 - 631.

Fine, Kit (2007). *Semantic Relationism* Oxford University Press

Frege, Gottlob (1956). 'The Thought' in *Mind* (65)287-311

- Frege, Gottlob (2000). Ueber sinn und bedeutung (on sense and reference). In Robert J. Stainton (ed.), *Philosophical Review*. Broadview Press. 45
- Kripke, Saul (1980). *Naming and Necessity*. Harvard University Press.
- Larson, Richard & Segal, Gabriel (1995). *Knowledge of Meaning*. The Mit Press.
- Lasnik, Howard. (1976) 'Remarks On Coreference' *Linguistic Analysis* (2) 1-22
- Lawlor, Krista (2010). Varieties of Coreference. *Philosophy, Philosophy and Phenomenological Research* 81 (2):485-495
- Millikan, Ruth (1997). Images of Identity: In Search of Modes of Presentation. *Mind* 106
- Ninan, Dilip (2014). On Recanati's Mental Files. *Inquiry* 58 (4):368-377.
- Onofri, Andrea (2014). Mental Files and Rational Inferences. *Inquiry* (4):1-15.
- Papineau, David (2006). Phenomenal and perceptual concepts. In Torin Alter & Sven Walter (eds.), *Phenomenal Concepts and Phenomenal Knowledge: New Essays on Consciousness and Physicalism*. Oxford University Press. 111--144.
- Pinillos, N.A. (2006) On Referring to the Same Thing. PHD Dissertation (Rutgers University).
- Pinillos, N.A. (2011) Coreference and Meaning. *Philosophical Studies* 154(2)
- Pinillos, N.A. (2015). 'Millianism, Relationism and Attitude Ascriptions' in *On Reference* Ed. Andrea Bianchi Oxford University Press
- Recanati, Francois (2012). *Mental Files*. Oxford University Press
- Recanati, Francois. (2015) Cognitive Dynamics. To appear in K. Korta and M. Ponte (eds.) *Reference and Representation in Language and Thought*.
- Schroeter, Laura (2007). Illusion of transparency. *Australasian Journal of Philosophy* 85 (4):597 – 618.
- Williamson, Timothy (2000). *Knowledge and Its Limits*. Oxford University Press

