

Discourse Content

Sam Cumming

University of California, Los Angeles

October 10, 2008

The notion of a discourse entity is fundamental to work on the semantics of discourse. Discourse entities are components of the discourse context rather than elements of the (concrete) world, and keeping track of them (i.e. determining the anaphoric dependencies in a text) is a necessary step in the comprehension of discourse.

In this paper, I propose to analyse intentionality as a relation to *discourse content* (i.e. content built out of discourse entities), rather than truth-conditional or referential content. As building blocks of content, discourse entities have distinct advantages over concrete referents. First, there can be discourse “reference” (i.e. anaphoric dependency) without reference; expressions that fail to refer may nevertheless be contentful. Second, two distinct discourse entities (corresponding to different patterns of anaphoric dependency) may be grounded in the same referent, and so substitution of co-referring expressions is not assumed to be content-preserving (cf. Frege, 1892). Finally, discourse reference, unlike reference, has a straightforward connection to comprehension: determining the discourse reference of a noun phrase – that is, resolving it anaphorically – is necessary and sufficient for understanding it.

Discourse semantics must be clarified and extended before it can provide a satisfactory basis for discriminating intentional states. First, discourse entities must be distinguished from mental symbols (Kamp, 1985; Asher, 1986). In this paper, I follow DeVault and Stone (2006) in treating discourse entities as objective social constructs. Second, the “lifespan” of discourse entities (Karttunen, 1976) – and so the range of anaphora – must be allowed to extend beyond the limits of a given discourse, to enable the transmission of content between discourses. I argue that the proper name is an example of a discourse-crossing anaphoric expression.

1 Discourse Entities

In a famous passage, Lauri Karttunen sets the stage for a denotational semantics of anaphora:

Let us say that the appearance of an indefinite noun phrase establishes a discourse referent just in case it justifies the occurrence of a coreferential pronoun or a definite noun phrase later in the text. (Karttunen, 1976, 366)

Karttunen’s ontology provides the means to encode anaphoric links by way of denotation. Noun phrases denote discourse entities (a.k.a. discourse referents, or drefs for short). If two noun phrases denote the same dref, then they are anaphorically linked: either one is the antecedent of the other, or both have the same antecedent. We can make a few general remarks about this denotational semantics. Every indefinite noun phrase serving as an antecedent will have a different denotation. Since cataphora is rare, we can expect an indefinite to denote a “new” dref (i.e. one that is not the denotation of any previous expression in the discourse) and a pronoun or other anaphoric expression to denote a dref that is “old” (see also Heim, 1982).

If two referring expressions denote the same dref, then they refer to the same object. Nevertheless, it is possible to distinguish denotation from reference. Firstly, because there can be anaphora (and so discourse entities) without reference.

(1) Bill saw a unicorn. It had a golden mane.

The expression ‘a unicorn’ in (1) above is the antecedent for the pronoun ‘it’. In Karttunen’s terminology, the indefinite establishes a dref which is retrieved by the pronoun. On a denotational semantics of anaphora such as I have sketched, both expressions denote the same dref. However, there is no reference here. Neither the indefinite ‘a unicorn’ nor the pronoun ‘it’ refers to anything.

Furthermore, expressions that co-refer do not necessarily denote the same dref (cf. Heim, 1983).

(2) Al went to work and Bob went to the bar. One of them took a gun from the drawer of the bureau and brought it with him.

The noun phrase ‘one of them’, while it contains the anaphoric expression ‘them’, is not itself anaphorically linked to either ‘Al’ or ‘Bob’. Instead, it

creates a new option for subsequent anaphora—it establishes a novel dref. The pronoun ‘him’ in the same sentence is anaphoric to ‘one of them’, rather than to either ‘Al’ or ‘Bob’, and so differs in denotation from both. Note that, due to the logic of the sentence, this pronoun (if it is understood as referring at all) must refer to either Al or Bob—whoever brought the gun. Thus, the denotation of ‘he’ in (2) cannot be the same as its reference.

Discourse entities are not the same thing as the concrete referents of noun phrases. So what are they? On the account I favour, they are socially-constructed abstract objects. David Lewis offers a view (nestled among some others in Lewis 1979) of the context of interpretation that is like this. He sees the context as an abstract representation – a scorecard – whose state is updated in a rule-governed way by acts of the conversational participants. He distinguishes this view from the view (held by Stalnaker [1978] and Clark [1996]) that the state of the context is derived from subjective mental states of the participants. DeVault and Stone (2006) compare the context to the abstract “chess-board” representing the state of play in a correspondence game of chess (in which the players do without a physical gameboard). This complex representation consists of parameters (or “variables”) set to different values. For instance, the parameter for who has the next move is initialized at *white* and toggles between players after each move. These parameters are analogous to individual discourse entities (and in some cases, the parameter values could be thought of as the concrete referents of those entities).¹

2 Discourse Reference and Comprehension

In the spirit of Turing’s writings on machine intelligence, Karttunen connects a vague concept – “comprehension” – with a precise engineering task: constructing a device that, if given a text, will “build a file” containing “records of all the individuals . . . mentioned in the text” (Karttunen, 1976). The device must recognize both when a new individual is meant and when the same individual is meant. In other words, it must resolve all anaphoric connections in the text.

In a different context, Kit Fine makes the same connection between comprehension and anaphoric coordination:

¹We need not assume that all the information in the context is even in principle available to the conversational participants. Indeed, the contention of this paper is that content is couched in terms of discourse entities rather than concrete individuals, which would preclude a particular parameter setting (in the sense above) from being a potential object of belief.

I take it that we all have some intuitive grip on this notion of coordination or *representing as the same*. But a good test of when an object is represented as the same is in terms of whether one might sensibly raise the question of whether it *is* the 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. 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 thereby betrays his lack of understanding of what you meant. (Fine, 2007, 40)

We might summarize this, in the terms of the preceding section, as follows: determining the denotation (i.e. discourse reference) of the expressions occurring in a text is necessary for textual comprehension.

By contrast it is not a requirement on comprehension that the interpreter recognize when the same individual is being referred to. For a long time, people did not know that the Morning Star was the Evening Star. Nevertheless, they understood each other perfectly well when speaking of those heavenly bodies.

The generative syntax tradition (most notably Chomsky), while criticizing the philosophical notion of *reference*, has made considerable use of *coreference* (understood, like discourse reference, not to imply true reference), most famously in Chomsky’s Binding Theory (Chomsky, 1982). While research in this tradition tends to be confined to isolated sentences, and while its substantial claims are about syntactic representation rather than communication, its generalizations connect recognitions of consistent patterns of “co-indexing” (the syntactic correlate of coreference) with linguistic competence.

3 Content

According to Brentano (1973, 88-89), encumbrance with a content (“aboutness” or “intentionality”) is the defining characteristic of the mental. More reasonably, it is only a portion of mental phenomena – e.g. what Russell called “propositional attitudes” – that are properly speaking contentful. The semantic level of explanation, at which mental states are characterized by their content, is widely considered indispensable to psychology and related fields (Marr, 1982; Newell, 1982; Fodor, 1975).

Linguistic utterances are a sophisticated means of manipulating the intentional states of others (Grice, 1957). In speech act theory, utterances are classified by illocutionary type (assertion, imperative, etc.) and content (Searle, 1975), which together determine the mental states necessary and sufficient for a felicitous, sincere utterance. (For instance, an assertion with content p is sincere iff the speaker believes p).

Philosophical discussion of content usually prescind from the details of speech act theory, and pronounces that a (declarative) sentence type encodes a content, conceived of as a package of information.

To begin with, pretheoretically it seems that sentences encode pieces of information and that distinct sentences may encode the same piece of information, as is perhaps the case with ‘Snow is white’, and ‘Schnee ist weiss’. (King, 2007, 1).

This suggests a further role for content.² If an utterance encodes a certain content, then it is natural to treat interpretation as *decoding* the utterance and recovering its content. It is standard to take the connection between content and comprehension as follows: one has understood an utterance just in case one has assigned the correct content to it.

Attitude reports are linguistic utterances that “close the circuit” of mental and linguistic content. The subordinate clause of an attitude report linguistically encodes a content that is attributed (in the report) to a cognitive state of the subject of the report. For instance, the sentence ‘John thinks that snow is white’ encodes the content that John is in a mental state (of belief) with the same content as the sentence ‘Snow is white’.

Just as the sentence is the primary notion in syntax, so the *proposition* (the content of a sentence) is the key unit in semantics.³ The contents of mental states are thought to be propositional (despite curtailed intentional attributions like ‘John wants a sloop’ and ‘Bill seeks a unicorn’). In the literature, there are two standard ways of conceiving of the content of a sentence: top-down and bottom-up.

The top-down thinker (or *truth-conditionalist*) conceives of the nature of propositions as follows. A sentential utterance is a complete representation; it represents the world as being a determinate way. In possible worlds talk, it determines a boolean value - true or false - at every possible world. Furthermore, the content of a representation consists in how it represents the

²For this role, see Frege 1892 and Burge 1977.

³I follow one standard interpretation of ‘semantics’ as pertaining to content (King and Stanley, 2005).

world. Two representations differ in content only if they differ in the way they represent the world as being (i.e. there is some possible world to which they assign opposite boolean values). The truth-conditionalist thus identifies propositions with conditions that apply to the world as a whole (or, more compactly, with the sets of possible worlds satisfying the conditions).

The *referentialist* about content is a bottom-up thinker.⁴ The referentialist takes what the parts of the sentence represent (i.e. refer to) and retains them as parts of the content of the sentence as a whole. The referentialist treats a proposition as a tree decorated with a semantic object (the referent of the corresponding subtree) at each of its nodes.⁵ Note that these “Russellian” propositions are strictly more fine-grained than sets of possible worlds.⁶ For two sentences to be a semantic match, according to the referentialist, they must match in referential content at every node on their syntactic analysis (not just at the root).

4 Discontent

Brentano’s reclamation of the medieval term “intentional inexistence” in his discussion of content suggests perhaps the most notorious problem to beset the truth-conditional and referential accounts.

Olentzaro is the Basque counterpart of Santa Claus, with a twist. He is a famous *jentillak* or giant and charcoal burner who carries a sickle and is rumoured to slit the throats of disobedient children. In place of a beard, his chin is besmirched with soot. Little Aritz, we can suppose, thinks Olentzaro is mean. In other words, he has a belief with the same content as is encoded in the sentence ‘Olentzaro is mean’. Of course the name ‘Olentzaro’ doesn’t refer to anything—the story is a myth. It follows that there is no state-of-affairs of Olentzaro’s being mean that we can check against the word to see if it obtains, and so (it would appear) no truth-condition corresponding to

⁴For the origins of this approach, see Carnap’s notion of “intensional isomorphism” (1947).

⁵Note that if there is just one mode of composition – e.g. function application – associated with every branching node, then the decoration of the whole tree is determined by the values at the terminal nodes. Under this assumption, we can employ the compact (and customary) practice of representing the referential content of a (binary-branching) complex expression as an ordered pair. Thus the referential content of the NP ‘Plato’s father’ might be represented as $\langle \text{father-of, plato} \rangle$, while the content of ‘Plato’s father is bald’ might be represented as $\langle \text{is-bald, } \langle \text{father-of, plato} \rangle \rangle$.

⁶Assuming a version of the referentialist semantics on which each proposition determines a set of possible worlds.

the sentence. Certainly there is no tree decorated with the referents of the expressions in the sentence.

Another familiar objection to both accounts is Frege's (1892) puzzle. Sentence (3) has, intuitively speaking, a non-trivial content. It contains, as Frege puts it, a valuable extension of our knowledge.

(3) The Morning Star is the Evening Star.

The sentence (4), on the other hand, does seem to have a trivial content.

(4) The Morning Star is the Morning Star.

Since the truth conditions of (3) and (4) are the same, the truth-conditionalist must conclude that their contents are the same. The referentialist is no better off, as the parts of (4) refer to the same entities as the corresponding parts of (3), and so it passes the more stringent referentialist test for sameness of content.

It might be tempting simply to deny Frege's intuition and say that, after all, sentences that differ only by the substitution of coreferring expressions (like (3) and (4)) have the same content. However, as Brian Loar (1976) points out, this leaves us with a notion of content that is too coarse-grained to figure in an account of communication.

Imagine Ralph has been talking about two men, one he saw at the beach, and one in a brown hat whom he noticed in ambiguous circumstances.⁷ Ralph utters (5), intending by 'he' the man in the brown hat.

(5) He is a spy.

You recoil in surprise, thinking he means the man from the beach, a pillar of the community. Intuitively, you have misunderstood the utterance. This is apparent in your reaction, and is the reason he hastens to set you straight. As it happens, the man at the beach and the man in the brown hat are the same man, and so the truth-conditional (and referential) content of your construal *does* match the content of my utterance. Contrary to our intuitions about the case, the referential and truth-conditional accounts of content predict that your initial interpretation of my utterance is correct.

⁷The example is from Quine 1956.

5 Discourse Content

The puzzles faced by truth-conditional and referential accounts of content seem related to the mismatch between reference and discourse reference (denotation) discussed in the first section of this paper. It is not unreasonable to hope, then, that Karttunen’s denotational semantics might be of use in developing an account of content that resists these puzzles. My suggestion is to let the content of a noun phrase be the discourse entity it denotes, and the content of a sentence be its syntactic analysis decorated with the denotation (rather than the reference) of each node. I call this composite entity a *discourse proposition*.⁸

Consider the puzzles once more. Since there can be discourse reference without reference, expressions that don’t refer to anything (e.g. ‘the unicorn’) can nevertheless have a (discourse) content. Furthermore, since we can have two (or more) discourse entities where there is only one real-world referent (consider the reference and discourse reference of ‘the man in the brown hat’ and ‘the man at the beach’), it seems feasible to have expressions that corefer but do not codenote (i.e. don’t have the same content). If this is true, then substitution of coreferring expressions will not always preserve content.

On Karttunen’s scheme, if two expressions have the same discourse reference, an interpreter is required to recognize that they do. An interpreter is not tasked with determining which expressions (among those with differing discourse reference) share a referent. In fact, coming up with an interpretation that gets reference right is, as we have seen, neither necessary nor sufficient for comprehension.

Discourse content, on the other hand, is the right “grain” to serve in an account of communication. Indeed, discourse propositions are the pieces contributed by individual sentences to the discourse model. If constructing an accurate discourse model amounts to understanding the discourse (Karttunen’s working hypothesis), then it would seem that building the discourse proposition corresponding to a sentence must amount to understanding the sentence.

Some work must be done to ensure that discourse content will satisfy all of our requirements on a notion of content. In particular, existing theories of the nature and “lifespan” of discourse entities must be revised.

⁸An alternative account on which propositions are unstructured (and so closer to those of the truth-conditionalist) is possible. It requires us to think of compositional operations like function application as applying at the level of denotation/discourse reference rather than reference (see Muskens 1996 for details).

Kamp (1985) and Asher (1986) have a traditional (truth-conditionalist) notion of content (Asher, 1987, 145), but model intentional states with representational structures analogous to Karttunen’s file. Indeed, on their account the discourse model is a subjective representation (of the interpreter) and drefs are mental symbols. They conclude that representational structure is needed in addition to content to give a theoretically satisfying account of propositional attitudes.⁹

Discourse content, if they had such a notion, would be *syntactic* on Kamp and Asher’s view. A mental state with the content $\langle F, u \rangle$ would simply be one that tokened the mental sentence Fu . On the assumption that every agent has an idiosyncratic mental syntax,¹⁰ this view individuates contents much too finely, as it entails that two agents could never be in mental states with comparable content.

Heim (1982, chap. 3) and van Rooy (2000, 186-7) prefer an account of drefs on which they are part of a shared common ground (and so available intersubjectively). They nevertheless follow Kamp and Asher in treating drefs as syntactic objects (Heim identifies drefs with numerical indices in logical form, while van Rooy speaks of two agents storing information “under the same variable”). Their accounts remain obscure, however (are they appealing to the fabled universal syntax of thought?), and van Rooy abandons (this version of) his account for this reason (168).

As mentioned before, I subscribe to the view of DeVault and Stone (2006) that discourse entities are abstract objects, rather than mental symbols. Just as both players in a game of correspondence chess have access to the same abstract game board, so too do conversational participants manipulate the same discourse context with its population of drefs.

On Karttunen’s original account, drefs have a limited “lifespan,” corresponding to the anaphoric reach of their establishing indefinite (for instance, drefs introduced under the scope of a negation or inside a conditional do not survive outside that connective’s scope). Karttunen doesn’t comment on the lifespan of drefs that are not introduced in the scope of another connective. However, a popular upper limit to the lifespan of a dref is the length of the discourse in which it is established. For Kamp and Reyle (1993), a discourse model must be “proper” (all drefs in the model were introduced in the discourse) to be evaluable for truth. Similarly, Groenendijk and Stokhof (1991) allow anaphora to extend across clauses that are conjoined to the right of

⁹Strawson (1974) adopts a similar expedient in his (earlier) solution to Frege’s puzzle.

¹⁰“On one plausible understanding of cognitive states, the cognitive states of two agents are completely distinct and cannot share reference markers [drefs]” (Asher, 1987, 151).

the one containing the antecedent, but no further. A corollary of this limitation is that discourse content is *discourse bound*: it cannot be expressed outside of a particular discourse.

These restrictions on drefs contrast with the freedoms of content. We require some sentences in different discourses to count as having the same content since we want something learned in one discourse to be capable of transmission to another.

In the next section, I will argue that anaphora is not confined to a single discourse, but indeed can cross discourse boundaries.

6 Proper Names

Names are thought to belong to the category of “definite” noun phrases, along with pronouns, demonstratives and definite descriptions (Geurts, 1997). While indefinite noun phrases standardly function as anaphoric antecedents, definites are, on the contrary, expressions that *require* antecedents. Sometimes, as in (6), an explicit antecedent for a name is given in the discourse.

- (6) Tampa was home to a serial killer named Bobby Joe Long. Long was known as “the Classified-Ad Rapist.”

At other times, however, names appear discourse-initially, for instance when people converse about a mutual friend or a famous personage. Generally, in such cases, an indefinite antecedent would sound odd.

One might suppose that, in these cases, the antecedent for the discourse-initial name is *accommodated* (Lewis, 1979). That is to say, the presupposition on the utterance – in this case a pre-established dref – “comes into existence” as soon as the sentence is uttered. Once this accommodation is taken into account, discourse-initial names will effectively function as *indefinites*.

To prevent accommodation from becoming too powerful (and erasing the distinction between presupposition and assertion, definite and indefinite), some curbs must be placed on it. Accommodation runs smoothly so long as the matter to be accommodated is uncontroversial (so the account goes). In our case, the action to be accommodated (providing an antecedent for a name) would otherwise be accomplished by a statement such as ‘There is someone named $\langle name \rangle$ ’ (as in (6) above). The usual picture of accommodation would therefore predict that a discourse-initial name is felicitous (its antecedent accommodated) whenever the corresponding statement would pass unchallenged (see, for instance, Geurts, 1997).

Since one would rarely take exception to the claim that someone with a particular name exists, the accommodation account predicts that virtually all discourse-initial names are felicitous, a prediction that is empirically false. For instance, if you and I know more than one person with the first name ‘Gottlob’, then I will not object to the claim that there is someone called Gottlob. However, it will not be felicitous for you to utter ‘Gottlob’ discourse-initially (unless some one Gottlob is most salient).

We might stipulate that accommodation of a discourse-initial name takes place whenever the audience believes there is a *unique* (or anyhow uniquely salient) bearer of that name. This is truly a stipulation—I don’t see that it follows from the earlier “no controversy” constraint. There is a problem even for this account, however. It predicts that accommodation will proceed even though the (unique) individual the hearer has in mind differs from the one the speaker intended (for instance, suppose I utter the name ‘Gareth Evans’ and mean the British philosopher, but my interlocutor has only heard of the former Australian foreign minister of that name). Something goes wrong in this case, but the accommodation theorist can’t say what. The conditions for accommodation are met, and so an antecedent is duly (though tacitly) introduced into the discourse. There is no diagnosis of misinterpretation, so long as the name is interpreted by my interlocutor as being anaphorically dependent on the accommodated antecedent.

As a final objection, there seem to be genuine cases of accommodation which are intuitively different from run-of-the-mill discourse-initial names. For example, if you speak in a certain tone of “Agnes,” I will assume you have a significant other of that name (and accommodate a dref for her). This seems a different sort of case – and a better candidate for bona fide accommodation – than a situation in which I am already fully aware (at the beginning of a new discourse) that your partner is called Agnes.

I propose instead that the familiarity presupposition on discourse-initial names (and other definites) can be resolved without accommodation. In such cases, the dref in question is *already familiar* at the beginning of the discourse, because the discourse participants have encountered *the very same dref* in a previous context. For this to be so, drefs must sometimes outlive their originating discourse, which is the same as saying that anaphora can pass over discourse boundaries.

Suppose for a moment that the limits of a (spoken) discourse are defined roughly by temporal continuity. A case of discourse-crossing anaphora would then be the following. Two workmates begin a conversation over lunch one day (constituting one discourse) which they continue (while commencing a new discourse) at lunch on the next day. Suppose on the first day *A* mentions

a new love interest. On the second day, *B* asks ‘So, did *he* call?’ (adverting to the love interest). A natural way to understand the case is to treat the pronoun at the start of the second discourse as anaphoric to an indefinite expression in the first discourse. However, such an analysis is impossible if discourse boundaries are impervious to anaphora.

The argument for trans-discourse anaphora does not depend on any controversial account of the size and shape of discourses. If Grantham is a mutual friend of ours, then it is possible for me to use his name in conversation with you even if there was no previous conversation between us at which the name was introduced. We might, after all, have known Grantham before we knew each other, in which case our respective introductions to the name would have been different. One couldn’t, in this case, simply extend the boundaries of the discourse to include *the* antecedent for the name, at least not straightforwardly.

Indeed, we must be careful to distinguish establishing events from introductions. Though we were each introduced to the name ‘Grantham’ on a different occasion, the name itself (as we are using it) cannot have two distinct establishing episodes, as this would amount to an anaphoric expression having two antecedents. I would suggest that the name ‘Grantham’ (as we are using it) has only one establishing event: Grantham’s baptism.

On Kripke’s (1972) “causal chain” theory of reference (see also Donnellan 1972 and Sommers 1982), a token of a name is anaphorically dependent on an initial baptism from which it receives its reference. This suggests that the baptism, like an indefinite noun phrase, establishes a dref that is retrieved by subsequent use of the name. Moreover, that dref is grounded (or “anchored”) in a particular real-world object (what we might call the dref’s referent) by its establishing baptism,¹¹ and any expression that denotes that dref (i.e. is anaphorically dependent on the baptism) also *refers* to the object.

Of course, one need not have been present at Grantham’s baptism to retrieve the dref established there. As in Kripke’s theory of reference, we can acquire the ability to denote a dref from those already in possession of it. Naturalistic criteria for passing this baton are elusive,¹² but the basic outline goes as follows: a trainee comes in contact with an individual already in possession of a concept (mental symbol) denoting the dref. The trainee acquires a new concept through interaction with this individual and, if the feat succeeds, the trainee’s new concept also denotes the dref.¹³ Note that

¹¹See Evans 1982 for an alternative theory of such grounding.

¹²See Devitt 1974 and more recently DeVault, Oved and Stone 2006 for some attempts.

¹³Note that it might not succeed, as evidenced, I would argue, by Kripke’s cases of Pierre and Peter (Cumming, 2007, chap.4).

the exchange does *not* occur in the following way: the trainer utters an expression that denotes the dref and the trainee thereby acquires it. It is not possible to express a (nonnovel) dref to an audience that is not already familiar with it.

7 The New Limits of Discourse Content

There is a denotation, though no referent, of the name ‘Olentzaro’ (we might write it u_o). It follows that the sentence ‘Olentzaro is mean’ expresses a discourse content (though not a referential or truth-conditional content). The names ‘Hesperus’ and ‘Phosphorus’, while they refer to the same planet, originate in separate baptisms and so denote different drefs (u_h and u_p). We can thus distinguish the discourse contents of sentences (7) and (8):

- (7) Hesperus is Phosphorus.
 $\langle \langle =, u_p \rangle, u_h \rangle$
- (8) Hesperus is Hesperus.
 $\langle \langle =, u_h \rangle, u_h \rangle$

While discourse content is no longer confined to a single discourse, its range is still limited. A discourse proposition containing a dref u is only expressible in contexts in which anaphora is possible to u ’s originating event. Transmission of discourse content is, naturally enough, only possible between contexts that are anaphorically connected. Rather than detracting from the account, this fact actually provides further confirmation of it, as I will proceed to explain.

Consider our motivating cases once more. It seems clear that one cannot entertain the thought that *Olentzaro is mean* unless one is connected (probably via a long chain of speakers) to the origins of the Basque myth. A member of a yet-undiscovered tribe in New Guinea, completely sealed off from Western culture, just could not have a thought with this content. The same goes, if you think about it, for the thought that *Hesperus is Phosphorus* (or, perhaps easier to hear, *Hesperus is Venus*). Even if the tribe has two words for the planet, and even if, by some miracle, those words are phonetically identical to the English words, intuitively the content expressed by the English sentence ‘Hesperus is Venus’ is distinct from the content expressed by the corresponding (anaphorically isolate) sentence of that tribe’s language.¹⁴

¹⁴Devitt (1981, 239, my emphasis) runs a similar thought experiment:

By contrast, we do think that Agamemnon (for example), who is connected to us by a continuous cultural tradition, was capable of entertaining the thought that *Hesperus is Phosphorus*, and that a present day Italian monoglot, whose names for heavenly bodies (despite differing in form) have the same cultural origins as ours, might well believe that *Hesperus is Venus*.

But surely a member of the isolated tribe can think that Venus is bright, and that Venus is not Mars, even though they cannot entertain the discourse content expressed by the corresponding English sentences. The tribe, we are assuming, has its own words (with different discourse contents) that refer to these heavenly bodies. It would appear, then, that true attitude reports involving the name ‘Venus’ make a claim about the *referential* or *truth-conditional* contents of the tribespeople’s thoughts (specifically, that it matches that of the English subordinate clause).

Are both (referential and denotational) notions of content then necessary, perhaps respectively for *de re* and *de dicto* readings of attitude reports? On the contrary, David Kaplan has shown us how to make do with only one kind of content – the sort that figures in the *de dicto* reports. According to his 1968 paper “Quantifying In,” *de re* reports do not relate agents to the content of the subordinate clause *at any grain*. Instead, they involve (approximately speaking) existential quantifications over *de dicto* reports. More carefully, an attitude report that is *de re* in some position in the subordinate clause contains an existential quantification (over appropriately shaped pieces of *de dicto* content) into that position. He uses this analysis to explain the fact (among others) that *de re* beliefs are not closed under certain boolean operations (pointed out by Quine [1956]).

Is anything more required [of an accurate attitude report involving a name] than that [the report and the cognitive state] concern the same name and same object? The discussion suggests there is; they must be linked into networks arising from the same groundings (including the same naming ceremony or suitable substitute). *And this seems intuitively correct here*. Suppose that there are two distinct communities which never communicate with each other, both communities by chance giving the same name to an object. Could an opaque belief statement, including a name token arising from one naming, be confirmed by an expression of belief arising from the other? The situation is so unusual that we may have no clear pretheoretical intuitions about it. However it does seem that if we are to preserve the *point* of opaque belief, we must rule that the statement is not confirmed by such an expression.

He is guarded about his intuitions, but I would guess that this is because he didn’t consider a case like mine where the *de dicto* reading is encouraged by the use of an identity sentence (see the discussion).

On Kaplan’s analysis, with “*de dicto* content” set to discourse content, the *de re* reading of ‘Ralph thinks Ortcutt is astute’ is true iff there is a dref u that refers to Ortcutt and Ralph believes the discourse proposition $\langle \text{astute}, u \rangle$. The dref denoted by the noun phrase ‘the man at the beach’ might qualify as a witness to this claim. Similarly, ‘Bob thinks Venus is bright’ (on its *de re* reading) is true iff there is a dref u that refers to the planet Venus such that Bob believes the discourse proposition $\langle \text{bright}, u \rangle$. The witness to the existential claim is the dref Bob’s tribe uses to talk about Venus. If we attempt to read ‘Bob thinks Hesperus is Venus’ *de re* (in both names), the result is the almost trivial claim that there are—not necessarily distinct—drefs u and u' that refer to Hesperus/Venus and Bob believes $\langle \langle =, u \rangle, u' \rangle$. While we may be able to hear this reading (and judge it true), the untrue *de dicto* reading is far more prominent.

Kaplan’s analysis is made inevitable by the existence of mixed *de re/de dicto* attributions (Cumming, 2008). Suppose Ralph is in the following state: he doesn’t know that the man in the brown hat is Ortcutt and, additionally, he suspects him (the hat) of being the mortal enemy of Sherlock Holmes (he thinks Conan Doyle’s stories are all true accounts of actual criminal investigations). Now take the following report:

(9) Ralph thinks Ortcutt is the mortal enemy of Sherlock Holmes.

Intuitively, (9) has a true reading. On this reading, ‘Ortcutt’ is interpreted *de re* (since it would be false otherwise) while ‘Sherlock Holmes’ is (of necessity) read *de dicto*. Kaplan’s mechanism can account for this smoothly by quantifying into the first position but not the second. It cannot, however, be explained by one who maintains that there are two (and only two) senses of ‘thinks’, one expressing a relation to the discourse content of the subordinate clause and one expressing a relation to its referential content.

8 Coda: Under Modals

I would like to briefly consider the interaction of attitude reports and metaphysical modals. The intuitive interpretation of the claim

(10) Bob could have believed that Olentzaro is mean.

is that it is metaphysically possible that Bob have a belief with the content that the sentence ‘Olentzaro is mean’ actually has. Under analysis, (10) states that there are metaphysically accessible circumstances under which Bob believes the discourse proposition $\langle \text{mean}, u_o \rangle$. Is this statement true?

Suppose Bob is our tribesman, who in actuality has no such belief and moreover is unfamiliar with its component dref. If we allow ourselves enough leeway in stipulating Bob’s destiny, the requirement can be met. After all, Bob could have left his hidden valley and moved to Port Moresby (where he studied English and learned of the Basque Santa).

In the hypothetical situation just outlined, Bob has a belief that he would express with the sentence ‘Olentzaro is mean’. This fact on its own is neither sufficient nor necessary for his having the belief that Olentzaro is mean. We can imagine a possible world much like the actual world, but in which ‘Olentzaro is mean’ means the sky is blue in Bob’s native tongue; in that world, Bob has a belief he would express with the sentence ‘Olentzaro is mean’, but does not believe that Olentzaro is mean. Similarly, we can imagine a world like our hypothetical scenario in which Bob leaves his homeland except that the sentence ‘Olentzaro is mean’ means snow is chartreuse in English (perhaps the words ‘Olentzaro’ and ‘snow’, and ‘mean’ and ‘chartreuse’ are all switched). In the latter world, Bob believes that Olentzaro is mean, but does not have a belief that he would express with the sentence ‘Olentzaro is mean’.

For Bob to believe that Olentzaro is mean at a world, he must have a belief that, were he to express it, he would do so with a sentence that means *Olentzaro is mean* (i.e. $\langle \text{mean}, u_o \rangle$) in that world. In our hypothetical scenario, Bob expresses his belief with the sentence ‘Olentzaro is mean’ and moreover that sentence means *Olentzaro is mean* (and not *snow is chartreuse*).

In the hypothetical scenario, the facts about Bob have changed—he moves away from his valley, etc.—but so have the facts about the dref u_o . For instance, it is “acquired” by Bob, and occurs in a belief of his. On my understanding of a dref as a socially constructed, abstract object, such a scenario is indeed possible. There certainly are worlds in which the relations between us and abstract objects are different. For instance, there is a world in which I weigh 175 kg, rather than 175 lb, and so stand in a different relation to the (abstract) number 175. Suppose, rather than treating a dref as an abstract object with a contingent history, one were to *reduce* a dref to its (actual) historical trace, a connected graph of utterances emanating from an initial baptism.¹⁵ In this case it seems problematic to treat the dref’s “history” as contingent. The dref *is* such a history, and so if the pertinent historical facts were different, it seems we would be dealing with a different dref. Such an account will have trouble verifying (10).

¹⁵Devitt (1981) and Kaplan (1990) both have proposals of this nature.

Conclusion

Making the connection between theories of discourse and theories of content has led us to view *de dicto* content as inextricable from the social milieu. While Putnam and Burge have taught us that an individual needs society to ground their concepts and utterances in the world, this paper offers the surprising claim that content *itself* is the product of social interaction, a claim which, as we have seen, has testable consequences in the semantics of attitude attributions. Following through on the connection has also forced us to revise the picture of discourses as anaphoric “islands,” and provided additional support for the view of proper names as trans-discourse anaphors (Cumming, 2007).

It is not enough to say that the building blocks of content are discourse entities and leave it at that. Discourse entities have a precise aetiology and behavior within theories of the semantics and pragmatics of discourse, but are much harder to spot “in the wild” (so to speak). My preliminary remarks about the metaphysics of drefs are unlikely to seem an advance on what has already been conjectured about Fregean *Sinn*.

The immediate benefits of the connection pursued in this paper therefore remain at the “semantic level.” By fleshing out the relationships between communication and familiarity, between anaphora and attitude reports, I have widened the theoretical domain furnishing the “conceptual role” of discourse entities, and hopefully removed one of the veils shrouding intentionality.

References

- Asher, N. (1986). Belief in discourse representation theory. *Journal of Philosophical Logic*, 15(2):127–89.
- Asher, N. (1987). A typology of attitude verbs. *Linguistics and Philosophy*, 10:125–97.
- Brentano, F. (1973). *Psychology from an empirical standpoint*. Routledge and Kegan Paul, London.
- Burge, T. (1977). Belief *de re*. *Journal of Philosophy*, 74:338–62.
- Carnap, R. (1947). *Meaning and Necessity*. University of Chicago Press, Chicago.

- Chomsky, N. (1982). *Some Concepts and Consequences of the Theory of Government and Binding*. MIT Press, Cambridge, MA.
- Clark, H. (1996). *Using Language*. Cambridge University Press, Cambridge.
- Cumming, S. (2007). *Proper Nouns*. PhD thesis, Rutgers, The State University of New Jersey.
- Cumming, S. (2008). Variabilism. *Philosophical Review*, 117(4):??
- David DeVault, I. O. and Stone, M. (2006). Societal grounding is essential for meaningful language use. In *Proceedings of the 21st National Conference on Artificial Intelligence (AAAI)*, Boston, MA.
- DeVault, D. and Stone, M. (2006). Scorekeeping in an uncertain language game. In Schlangen, D. and Fernandez, R., editors, *BRANDIAL 06: Proceedings of the 10th Workshop on the Semantics and Pragmatics of Dialogue*, pages 139–146.
- Devitt, M. (1974). Singular terms. *The Journal of Philosophy*, 71(7):183–205.
- Devitt, M. (1981). *Designation*. Columbia University Press, New York.
- Donnellan, K. (1972). Proper names and identifying descriptions. In Davidson, D. and Harman, G., editors, *Semantics of Natural Language*, pages 356–379. D. Reidel, Dordrecht.
- Evans, G. (1982). *Varieties of Reference*. Clarendon Press, Oxford.
- Fine, K. (2007). *Semantic Relationism*. Blackwell Publishing, Oxford.
- Fodor, J. A. (1975). *The Language of Thought*. Thomas A. Crowell Co., New York.
- Frege, G. (1892). Uber Sinn und Bedeutung. *Zeitschrift fur Philosophie und philosophische Kritik*, pages 25–50. English Translation: *On Sense and Meaning*, in Brian McGuinness (ed.), *Frege: Collected Works*, pages 157–177. Basil Blackwell, Oxford.
- Geurts, B. (1997). Good news about the description theory of names. *Journal of Semantics*, 14:319–348.
- Grice, P. (1957). Meaning. *The Philosophical Review*, 66:377–88.

- Groenendijk, J. and Stokhof, M. (1991). Dynamic predicate logic. *Linguistics and Philosophy*, 14(1):39–100.
- Heim, I. (1982). *On the Semantics of Definite and Indefinite Noun Phrases*. PhD thesis, University of Massachusetts at Amherst.
- Heim, I. (1983). File change semantics and the familiarity theory of definiteness. In R. Bäuerle, C. S. and von Stechow, A., editors, *Meaning, Use and Interpretation of Language*, pages 164–89. de Gruyter, Berlin.
- Kamp, H. (1985). Context, thought and communication. *Proceedings of the Aristotelian Society*, 85:239–61.
- Kamp, H. and Reyle, U. (1993). *From Discourse to Logic*. Kluwer, Dordrecht.
- Kaplan, D. (1968). Quantifying in. *Synthese*, 19(1-2):178–214.
- Kaplan, D. (1990). Words. *Proceedings of the Aristotelian Society, supplementary volume 64*, pages 93–119.
- Karttunen, L. (1976). Discourse referents. In McCawley, J., editor, *Syntax and Semantics: Notes from the Linguistic Underground*, volume 7, pages 363–386. Academic Press, New York.
- King, J. (2007). *The Nature and Structure of Content*. Oxford University Press, Oxford.
- King, J. and Stanley, J. (2005). Semantics, pragmatics and the role of semantic content. In Szabó, Z., editor, *Semantics versus Pragmatics*, pages 111–64. Oxford University Press, Oxford.
- Kripke, S. (1972). Naming and necessity. In Davidson, D. and Harman, G., editors, *Semantics of Natural Language*, pages 253–355. D. Reidel, Dordrecht.
- Lewis, D. (1979). Scorekeeping in a language game. *Journal of Philosophical Logic*, 8(3):339–359.
- Loar, B. (1976). The semantics of singular terms. *Philosophical Studies*, 30:353–377.
- Marr, D. (1982). *Vision*. Freeman, New York.

- Muskens, R. (1996). Combining Montague semantics and discourse representation. *Linguistics and Philosophy*, 19:143–186.
- Newell, A. (1982). The knowledge level. *Artificial Intelligence*, 18(1):87–127.
- Quine, W. V. (1956). Quantifiers and propositional attitudes. *Journal of Philosophy*, 53:101–111.
- Searle, J. (1975). A taxonomy of illocutionary acts. In Gunderson, K., editor, *Language, Mind and Knowledge: Minnesota Studies in the Philosophy of Science*, vol. 7, pages 344–369. University of Minnesota Press, Minneapolis.
- Sommers, F. (1982). *The logic of natural language*. Clarendon Press, Oxford.
- Stalnaker, R. (1978). Assertion. In Cole, P., editor, *Syntax and Semantics*, volume 9: Pragmatics, pages 315–32. Academic Press, New York. Reprinted in Stalnaker 1999, 78–95.
- Stalnaker, R. (1999). *Context and Content*. Oxford University Press, Oxford.
- Strawson, P. (1974). *Subject and Predicate in Logic and Grammar*. Methuen, London.
- van Rooij, R. (2000). Anaphoric relations across attitude contexts. In von Stechow, P. and Egli, U., editors, *Reference and Anaphoric Relations*, pages 157–181. Kluwer, Dordrecht.