In: Simon Botley and Tony McEnery (eds.) (1998). *New Approaches to Discourse Anaphora: Proceedings of the Second Colloquium on Discourse Anaphora and Anaphor Resolution (DAARC2).* Lancaster University Centre for Computer Corpus Research on Language, Vol. 11. Pp. 212–221.

Towards a discourse-semantic account of donkey anaphora

Pieter A.M. Seuren

1 Donkey anaphora: the problem

Donkey anaphora is a form of pronominal or nominal anaphora which cannot be adequately rendered in the language of modern predicate calculus as devised by Russell at the start of this century. The problem of donkey anaphora consists precisely in the inability of modern predicate calculus to account for it, and is thus entirely theory-dependent. Sentences containing such anaphoric expressions are known as donkey sentences, after Geach's little but influential book *Reference and Generality* of 1962, where the relevant example sentences all involve mention of donkeys.

One may surmise that Geach had taken his cue from the English nominalist philosopher Walter Burleigh, who lived from ±1275 till after 1344, even if Geach does not refer to him but merely speaks (1962:116) of 'another sort of medieval example'. In Burleigh (1988:92), which was probably written around 1328, one finds the example:[1]

(1) Omnis homo habens asinum videt illum
(every man owning a donkey sees it)

Obviously, the modern problem of donkey anaphora could not be known to Burleigh, as Russellian predicate calculus did not exist in his day. Burleigh's problem was of an entirely different nature. Having noticed first that there exist what we now call quantifier-bound pronouns, as in:

(2) Nobody expected that the dog would bite *him* (*them*)

and having stated that such pronouns may never take as antecedent a constituent of the same clause ('propositio categorica'), he presents (1) as an apparent counterexample, since the pronoun *illum* ('it') takes as antecedent *asinum* ('a donkey'), which stands under the same finite verb (*videt*) and is thus in the same clause. He then solves this problem by observing that the antecedent of *illum*, i.e. *asinum*, is not a main constituent of the same clause but a constituent of a subordinate predication, i.e. *habens asinum* ('owning a donkey').

Geach's treatment of (1) and similar cases in his (1962) is, on the whole, rather circumstantial and not always entirely perspicuous, due, it seems, to his double perspective of medieval and modern theories of reference. But one must give him credit for having seen, or at least intuited, the relevance of cases like (1) in the light of modern predicate calculus. Upon reading Geach (1962), logicians and philosophers of language quickly saw the point and have since treated the donkey anaphora problem as a serious threat to the status of modern predicate calculus as a means of rendering and analysing the meanings of all natural language sentences.

The reason why the language of modern predicate calculus (LPC) is unable to render donkey anaphora lies in the fact that LPC, in so far as it remains purely extensional (i.e. allows for substitution of co-referring terms salva veritate), knows only two kinds of argument terms for predicates: constant or referring terms, which refer to a precisely defined and actually existing entity or set of entities, and bound variables. In fact, the first category, that of referring terms, came under attack early on in the piece. In his (1905) Russell tried to tackle the riddle of reference to non-existing entities, as in his celebrated example sentence *The present king of France is bald.* He did so by dissolving definite descriptions into an existential quantifier and a few propositional functions containing the variable bound by it. The American philosopher Quine, in his *Word and Object* (1960:181-6), supported and generalized Russell's analysis of definite terms, proposing that the regimented 'canonical' form of sentences, which displays just their logical properties without the clutter accrued from the impure conditions of usage, should contain no referring expressions at all, so that all statements about the world could be expressed with the help of the two standard quantifiers and propositional functions containing the variables bound by them. Hence his slogan: 'To be

[1] I owe this reference to Joachim Ballweg.

is to be the value of a variable'. A systematic application of this program of 'elimination of singular terms' should, in Quine's eyes, provide answers to all the many problems connected with reference. Later on, when Kripke proposed his analysis of proper names as rigid designators (Kripke 1972), mainstream philosophy of language and formal semantics allowed proper names back in as the only admissible class of referring expressions in the semantic analysis of sentences. On the whole, the main burden of reference came to lie with the bound variables. Yet, with or without Russell's analysis of definite descriptions, or Quine's program of elimination of singular terms, or Kripke's analysis of proper names, LPC still remains unable to cater for donkey anaphora.

Let us have a look at the details. Donkey anaphora in its simplest form, as diagnosed by Geach, occurs in sentences of the following types:

(3)a. If Smith owns a donkey, he feeds *it*
b. Every farmer who owns a donkey feeds *it*
c. Either Smith owns no donkey, or he feeds *it*

The pronoun *it* in these sentences cannot be rendered in LPC as a constant or referring term, since, in the semantics standardly provided with LPC, these sentences must be considered true in worlds not containing any donkeys at all (in which case the clause containing the troublesome pronoun *it* no longer counts for the truth calculus). So the *it* must be rendered as a bound variable. However, as is shown in (4a-c), the normal and most straightforward translation method of natural language sentences into LPC leads to trouble:

(4)a. $\exists x[\text{Donkey}(x) \wedge \text{Own}(\text{Smith},x)] \rightarrow \text{Feed}(\text{Smith},\mathbf{?})$
b. $\forall y[\text{Farmer}(y) \wedge \exists x[\text{Donkey}(x) \wedge \text{Own}(y,x)] \rightarrow \text{Feed}(y,\mathbf{?})]$
c. $\neg \exists x[\text{Donkey}(x) \wedge \text{Own}(\text{Smith},x)] \vee \text{Feed}(\text{Smith},\mathbf{?})$

The problem is, obviously, that the variable *x*, which one would like to see in the position of the question mark in (4a-c), is not allowed to occur there because it falls outside the scope of the (existential) quantifier binding it.

Quine (1960) shows no awareness of the donkey anaphora problem. He does, however, deal with a similar problem posed by sentences of the type (1960:138):

(5) If any member contributes, he gets a poppy

If the word *any* is taken to represent the existential quantifier, the pronoun *he* is left stranded, or, as Quine says (1960:139), 'left high and dry':

(6) $\exists x[\text{Member}(x) \wedge \text{Contribute}(x)] \rightarrow \text{Get a poppy}(\mathbf{?})$

Quine then proposes not to use the existential quantifier and get the universal quantifier to do all the work, stipulating that 'by a simple and irreducible trait of English usage', *every* always takes the smallest and *any* the largest possible scope. This proposal was used later to solve the problem of the stranded variables in cases like (4a-c). The idea was to translate (3a-c) with the help of the universal quantifier only:

(7)a. $\forall x[\text{Donkey}(x) \rightarrow [\text{Own}(\text{Smith},x) \rightarrow \text{Feed}(\text{Smith},x)]]$
b. $\forall x \forall y[[\text{Donkey}(x) \wedge \text{Farmer}(y) \wedge \text{Own}(y,x)] \rightarrow \text{Feed}(y,x)]]$
c. $\forall x[\text{Donkey}(x) \rightarrow [\neg \text{Own}(\text{Smith},x) \vee \text{Feed}(\text{Smith},x)]]$

This does indeed eliminate the scope problem raised by (4a-c). There are, however, other problems that now raise their head. First, one wonders why natural language chooses to use surface structure representatives of the existential quantifier (*a donkey*, *no donkey*), allowing unbound variables to dangle, clearly without any problem for natural interpretation processes. In other words, one wonders what could justify the sudden change in the translation or mapping relation between the logical and the grammatical form of the sentences involved.

Secondly, Quine forgot to mention that in a sentence like *If John has bought **some**thing, I'll take **it** away* the same scope problem occurs for *it* as in *If John has bought **any**thing, I'll take **it** away.* But he does not put forward the daring proposal to treat the word *some* as a representative of the universal quantifier.

Thirdly, and perhaps most seriously, Quine's universal quantifier ploy fails to work in certain other, less simple cases, in particular those where the clauses in question are placed under different sentential operators. Consider, for example, the following sentences:

(8)a. If it's a bad thing that Smith owns a donkey, it's a good thing he feeds *it*
b. Every farmer who is thought to own a donkey is expected to feed *it*
c. Either Smith no longer owns a donkey, or he still feeds *it*

Now try Quine's ploy. Any form of binding the pronoun *it* under the universal quantifier results in an analysis that is incompatible with what the sentence in question means. Consider such binding for (8a-c). The only way of binding *it* under the spurious universal quantifier is to place the quantifier over the whole structure, as in (9a-c), respectively:

(9)a. $\forall x[Donkey(x) \rightarrow [BAD[Own(Smith,x)] \rightarrow GOOD[Feed(Smith,x)]]]$

b. $\forall x \forall y[[Donkey(x) \wedge Farmer(y) \wedge THOUGHT[Own(y,x)]] \rightarrow EXPECTED[Feed (y,x)]]$

c. $\forall x[Donkey(x) \rightarrow [NO\ LONGER[Own(Smith,x)] \vee STILL[Feed(Smith,x)]]]$

(9a), however, fails because it is gratuitously true if there are no donkeys, whereas (8a) cannot be true in such a case. Since both operators *it's a bad thing* and *it's a good thing* are factive, (8a) presupposes that Smith owns and feeds a donkey, and presupposition failure cannot lead to truth. Moreover, (8a) is not a statement about all donkeys, but rather says that if it's a bad thing that Smith is a donkey-owner, it's a good thing that he feeds the donkey he owns. Likewise for (9b), which is true in all cases where there is no specific donkey thought to be owned by any specific farmer. In that case, no donkey and no farmer will satisfy 'THOUGHT [Own (x,y)]', which makes (9b) gratuitously true. But (8b) can be false in such a case, namely when there is at least one farmer thought to be a donkey-owner (though no-one has an idea about which donkey he owns) but not expected to feed the animal he is thought to own. For (8b) is not about all specific donkeys but about farmers who are thought to be donkey-owners. Likewise again for (9c), which is true when there are no donkeys at all, and also, if there are donkeys, when Smith still feeds them all even if he has never owned any of them. But (8c) cannot be considered true in such cases.

One is, therefore, forced to conclude that pronouns occurring in a position like that of *he* in (5) above are also instances of donkey anaphora, in that LPC is unable to account for them.

A few years after his (1962), Geach discovered a further category of anaphora that seemed to elude LPC. Consider (Geach 1969:149):

(10) Socrates owned a dog, and *it* bit Socrates

What we have here is an anaphoric pick-up of an entity just introduced into the discourse by means of an existential quantifier. This form of anaphora is called 'primary anaphora' in Seuren (1998:448). One would, says Geach, be inclined to treat such cases as ordinary conjunctions of propositions, and thus of the form '$p \wedge q$'. But this won't do, says Geach, for $p \wedge q$ is incompatible with $p \wedge \neg q$, whereas (10) is perfectly compatible with:

(11) Socrates owned a dog, and it did not bite Socrates

All that is needed for both (10) and (11) to be true at the same time is for Socrates to own two dogs, one of which bit him and the other did not. What is needed, says Geach, is a different logical analysis of (10) and, for that matter, of (11), where the pronoun *it* is treated as a variable bound by one single existential quantifier that takes scope over propositional functions, as in (12a) and (12b), respectively:

(12)a. $\exists x[Dog(x) \wedge Own(Socrates,x) \wedge Bite(x,Socrates)]$

b. $\exists x[Dog(x) \wedge Own(Socrates,x) \wedge \neg Bite(x,Socrates)]$

Curiously, again, Geach appears to have taken his cue from Walter Burleigh, who adds the following comment to his discussion of (1) (1988:92-3):

> It follows that the following are compatible: 'Every man owning a donkey sees it' and 'Some man owning a donkey does not see it'. For assuming that every man owns two donkeys, one of which he sees and one of which he does not see, then it is not only true to say 'Every man owning a donkey sees it', but also to say 'Some man owning a donkey does not see it'. In the same way, suppose that every man who has a son also has two sons, and that he loves the one but hates the other, then both the following are true: 'Every man who has a son loves him' and 'Some man who has a son does not love him'.

Apparently, Geach found a great deal of inspiration in medieval authors, whose ideas and examples he managed to mould into a modern discussion.

Geach's analysis of (10) and (11), however, as given in (12), does not cut ice. In Seuren (1977) it is shown that this analysis founders on scope problems when sentential operators are introduced. Consider, for example:

(13) I know that Socrates owned a dog, and I hope *it* bit Socrates

There are two ways of binding the pronoun *it* under an existential operator:

(14)a. $\exists x[Dog(x) \wedge Know(I, [Own(Socrates,x)]) \wedge Hope(I, [Bite(x,Socrates)])]$

b. $Know(I, [\exists x[Dog(x) \wedge Own(Socrates,x) \wedge Hope(I, [Bite(x,Socrates)])]])$

But neither corresponds with what (13) means. (14a) speaks of a specific dog such that I know Socrates owned it and such that I hope it bit Socrates. But (13) is not about a specific dog.

(14b) says that I know the following: (a) that Socrates owned a dog and (b) that I hope it bit him. But again, (13) does not say that I know that I hope that Socrates' dog bit him. What one would like is an analysis of, indeed, the form 'p ∧ q', but with a stranded variable, just as in the standard cases of donkey anaphora:

(15) Know(I, [∃x[Dog(x) ∧ Own(Socrates,x)]]) ∧ Hope(I, [Bite(**?**,Socrates)])

We conclude, therefore, that primary anaphora is a further instance of donkey anaphora, which cannot be handled in terms of LPC.

To compound the difficulty further one may observe that the pronoun *it* in the sentences in question allows for replacement with a so-called epithet anaphor:

(16)a. If Smith owns a donkey, he feeds *the wretched animal*
b. Every farmer who owns a donkey feeds *the poor creature*
c. Either Smith owns no donkey, or he feeds *his precious treasure*
d. If it's a bad thing that Smith owns a donkey, it's a good thing he feeds *the creature*
e. Every farmer who is thought to own a donkey is expected to feed *the animal*
f. Either Smith no longer owns a donkey, or he still feeds *his treasure*

This substitution is typically allowed for referring pronouns but impossible for bound variable pronouns, as one sees when an epithet is put in the place of the bound variable pronoun *him*:

(17) Nobody expected that the dog would bite *the poor sod*

Now the expression *the poor sod* must refer to an outside person and cannot be interpreted as being bound by the quantifier *nobody.* It thus seems highly unlikely that the *it* in (3a-c) and (8a-c) should be taken to represent a bound variable. This *it* behaves like a referring expression, and not like a bound variable pronoun. Geach's proposal, therefore, lacks generality, which makes it unsatisfactory. And we are still saddled with the donkey anaphora problem.

The problem of donkey anaphora should not be dismissed too lightly. Modern predicate calculus has become the backbone of all of logic and most of philosophy of language. And even if one feels, as the present author does, that it fails to do justice to natural language on many counts, one cannot deny its enormous analytical power. One simply has to take modern predicate calculus seriously, and it seems sound policy to attempt to remedy any faults one may find with it by revision and extension, rather than think of wholesale rejection.

The **language** of modern predicate calculus, in particular, has proved to be of great empirical relevance to the study of syntax, in that it is becoming increasingly clear that the best choice for the format of syntactic deep structure is one that follows the structural principles of LPC. One may have one's doubts about the ways in which modern predicate calculus has been applied to natural language by philosophers and logicians. One may, for example, reject Quine's program of elimination of singular terms or Kripke's theory of proper names, or claim full rehabilitation for referring expressions. One may wish to change the standard semantics for LPC from a truth-conditional calculus to a discourse-incremental procedure. But the central position of LPC in the business of analysing and describing natural language sentences will remain untouched.

2 Towards a solution

When faced with the task of solving a serious problem it seems good policy to try and see where the fault lies. In the case of donkey anaphora, the fault lies not so much in the language of predicate calculus as in the semantics associated with it. Once the semantics has been repaired, the language will easily be adapted accordingly.

The semantics of LPC consists in the application of model theory: a sentence S in LPC is true in a 'world' **W** of a model **M** just in case the n-tuple of its term extensions in **W** is a member of the extension of the predicate in **W**.[2] The set of conditions for the truth of S in all possible worlds is taken to define the meaning of S.

This is so not only for purely extensional sentences, but also for those that contain intensional predicates or operators, as these are defined in such a way that it is the intension of the intensional term that serves as its extension. The systematic application of this idea, originally due to Frege (1892), is called the program of extensionalisation of intensions.

In fact, this model-theoretic account of truth is nothing but a systematic application of the common version of Aristotle's definition of truth as correspondence between what is said

[2] Standard model-theoretic semantics will add that the truth-conditional operators as well as the quantifiers are to be defined syncategorematically, i.e. by special stipulation. This, however, is not essential, since these operators can all be regarded as predicates – a treatment that would seem preferable. The truth-conditional operators can be interpreted as predicates over (pairs of) truth-values, while the quantifiers are, in fact, nothing but predicates over (pairs of) sets.

and what is the case: model-theory simply makes this notion of correspondence explicit. However, as shown in Stegmüller (1957:16-17) and in Seuren (1998:12-18), Aristotle entertained two distinct notions of truth, a **verbal** notion defining truth as correspondence between what is **said** and what is the case, and a **cognitive** notion which defines truth as correspondence between what is **thought** and what is the case. The verbal notion has become standard and has been universally applied in formal logic, whereas the cognitive notion has not had much of a career. The reason for this is, as is pointed out by Stegmüller (ib.), that it has proved less hard to analyse sentences than to analyse thoughts in an empirically justifiable way, since thoughts are much farther removed from perceptibility. For that reason, it was easier for the verbal notion of truth to achieve demonstrable success. The cognitive notion was doomed to remain a purely philosophical construct, without much in the way of formal underpinning. Now, however, as the formal underpinnings of the verbal notion have come to stand out in great clarity and precision, this chicken has come home to roost.

For in current standard semantic theory the semantics of LPC has become the semantics of natural language: every natural language sentence is assigned its logical form in terms of LPC, and it is then the LPC semantics that is taken to be the appropriate semantics for the sentences in question. The LPC logical form is Quine's 'regimented' or 'canonical' form mentioned above, meant to express just the logical properties, leaving out the clutter that is due to the exigencies of communicative usage.

Over the past twenty or so years, however, a number of serious, and, it would seem, insurmountable problems have presented themselves for this model-theoretic paradigm, which has been inspired by mathematical logic (see Seuren 1998:388-404).[3] We are now beginning to realize that this paradigm may well work for the language associated with standard mathematics, which is meant to be directly applicable to a real world domain, but that it will have to fail for natural human language, which functions through the intermediary of processes of production and perception and processes of thought. For the past twenty or twenty-five years, a movement has been gaining ground that no longer supports the rarified, logic-inspired view of natural language semantics.

The various strands composing this new movement share the view that, for speakers of natural languages, truth and falsity are primarily properties of propositional thoughts, and only in a derived sense of sentences as their verbal expressions, whose linguistic material is normally no more than a hint as to what is meant and understood. Yet the consequences of this view are rarely fully realized. In practice, what one sees is mainly a series of attempts to adapt the mathematical-logical paradigm to the demands of a more cognition-oriented approach. In other words, the spirit too often remains mathematical, whereas what one is doing is a variety of cognitive science. If the consequences are drawn to the full, it will become clear that one should support the **cognitive** notion of truth, rejecting the **verbal** notion, and favour a cognitive semantics, rather than one based on mathematical logic. It is to be expected that this new approach, which in fact harks back to very traditional but non-formal notions, will make for a much better fit with respect to the facts of language, and will avoid the pitfalls of model-theoretic semantics, including donkey anaphora.[4]

What, then, is to be done? What policy is to be outlined for research along the lines sketched above? First of all, one should avoid undue haste. Rather than jump to formalizations, one should reflect first on the nature of the problem, and make a preliminary problem analysis. In rough outline, such an analysis reveals a few general and essential aspects where cognitive semantics is bound to be at odds with model-theoretic semantics.

First, it appears that uttered sentences are interpreted in terms of a usually highly **restricted verification domain**, and not, as in standard formal semantics, in terms of one or more whole 'worlds'. This particular organization of linguistic interpretation demands the specification for each uttered sentence of its domain of interpretation, a condition which is alien to standard formal semantics. The processes whereby restricted verification domains are delimited are essentially of a cognitive nature. They must, therefore, be analysed and described in terms of, and in interaction with, cognitive science.

---

[3] The best-known, and universally recognized, problem is that of propositional attitudes. See for example Dowty et al. (1981:175).

[4] The more or less radically discourse-related strands in semantics are represented mainly by Barwise & Perry (1983), Fauconnier (1985), Kamp & Reyle (1993, which was set up primarily to deal with donkey anaphora), Seuren (1985). Developments like data or file semantics in various forms are less radical and represent adaptations of existing model-theoretic (possible world) semantics. The views defended here are those of, or agree with, Seuren (1985).

Secondly, there is the factor of **implicit information**. It is now generally recognized that large and essential elements of linguistic comprehension are supplied by situational and world knowledge, supposed to be available to the speaker-hearer. Although linguistic comprehension is guided and restricted by the incoming verbal material, it is assisted by appeals to available knowledge and natural inductive and deductive powers.

Thirdly, natural language comprehension is, by constitution, **incremental**. This means that for each coherent discourse there is a specific cognitive working space, called 'discourse domain', where the information supplied by previous utterances is stored in a specific format. A discourse domain (DD) will contain representations of entities ('addresses'), of putative facts ('increments'), subdomains representing what may not be real facts but are thought, imagined, hoped, etc. to be facts, and instructions with regard to the further development of the domain. Thus, negation results in a prohibitive instruction: 'do not increment the negated proposition'. Likewise with the universal quantifier: a sentence like *Every farmer owns a donkey* will result in an instruction to set up a separate donkey address for every farmer address that may crop up in the DD. In general, we say that a universal quantifier is incremented as an instruction to add a specific property to any address of the appropriate kind that need may be felt to set up.

The information stored in a DD relates not only to possible or real factual situations, but also to the commitment the speaker has entered into with regard to each sentence uttered by him or her. Discourse semantics thus contains speech act theory as an essential part. We shall see in a moment that this is a crucial feature of discourse semantics.

DDs are to be seen as special, purpose-built modules, acting as a link between verbal messages and their cognitive integration. They are an essential element in the cognitive function of delimiting the verification domain of each uttered sentence.

It looks, therefore, as if the actual, practical processing of linguistic messages in situations of linguistic communication (and probably also, in a much wider sense, of all information entering cognition) is *per se* domain-restricted, and thus requires specific mechanisms that filter out the relevance area of the information concerned. This is an essential modification of the Aristotelian and Medieval (Modist) paradigm of the one-to-one mapping of the categories of thought, being, and speaking, and, therefore, likewise of the modern paradigm of model-theoretic semantics. It also poses a serious and extremely interesting challenge to the semanticist's powers of formalization.

A new semantic theory, based on the considerations given above, will differ from model-theoretic semantics in several essential respects. First of all, semantic interpretation must be seen as a process that is partly dependent on available situational and world knowledge, which is a serious complication for the principle of compositionality, the hallmark of model-theoretic semantics. Then, the theory will have to reckon with discourse-dependent mechanisms, such as presuppositions, definite reference, including deixis and anaphora, and mechanisms of lexical adaptation, such as polysemy and metaphor — all of them much trodden but hitherto not very productive grounds. One would intuitively expect that these mechanisms will be seen to be highly functional for the actual, everyday functioning of language in communicative situations. Presupposition and definite reference appear to be essentially energy-saving devices, as they allow the speaker to leave a great deal of what he has to say unsaid, relying on properties of lexical meaning and discourse focussing.[5] The same goes for polysemy and metaphor, which depend on hitherto little understood flexibilities of lexical meaning. The theory will, moreover, provide room for a semantic account of discourse structure organization, in particular organization into question-answer pairs (topic-comment modulation). All in all, the dividing line between semantics and pragmatics (which typically deals with the 'clutter' of communicative language use not accounted for in model-theoretic semantics) will be drawn differently, as semantics will be taking over a large part of the explanatory task assumed by pragmatics.

What this amounts to is a program of intensionalization of extensions: all semantics is intensional. There is a 'commitment domain', which represents what speaker-hearers accept as binding commitments with respect to the facts (verification domain) under discussion. Necessarily, however, this commitment domain is itself a mental construct, just like the mental constructs that are considered intensional (and hence block substitutivity). The commitments laid down in the commitment domain often apply to imagined situations whose relation to what is considered the real world is defined by the governing predicate. Thus, when I say that John thinks that the moon is made of green cheese, I have made myself responsible

---

[5] For the discourse-functionality of presupposition and the mechanism of post hoc insertion (accommodation), see Seuren (1985:284-95).

for the truth of the statement that John thinks that the moon is made of green cheese, not for the truth of the statement that the moon is made of green cheese. This latter putative fact is stored in a subdomain under the heading of what John thinks. Thus there are intensional subdomains for what the speaker, or a person mentioned in the discourse, hopes, realizes, has forgotten, etc., or for what he or she considers possible, probable, fortunate, etc.

Likewise for the classical truth-conditional operators *or* and *if.* When I say that John is either a physicist or a chemist, I establish two subdomains and commit myself to a choice between these two, each excluding the other. Comprehensibility demands that both subdomains link up normally with the commitment domain already established. And when I say that if John is a physicist he is not a very good one, I establish a little conditional subdomain where the antecedent clause *John is a physicist* is held up as a possible increment to the commitment domain, but with the condition that it cannot be incremented without adding the consequent clause *he is not a very good one.* An *if A - then B* construction thus embodies the commitment that in case A is incremented it has to be followed by B.

This point is important, because in standard logic-based semantics, the operators *or* and *if* are truth-functional and extensional, whereas in discourse semantics they are neither truth-functional nor extensional but intensional. Consider, for example, the following two sentences (Seuren 1985:396):

(18)a. Either the morning star is the evening star, or there are ten planets
b. If the morning star is not the evening star, there are ten planets

According to standard propositional calculus, substitution *salva veritate* is guaranteed in these cases: if, for example, the term *evening star* is replaced with *morning star*, the result may be weird but in the actual world, where the two names refer to the same entity, the planet Venus, the truth-value does not change. This may be so in standard logic, yet many speakers who accvept the truth of (18a,b), will have doubts about the truth of (19a,b), which, if true, are so for quite different reasons They certainly express quite different speaker commitments:

(19)a. Either the morning star is the morning star, or there are ten planets
b. If the morning star is not the morning star, there are ten planets

If this observation is correct, it follows that natural language semantics must treat the operators *or* and *if* differently from standard logic. In fact, I claim that they correspond to intensional subdomains and are not truth-functional at all. They can be treated as truth-functional operators only in contexts where speaker-commitments do not count and all the emphasis is on a calculus of entailments, i.e. logical and mathematical contexts. Natural language semantics, in other words, is not truth-conditional but commitment-conditional.

The case of *or* and *if* not only illustrates the profound difference between a logic-based and a cognition-and-communication-based discourse semantics, it is also directly relevant to the discourse-semantic treatment of donkey anaphora, as will become clear in a moment.

The aspect we are most directly interested in, in the present context, is the relation of reference. Definite referring terms in a sentence are no longer considered to refer directly to their reference object, by some Kripkean flight of metaphysics, but indirectly, through the intermediary of cognitive constructs or representations of reified entities, denizens of discourse domains, which correspond to entities presumed to be there in the real world in all sorts of complicated ways.[6] That is, linguistic interpretation involves first and foremost a relation between any given definite term and a cognitive representation, an ‘address’, in the DD at hand. Any relation with really existing entities is a matter of cognitive, rather than linguistic, interpretation.

We thus distinguish (Seuren 1985:456) between the **denotation** of a definite referring term occurring in an uttered sentence, which selects a discourse address (or set of addresses) for it, and the **designation** of a given address (or set of addresses), which defines the cognitive-intentional relation between the address or addresses in question and possible or actual world entities as they are managed by general cognition. Definite referring terms **denote** their addresses, and the addresses **designate** cognitively manageable entities. Taking

[6] One only has to think briefly in order to realize the complicated relationship between reified cognitive constructs on the one hand, and bare minimal reality on the other. Reifying terms like *the military-industrial complex*, *the average Boston cabdriver*, *the difference between A and B*, *the Prime Minister’s predilection for bondage*, etc. etc. certainly function as referring terms, yet what they actually refer to is less than obvious. Discourse semantics best treats these terms as corresponding to discourse entities of a certain category, whose relation to existing world entities requires a complex defining formula in terms of a given ontology. In practice, such a complex formula of interpretation is often not available when people use language. Speakers, more often than not, make use of vague and suggestive reifications that lack precise definitions.

denotation and designation together, we may say that definite referring terms **refer to** cognitively manageable entities. What counts for strict linguistic interpretation is the denotation relation, not the relation of designation, which is, though of course connected with linguistic interpretation, less bound up with the mechanisms of language and more with general cognitive functions of intentionality and focussing.

So to explain donkey anaphora we need a theory that sets up discourse addresses for the anaphoric pronoun to denote. For primary anaphora this is, in principle, simple and also generally accepted in discourse-related semantic theories. For a sentence like:

(20) John bought a painting and he sold **it** immediately

the DD will react to the existential quantifier represented by the indefinite article *a* by setting up a new address for the painting John bought, and assign to John the property of buying it. The next increment will process *he sold it immediately*, which now has the new painting address at its disposal.[7] This account of primary anaphora provides an immediate solution to the case brought to bear by Geach in (10) and (11) above:

(10) Socrates owned a dog, and it bit Socrates
(11) Socrates owned a dog, and it did not bite Socrates

All that is needed to account for such cases with higher operators, such as (13) quoted above:

(13) I know that Socrates owned a dog, and I hope it bit Socrates

is that referring expressions, including anaphoric pronouns, be allowed to denote addresses in a different subdomain. But this is needed on independent grounds , as one sees from e.g.:

(21) John claims that he owns a Ferrari, but I have never seen *it*

We are now close to a principled solution to the donkey anaphora cases of example sentences (3a) and (3c) above, repeated here for convenience:

(3)a. If Smith owns a donkey, he feeds *it*
c. Either Smith owns no donkey, or he feeds *it*

In (3a) the DD sets up an *if*-domain as described above for the consecutive conjunction of, first, *Smith owns a donkey*, and next *he feeds it.* Both increments are in the same domain and members of a consecutive conjunction. The anaphoric pronoun *it* is a simple case of primary anaphora, but embedded in an *if*-subdomain. The same solution applies to Quine's case (5) quoted above: *If any member contributes, he gets a poppy.*

In (3c) the two sentences are processed in alternative *or*-subdomains, as described above. The first subdomain, which processes the negative sentence *Smith owns no donkey*, will contain a prohibition for the incrementation of *Smith owns a donkey*. The second subdomain, however, now appears to lack an antecedent for *it.* This constitutes a considerable problem, for which we propose, in principle, the following solution.

Let us introduce the notion of **virtual increments** and **virtual addresses**. We say that in incrementing a disjunction (with *or*) the second alternative subdomain contains first the virtual increment of the negation of what was incremented in the first subdomain. Thus:

(22) John is at home or his lights are out

will be incremented as two alternative subdomains, as described, but the second subdomain will contain the virtual increment of *John is not at home*: "<John is not at home and> his lights are out" (with the virtual increment between angled brackets). This, in general, will account for the exclusive character of natural language disjunction so often diagnosed by authors throughout the ages (cp. Seuren 1985:333-7). If this is accepted, we have the missing antecedent for *it* in (3c), which will then be incremented as:

(23) not[Smith owns a donkey] OR <Smith owns a donkey and> he feeds it

It must be observed that virtual increments and addresses are not an ad hoc invention. Moxey & Sanford (1986/7) found that a large majority of test subjects, who were given the task of completing sentences like:

(24) Few members were at the meeting. *They ...*

interpreted the pronoun *they* as referring to the members who were **not** at the meeting: they would, for example, complete the sentence as *They were out on the lake with their girlfriends.* The only way out of this dilemma seems to be the assumption of, indeed, a virtual address

[7] In event-describing discourse domains, successive increments represent successive events. Thus the difference between, for example, *He went to Spain and got rich*, and *He got rich and went to Spain.* We say that natural language conjunction is **consecutive**.

induced by the quantifier *few* and representing the complement of the set delimited by *few*. Moxey & Sanford, accordingly, speak of 'complement anaphora'.

So we are left with case (3b):

(3)b. Every farmer who owns a donkey feeds it

Here we have to do with the universal quantifier *every*, which, as has been said above, results in an instruction to add a specific property, the **matrix property**, to any address of the appropriate kind that may have to be set up, the **restrictor class** of addresses. Here the restrictor class consists of those addresses that are characterized by the properties of being a farmer and of owning a donkey, and the new matrix property to be added is that of 'feeding **it**'. The question is: what is the antecedent of *it*? The answer to this question can only be that the required antecedent must be found in the instruction that contains the definition of the restrictor addresses: "If you find, in the DD at hand, an address containing the properties 'Farmer(x)' and 'Own(x,a_n)', where a_n stands for a donkey-address, then add the property 'Feed(x,a_n)' ".

It thus appears that antecedents for anaphors can be found inside instructions as well. This means that we must make provisions for antecedents in the same (sub)domain, across (sub)domains, in virtual increments or addresses, and, finally, inside instructions. Further investigations will have to show the tenability of this view. Meanwhile, we may observe that antecedents within instructions have to be postulated also for cases like:

(25) Every farmer owns a donkey and feeds *it*

where the addresses in the restrictor class of farmers must be enriched with the property 'Own(x,a donkey) & Feed(x,**it**).

We conclude that, in discourse semantics, donkey anaphora does not constitute a separate problem, but is an integral part of the whole machinery of anaphora. It is not feasible, in the present context, to elaborate these ideas any further. We have tried to provide a sketchy and global idea of how donkey anaphora, and other forms of anaphora as well, can in principle be handled in the context of a cognition-oriented discourse semantics. We have not been able to present a fully formalized account of donkey anaphora, for the simple reason that such a formalization necessarily has to be part of an overall formal account of anaphora, which again must be part of an overall formal account of discourse semantics. And that is, for the time being, beyond our grasp.

References:

Barwise, J. & J. Perry 1983: *Situations and Attitudes.* MIT Press, Cambridge, Mass.

Burleigh, W. 1988: *Von der Reinheit der Kunst der Logik. Erster Traktat. Von den Eigenschaften der Termini.* (=De puritate artis logicae. De proprietatibus terminorum). Translated and edited by Peter Kunze, with introduction and commentary. Felix Meiner Verlag, Hamburg.

Dowty, D., R. E. Wall, S. R. Peters 1981: *Introduction to Montague Semantics.* Reidel, Dordrecht.

Fauconnier, G. 1985: *Mental Spaces. Aspects of Meaning Construction in Natural Language.* MIT Press, Cambridge, Mass.

Frege, G. 1892: Ueber Sinn und Bedeutung. *Zeitschrift für Philosophie und philosophische Kritik* 100:25-50.

Geach P.T. 1962: *Reference and Generality. An Examination of Some Medieval and Modern Theories.* Cornell University Press, Ithaca, NY.

Geach, P.T. 1969: Quine's syntactical insights. In: D. Davidson & J. Hintikka (eds.), *Words and Objections. Essays on the Work of W.V. Quine.* Reidel, Dordrecht.Pp 146-57. (Also in Geach 1972:115-27)

Geach, P.T. 1972: *Logic Matters.* Blackwell, Oxford.

Kamp, H. & U. Reyle 1993: *From Discourse to Logic. Introduction to Model-Theoretic Semantics of Natural Language, Formal Logic and Discourse Representation Theory.* Kluwer, Dordrecht.

Kripke, S. 1972: Naming and necessity. In: D. Davidson & G. Harman (eds.), *Semantics of Natural Language*, Reidel, Dordrecht. Pp. 253-355.
(Also 1980, *Naming and necessity.* Blackwell, Oxford.)

Moxey, L. M. and A. J. Sanford 1986/7: Quantifiers and focus. *Journal of Semantics* 5:189-206.

Quine, W.V.O. 1960: *Word and Object.* MIT Press, Cambridge, Massachusetts.

Russell, B. 1905: On denoting. *Mind* **14**: 479-93

Seuren, P.A.M. 1977: Forme logique et forme sémantique: un argument contre M. Geach. *Logique et Analyse* 79, vol. 20:338-47.

Seuren, P.A.M. 1985: *Discourse Semantics.* Blackwell, Oxford.

Seuren, P.A.M. 1994: Donkey sentences. In: R.E. Asher & J.M.Y. Simpson (eds.), *The Encyclopedia of Language and Linguistics.* Oxford, Pergamon Press. Pp. 1059-60

Seuren, P. A.M. 1998: *Western Linguistics. An Historical Introduction.* Blackwell, Oxford.

Stegmüller, W. 1957: *Das Wahrheitsproblem und die Idee der Semantik. Eine Einleitung in die Theorien von A. Tarski und R. Carnap.* Springer, Vienna.