Anaphora, accessibility, and bridging

Bart Geurts

Abstract

Dynamic theories of discourse interpretation seek to describe and explain antecedent-anaphor relations with the help of discourse referents. In a dynamic framework, it is the function of indefinite expressions to introduce new discourse referents, whilst anaphoric expressions serve to retrieve them. Dynamic theories provide a simple and intuitively appealing solution to a variety of problems. For instance, they explain how it is possible for an indefinite expression to "bind" a pronoun that isn't c-commanded by it, as in (1a,b), and why examples like (2a,b) are infelicitous:

(1) a. Pedro owns a donkey$_i$. It$_i$ is grey.
    b. Every farmer who owns a donkey$_i$ beats it$_i$.

(2) a. *Pedro doesn't own a donkey$_i$. It$_i$ is grey.
    b. *Every farmer who doesn't own a donkey$_i$ beats it$_i$.

While the dynamic turn has proved to be as fruitful as it is intuitive, it is not without its problems. One of the main worries has been to account for what I propose to call "piggyback anaphora" (examples by Karttunen):

(3) a. You must write a letter$_i$ to your parents. It$_i$ has to be sent by email.
    b. Harvey courts a girl$_i$ at every convention. She$_i$ always comes to the banquet with him.

The characteristic feature of this type of anaphora is that, intuitively speaking, the anaphoric link is enabled by the fact that the anaphor sits in the scope of an expression that quantifies over the same range of entities as the expression whose scope contains the intended antecedent. This is the guiding intuition underlying most accounts of the phenomenon, but although I agree that this is the right way to go, I also believe that these accounts are systematically flawed. The key to a more adequate solution to this problem, I will argue, is that the anaphors in (3a,b) also rely on bridging inferences.

1. Discourse referents

Classical logic provides us with two candidate models for analysing anaphoric pronouns. One is the bound variable: in (1), "she" may be construed as a variable bound by "every contestant":

(1) Every contestant is convinced that she will win the prize.

This captures one possible construal of the sentence, but there is another, in which the pronoun targets one particular individual. For this reading, the model of the free variable is more suitable: the value of "she" is fixed by a preselected assignment function, to be thought of as part of the context in which the sentence is uttered.

Unfortunately, there are plenty of pronoun occurrences that fit neither model; perhaps the simplest type of example is the following:

(2) Our vicar owns a Toyota$_i$. It$_i$ is grey.

We would like to say that "a Toyota" is the most natural antecedent for "it", but as long as we are sticking with classical logic we can say no such thing. On the one hand, the pronoun isn't bound by its intended antecedent, for this would entail that the scope of "a Toyota" extends beyond the sentence in which it occurs. On the other hand, "it" doesn't seem to function as a free variable, either: somehow or other the pronoun owes its referent to the indefinite, though without being bound by it.

Using a notion introduced by Karttunen in his trailblazing 1971 paper, the most intuitive way of describing what goes on in (2) is that the indefinite "a Toyota" serves to introduce a *discourse referent*, which is subsequently picked up by the pronoun. The purpose of this article is to trace how this notion developed since the 1970s, and then present and perhaps solve some of the problems it is up against, with emphasis on what I will call "piggyback anaphora".

One of the truly impressive achievements of Karttunen's paper is that although at the time he didn't have anything like a theory of discourse referents, his discussion evinces a clear sense of what such a theory would be able to do, and where it might run into trouble. Karttunen begins by showing how a discourse like (2) contrasts with:

(3) Our vicar doesn't own a Toyota$_i$. *It$_i$ is grey.

In a first stab at capturing this contrast, Karttunen suggests that, whereas in (2) the indefinite introduces a discourse referent for the pronoun to pick up, in (3) no such discourse referent is established. However, later

in the paper he refines the notion of discourse referent to accommodate examples like the following:

(4) a. John wants to catch a fish$_i$ and eat it$_i$ for supper. *Do you see the fish$_i$ over there?
    b. I don't believe that Mary had a baby$_i$ and named her$_i$ Sue. *The baby$_i$ has mumps.

In the face of these data, Karttunen settles for the idea that discourse referents have "lifespans" of varying lengths. In (4a) the lifespan of the discourse referent introduced by way of "a fish" is bounded by the scope of "wants", while in (4b) it is coterminous with the scope of "believe". However, as Karttunen points out, "the lifespan of a short-term [discourse] referent is not always so neatly bound as the above examples suggest." (p. 11) To illustrate the lack of neatness, Karttunen offers the following observations:

(5) a. You must write a letter$_i$ to your parents. It$_i$ has to be sent by email. The letter$_i$ must get there by tomorrow.
    b. Harvey courts a girl$_i$ at every convention. She$_i$ always comes to the banquet with him. The girl$_i$ is usually also very pretty.
    c. If Mary had a car$_i$, she would take me to work in it$_i$. I could drive the car$_i$ too.
    d. I wish Mary had a car$_i$. She would take me to work in it$_i$. I could drive the car$_i$ too.

In each of these discourses, a pronoun and then a definite NP unexpectedly manage to "refer back" to a narrow-scope indefinite in the opening sentence. What makes this possible, apparently, is that the anaphors are in the scope of an expression that somehow extends the mode of quantification (over worlds, events, or what have you) of the first sentence. For example, the first sentence of (5a) quantifies over worlds, and the subsequent sentences quantify over the same worlds. Similarly, the first sentence of (5b) quantifies over events of a certain kind, and the following sentences quantify over the same events. Similarly, mutatis mutandis, for the other discourses in (5). So, although the details of the process are very much unclear, it seems obvious enough that the anaphors in (5) are piggybacking on the interpretation of the expressions in whose scope they occur; I will therefore call them "piggyback anaphors".

I'm proposing "piggyback anaphora" as a name for a natural kind that, to the best of knowledge, has not been named before. To my mind, the examples in (5) are so similar that we should expect them to admit of

a unified explanation. However, the post-Karttunen literature suggests a different picture, with authors usually focusing on one class of facts to the exclusion of others. An important factor in this development may have been Roberts's (1987, 1989) influential work on "modal subordination". Concentrating her attention on data like Karttunen's (5a) and (5c), Roberts had very little to say about the others, and thus may have fostered the impression that the modal cases are special. Be this as it may, I see no reason to believe that this impression is correct, and as far as I know, it has never even been argued that it is, either. Hence, I will proceed on the assumption that the data in (5) exemplify a single phenomenon, and are therefore entitled to a unified explanation.

The following examples are evidently related to the ones in (5), but create further problems of their own:

(6) a. Suppose Mary had a car$_i$. She takes me to work in it$_i$. I drive the car$_i$ too. (Karttunen 1971)
    b. In each room there was a cat$_i$ and a goldfish$_j$. The cat$_i$ was eyeing the goldfish$_j$. (Stenning 1978)
    c. Every chess set comes with a spare pawn$_i$. It$_i$ is taped to the top of the box. (Sells 1985)
    d. Every director gave a present$_i$ to [a child from the orphanage]$_j$. They$_j$ opened them$_i$ rightaway. (Kamp and Reyle 1993)

Although the anaphoric expressions in these examples are perfectly felicitous, none of them are in the scope of an overt quantifier that could bind them. However, in all of these cases, it is plausible to assume that piggyback anaphora is enabled by *implicit* quantification. For example, the intended interpretation of the second sentence in (6c) is something like, "*In every chess set* the spare pawn is taped to the top of the box." The other examples in (6) are similar. The problem, then, is to explain where the implicit quantifiers come from. In the case of (6d) this is fairly clear (though the plural morphology of "them" raises some hairy problems): it is the distributive reading of "they" that calls for a universal quantifier (Kamp and Reyle 1993, Krifka 1996). The examples in (6a-c) are considerably less straightforward, but as the special problems caused by such examples will have to be addressed anyhow, we will leave them aside here (the problems as well as the examples), and confine our attention to those cases where piggyback anaphors are enabled by overt quantification over individuals, worlds, and so on.

Another point that should be mentioned but will not be addressed is that piggybacking is not confined to pronouns and definite descriptions,

but manifests itself with presuppositional expressions generally. As far as I know, this was first pointed out by Roberts (1995) and Geurts (1995) (though see McCawley 1981), but the relevant examples have been around since the early seventies:

(7) a. I wish Mary had a car. She would take me to work in it. I could drive the car *too*. (= (5d); Karttunen 1971)
b. Bill believed that Fred had been beating his wife and hoped that Fred would *stop* beating her. (Karttunen 1973)
c. Possibly Boris killed Louis and possibly Boris *regrets* killing Louis. (Gazdar 1979)

The factive verb "regret" in (7c) triggers the presupposition that Boris killed Louis. This presupposition is satisfied in the possible worlds selected by the first conjunct, and if we may assume that the second modal sentence resumes these possibilities, the construction is essentially the same as what we saw in (5). The same, mutatis mutandis, for (7a,b). (Sequenced attitude reports like (7b) are discussed at length by Heim 1992 and Geurts 1995, 1999.)

In the remainder of this paper we will focus our attention on the facts illustrated by the discourses in (5).

2. Neo-Russellian revisionism

Some people will say that the very idea of piggyback anaphora is completely off the track, as is the concept of discourse referent, for that matter. What we need instead, according to such people, is just the quantifiers and variables of classical logic, plus Russell's (1905) theory of descriptions. Pronouns are either bound variables or else go proxy for definite descriptions à la Russell. Hence, in (7a) "it" is a kind of shorthand for "Mary's car"; it is an "E-type pronoun" (Evans 1977, Heim 1990, Neale 1990, Elbourne 2005).

This sort of theory comes in a great many varieties, all of which come with their own epicycles, and I have neither the space nor the inclination to argue against all of them. The main purpose of this section is to bring out the signal differences between descriptivist theories and discourse-based theories of anaphora inspired by Karttunen's work.

Russell's theory of descriptions famously construes "The cup is empty" as "There is one and only one cup, and it is empty". That is to say, it claims that the semantics of the definite article entails descriptive unique-

ness. This claim is either false or vacuous. Consider the following state of affairs:

In this situation, "the square with the dot" is a perfectly felicitous and unequivocal means of referring to the square on the right (Haddock 1987). But, prima facie at least, this is *not* the unique square with the unique dot; for there are two dots. So if we take the uniqueness claim at face value, it is false.

Needless to say, various kinds of repair strategies are on offer. One is to claim that the occurrence of "the dot" in "the square with the dot" is actually short for a more elaborate description that does meet the uniqueness requirement. For example, if "the dot" abbreviates "the dot in the square on the right", then "the square with the dot" comes out meaning "the square with the dot in the square on the right", and everything is hunky dory. (I know this sounds preposterous, but I didn't make it up.) Another strategy is to slice situations so finely that, within some narrowly circumscribed situation, "the dot" will refer uniquely, after all (cf. Heim 1990, Elbourne 2005). The idea is that one of the minimal constituents of the situation pictured above is a situation containing just the square on the right, and in this situation there is only one dot.

The two strategies are obviously related, and they equally render the Russellian analysis vacuous: since all things are unique, it is *always* possible to find, for any given individual, some description that fits, and some situation that contains, that individual and nothing else.

It is somewhat of a mystery why one should want to insists on *making* definite descriptions uniquely referring when it is perfectly obvious that they don't have to be. It is never a problem if several individuals match a definite description as long as there is something—anything—that sets apart the intended referent from its detractors. In the example above, the rightmost dot is the right one because the referent of "the dot" is to help selecting between the three squares, so the one in the middle is so much as irrelevant—and that's all there is to it.

Given that Russell's theory of descriptions is seriously problematic, there is every reason to be suspicious of the notion that some pronouns go proxy for Russellian definite descriptions. Still, let us briefly consider how the story goes, using the following example:

(8) Last week a student$_i$ came to my office. She$_i$ is German but her Dutch is impeccable.

On the neo-Russellian view defended by Neale, for example, the pronoun "she" in the second sentence goes proxy for, say, "the student who came to my office last week". To begin with, it should be obvious that this proposals inherits all the problems attendant on Russell's treatment of definite description; in this case, the problem is how to avoid the entailment that only one student came to my office last week. But in addition the descriptivist view on pronouns gives rise to the question of how pronouns manage to select suitable descriptions, in the first place. As the discussion in the last part of Heim's 1990 paper makes abundantly clear, this is a very tough question indeed, and as far as I know there isn't even the beginning of a viable answer.

Ever since Evans (1977), the notion that some pronouns are definite descriptions in disguise has proved to be strangely appealing. The reason for this, I suspect, is that for *practical* purposes there may be no harm in saying that, for example, the pronoun "she" in (8) means "the student who came to my office last week"; a speaker may do this to clarify his meaning. But this observation should not be mistaken for the seed of a theory of pronouns.

Finally, I should like to note that, from a psychological point of view, the descriptivist approach has a pronounced cart-before-the-horse flavour. Once it has been linked to its antecedent it may indeed be possible to paraphrase a pronoun by means of a definite description. But since by then we have the antecedent, what point is there in doing so? Small wonder, therefore, that processing theories invariably adopt a Karttunen-style approach to anaphora.

3. Discourse representation theory

The theory prefigured in Karttunen's 1971 paper saw the light ten years later in the work of Kamp (1981) and Heim (1982). Kamp's Discourse Representation Theory and Heim's File Change Semantics were the first of a long and successful line of so-called "dynamic" theories of semantics developed since by Groenendijk and Stokhof (1990, 1991), Muskens (1996), Veltman (1996), Beaver (2001), and many others. In this article I will gloss over the differences between these various theories, and DRT will be my framework of choice. By and large, this will not matter much, though one caveat is in order. As is well known, DRT is a representational theory of discourse interpretation, which is to say that it postulates a level of semantic representations, called "discourse representation structures", which other dynamic theories prefer to do without. There has been a

somewhat tedious debate over this issue (see Geurts and Beaver 2007 for recent discussion), which I will pass over here, but it should be noted that in a non-representational framework it may be difficult to implement some of the ideas to be discussed in the following, especially Sections 6 and 7.

A discourse representation structure (DRS) is a mental representation built up by the hearer as the discourse unfolds. A DRS consists of two parts: a universe of discourse referents and a set of DRS-conditions which encode the information that has accumulated on these discourse referents. The following DRS represents the information that there are two individuals, one of which is a farmer, the other a donkey, and that the former chased the latter:

(9) [x, y: farmer(x), donkey(y), x chased y]

The universe of this DRS contains two discourse referents, x and y, and its condition set is {farmer(x), donkey(y), x chased y}.

A DRS like the one in (9) can be given a straightforward model-theoretic interpretation. In DRT this is done by means of embedding functions, which are partial functions from discourse referents to individuals in a given model M. An embedding function f verifies (9) in M iff the domain of f includes at least x and y, and according to M it is the case that f(x) is a farmer, f(y) is a donkey, and f(x) chased f(y).

The DRS in (9) is designed to reflect the intuitive meaning of:

(10) A farmer chased a donkey.

In the absence of any information about the context in which this sentence is uttered, the semantic representation of (10) is (9). So the indefinite expressions "a farmer" and "a donkey" are not treated as regular quantifiers; rather, they prompt the introduction of two new discourse referents, x and y, and contribute the information that x is a farmer and y a donkey; the verb adds to this that the former chased the latter.

If a discourse opens with an utterance of (10), the DRS in (9) is constructed, and this DRS forms the background against which the next utterance is interpreted, which might be (11a), for example:

(11) a. He caught it.
b. [$\underline{v}$, $\underline{\underline{w}}$: v caught w]

(11b) is the DRS that reflects the semantic content of (11a) before the pronouns are resolved. In this DRS, the anaphoric pronouns "he" and "it"

in (11a) are represented by the discourse referents v and w, respectively, which are underlined to indicate that they want to be identified with discourse referents that are given already. (The double underlining merely serves to distinguish between anaphors.) (11a) is uttered in the context of (9), so the next step in the interpretation of this sentence is to merge the DRS in (11b) with that in (9), the result of which is (12a):

(12) a. [x, y, v, w: farmer(x), donkey(y), x chased y, v caught w]
b. [x, y, v, w: v = x, w = y,
farmer(x), donkey(y), x chased y, v caught w]
c. [x, y: farmer(x), donkey(y), x chased y, x caught y]

Since (11a) is immediately preceded by (10), the most likely antecedents of "he" and "it" are "a farmer" and "a donkey", respectively. At DRS level, this is represented by equating v with x and w with y. These equations yield (12b), which is equivalent to (12c). Either DRS is verified in any model featuring a farmer who chased and caught a donkey.

Thus far, we have only considered DRSs with simple conditions, but in order to account for negated and conditional sentences, say, complex conditions are required.

(13) a. Pedro doesn't have a donkey.
b. [x: Pedro(x), ¬[y: donkey(y), x owns y]]

(13b) is the sentence DRS corresponding to (13a). This DRS contains a condition that consists of a DRS prefixed by a negation sign. An embedding function f verifies (13b) in a model M iff f maps x onto an individual in M which "is a Pedro", i.e. which is called "Pedro", and f cannot be extended to a function g which verifies the embedded DRS; that is to say, no such g should map y onto a donkey owned by Pedro.

The negated DRS in (13b) contains a token of the discourse referent x which is introduced in the main DRS. Apart from that, the embedded DRS also introduces a discourse referent of its own, i.e. y, which is associated with the indefinite NP "a donkey", and whose lifespan is delimited by the sub-DRS in which it is introduced. This explains Karttunen's observation that if (13a) were followed by (14a), for example, the pronoun could not be linked to the indefinite:

(14) a. It is grey.
b. [z: grey(z)]
c. [x, z: Pedro(x), ¬[y: donkey(y), x owns y], grey(z)]

If we merge (13b) and (14b), which is the sentence DRS associated with

(14a), we obtain (14c). In this representation, the discourse referent z does not have access to y, because y is introduced in a DRS that is not accessible to the DRS in which z is introduced, and therefore it is not possible to bind z to y. Thus, Karttunen's notion that a discourse referent may have a limited lifespan is explained in terms of accessibility. And as we will see in the next section, accessibility boils down to the standard notion of scope, albeit that the DRT logic is non-standard.

Like negated sentences, conditionals give rise to complex DRS-conditions, too:

(15) a. If Pedro owns a donkey, he beats it.
b. [x: Pedro(x), [y: donkey(y), x owns y] ⇒ [$\underline{v}$, $\underline{\underline{w}}$: v beats w]]
c. [x, v: Pedro(x), v = x,
[y, w: w = y, donkey(y), x owns y] ⇒ [: v beats w]]
d. [x: Pedro(x), [: [y: donkey(y), x owns y] ⇒ [: x beats y]]

(15b) is the sentence DRS associated with (15a), and assuming for convenience that this sentence is uttered in an empty context, it is also the initial DRS of the discourse. The complex condition in this structure is interpreted as follows: if f is to verify (15b) in the current model, then f(x) must be an individual called "Pedro", and every extension of f which verifies the antecedent-DRS must itself be extendable to a function that verifies the consequent-DRS. It follows from this that the main DRS in (15b) is accessible to the antecedent-DRS, which in its turn is accessible to the consequent-DRS, and therefore v may be linked to x (accessibility being a transitive relation) and w to y. The result is (15c), which is equivalent to (15d), with both DRSs expressing that Pedro beats every donkey he owns.

The interpretation of quantified sentences is very similar to what we have just seen:

(16) a. Every farmer who owns a donkey$_i$ beats it$_i$.
b. [: [x, y: farmer(x), donkey(y), x owns y]⟨∀x⟩[$\underline{u}$: x beats w]]
c. [: [x, y: farmer(x), donkey(y), x owns y]⟨∀x⟩[: x beats y]]

There are various ways of spelling out the interpretation of so-called duplex conditions of the form K⟨∀x⟩K′. Here we will settle for the weak interpretation, on which (16a) comes out meaning that every farmer who owns a donkey beats at least one of his donkeys.

It is still a matter of debate what the exact truth conditions of donkey sentences are. It is traditionally assumed that a sentence like (16a) is standardly interpreted as implying that every farmer beats every donkey

he owns, so the sentence's truth conditions would be strong by default, though this assumption is not borne out by experimental data (see Geurts 2002 for further discussion). But whether weak or strong, the interpretation of a condition of the form $K\langle Qx\rangle K'$, where Q may be any quantifier, makes K accessible to K', and in this respect conditionals and quantified sentences are alike. Consequently, the discourse referent x in (16b) is accessible to u, and the latter may be equated to the former. The resulting representation is equivalent to (16c).

The DRT analysis of quantified expressions like "all" or "most" is fairly standard. A quantifier binds a variable and delivers the truth conditions one should expect. Indefinites are different. An indefinite like "a donkey" is treated not as a quantifier but as a device for introducing a discourse referent and some descriptive material in the form of DRS-conditions; on the DRT account, indefinites have no quantifying force of their own. What quantifying force they seem to have is not theirs, but derives from the environment in which they occur (see also Lewis 1975). If the semantic material associated with "a donkey" is introduced in the main DRS, as in (12), the quantifying effect will be existential, owing to the fact that this DRS is verified in a model M iff *there is* a way of verifying it in M. If the semantic material associated with "a donkey" is introduced in the antecedent of a conditional, as in (15), the quantifying effect will be universal, owing to the fact that a condition $K \Rightarrow K'$ is verified in M iff *every* way of verifying K can be extended to a way of verifying K'. This context-dependent force is a consequence of the way DRT fleshes out Karttunen's idea that indefinites serve to introduce discourse referents. What Karttunen called the "lifespan" of a discourse referent is modeled by treating discourse referents as bound variables and extending the notion of scope in two ways: across sentence boundaries and within quantifying and conditional structures.

In order to make these ideas a bit more precise, we define the DRS language as follows:

DRSs and DRS-conditions

- A DRS K is a pair $\langle U_K, Con_K\rangle$, where U_K is a set of discourse referents, and Con_K is a set of DRS-conditions.
- If P is an n-place predicate, and $x_1, \ldots, x_n$ are discourse referents, then $P(x_1, \ldots, x_n)$ is a DRS-condition.
- If x and y are discourse referents, then $x = y$ is a DRS-condition.
- If K and K' are DRSs, then $\neg K$, $K \Rightarrow K'$, and $K \vee K'$ are DRS-conditions.

- If K and K′ are DRSs and x is a discourse referent, then K⟨∀x⟩K′ is a DRS-condition.

The truth-conditional semantics of the DRS language is given by defining when an embedding function verifies a DRS in a given model. An embedding function is a partial mapping from discourse referents to individuals. Given two embedding functions f and g and a DRS K, we say that g extends f with respect to K, or f[K]g for short, iff Dom(g) = Dom(f) ∪ U_K, and for all x in Dom(f): f(x) = g(x). Viewing functions as sets of pairs, this can be formulated more succinctly as follows:

f[K]g iff f ⊆ g and Dom(g) = Dom(f) ∪ U_K

We now proceed to define what it takes for an embedding function to verify a DRS or DRS-condition in a given model. As usual, a model M is a pair ⟨D,I⟩, where D is a set of individuals and I is an interpretation function that assigns sets of individuals to one-place predicates, sets of pairs of individuals to two-place predicates, and so on. To enhance the legibility of the definition somewhat the qualification "in M" is omitted throughout:

Verifying embeddings

- f verifies a DRS K iff f verifies all conditions in Con_K.
- f verifies $P(x_1, \ldots, x_n)$ iff $\langle f(x_1), \ldots, f(x_n)\rangle \in I(P)$.
- f verifies x = y iff f(x) = f(y).
- f verifies ¬K iff there is no g such that f[K]g and g verifies K.
- f verifies K ∨ K′ iff f there is a g such that f[K]g and g verifies K or f[K′]g and g verifies K′.
- f verifies K ⇒ K′ iff, for all g such that f[K]g and g verifies K, there is an h such that g[K′]h and h verifies K′.
- f verifies K⟨∀x⟩K′ iff, for all individuals d ∈ D, if there is a g such that f[K]g, g(x) = d, and g verifies K, then there is an h such that g[K′]h and h verifies K′.

A DRS is true in a model iff we can find a verifying embedding for it:

Truth

A DRS K is true in a model M iff there is an embedding function f such that Dom(f) = U_K and f verifies K in M.

Let us now have a closer look at the notion of accessibility, in terms of which we will formulate DRT's central claim about the interpretation of anaphoric expressions. Accessibility is a relation between DRSs that is transitive and reflexive, i.e. it is a preorder. More precisely, it is the smallest preorder for which the following holds, for all DRSs K, K′, and K″:

Accessibility

If Con_K contains a condition of the form …

- $\neg K'$ then K is accessible to K′
- $K' \vee K''$ then K is accessible to K′ and K″
- $K' \Rightarrow K''$ then K is accessible to K′ and K′ is accessible to K″
- $K' \langle \forall x \rangle K''$ then K is accessible to K′ and K′ is accessible to K″

The "accessible domain" of a DRS K contains all and only those discourse referents that are introduced in DRSs accessible to K:

Accessible domains

$A_K = \{x \mid$ there is a K′ such that K′ is accessible to K and $x \in U_{K'}\}$

Note that, since accessibility is a reflexive relation, it always holds that $U_K \subseteq A_K$. If x is a discourse referent introduced in K, i.e. $x \in U_K$, then we will say that all and only the members of A_K are accessible to x.

Consider now the following, somewhat abstract example:

(17) a. If there is an A, it is a B.
b. [: [x: A(x)] ⇒ [y: B(y), y = x]]
c. [: [x: A(x)] ⇒ [: B(x)]]

(17b) is the DRS for (17a), where x and y are the discourse referents associated with "an A" and "it", respectively. The model-theoretic interpretation of the DRS language guarantees that the occurrences of x in [x: A(x)] and [y: B(y), y = x] are "the same", in the sense that the range of possible values of x in [y: B(y), y = x] is fixed by [x: A(x)]; for the meaning of "⇒" entails that no g is eligible as a verifying embedding of [y: B(y), y = x] unless there is some $f \subseteq g$ such that f is a verifying embedding of [x: A(x)]. It is for this reason that (17b) and (17c) are equivalent.

As it turns out, two occurrences of any discourse referent x are guaranteed to covary in their values iff one of them is accessible to the other. Hence, the semantics of DRT entails the following constraint on the interpretation of anaphoric expressions:

Accessibility constraint

Let x be the discourse referent associated with a given anaphoric expression: then x must be equated with a discourse referent that is accessible to it.

Thus the following contrast, observed already by Karttunen (Section 1), is explained by the fact that, unlike (18a), (18b) cannot be interpreted in such a way that the accessibility constraint is satisfied:

(18) a. Our vicar owns a Toyota$_i$. It$_i$ is grey. (= (2))
b. Our vicar doesn't own a Toyota$_i$. *It$_i$ is grey. (= (3))

From a logical point of view, DRT construes anaphoric expressions as bound variables. But its apparatus of variable binding is heterodox: indefinites aren't treated as quantifiers, though in a sense they are variable binders, and anaphors may be bound across a quantifier, across "if", and even across a sentence boundary. Such is the logic of DRT and other dynamic theories of interpretation.

Compared to syntax-based theories, theories of the DRT family extend the binding domain of an indefinite in two ways. First, DRT makes it possible for an indefinite sitting in the restrictor of a quantifying expression Q, like "all" or "if … then", to bind an anaphor in Q's scope. Secondly, DRT enables indefinites to bind anaphors across a conjunctive expression or a sentence boundary (the latter may be seen as a special instance of conjunction). Empirically speaking, these are the defining features of DRT and its kin.

4. Accessibility, givenness, and bridging

The accessibility constraint imposes restrictions on the interpretation of anaphoric expressions. But what are anaphoric expressions? In the foregoing I assumed as a matter of course that personal pronouns and definite descriptions fall under this rubric, without even trying to define it. The accessibility constraint may be seen as a first shot at doing just this: we can define an anaphoric expression as any expression that is represented by a discourse referent x which must be equated to some accessible discourse referent $y \in A_K$, where K is the DRS in which x is introduced. The intuitive idea is that A_K contains the discourse entities that are *given* at the point where x enters the discourse representation. Hence, the function of an anaphoric expression is to anchor an utterance in given information. This is, of course, a view with a long pedigree, but note that the tradi-

tional view is extended by taking it below the sentence level. For example, in (17a) the information contained in the antecedent of the conditional is accessible, and therefore given, from the vantage point of the consequent.

In order to make this more precise, a lot more has to be said about how information comes to be given. Thus far, we have restricted our attention to anaphoric expressions that serve to retrieve a discourse referent previously introduced by an indefinite NP. But there are other ways in which discourse referent may become given. For example, "the chandelier" may be used to refer to a fixture in the room where the discourse takes place, and "the moon" is generally used to refer to the Earth's only natural satellite. In cases like these the intended referent need not have been mentioned in the previous discourse, and they are accounted for quite naturally by assuming that common knowledge is *ipso facto* given, i.e. represented in the main DRS.

Things become considerably murkier when definite descriptions are used to refer to entities that are not given in the strict sense of the word:

(19) I was at a wedding last week.
   a. The bride was pregnant.
   b. The mock turtle soup was a dream.

These are instances of "bridging" (Haviland and Clark 1974), in which a definite NP is used for identifying a referent that wasn't given previously, and that may be more (19a) or less (19b) expectable under the circumstances. However, even such data may be brought in line with the notion that anaphoric expressions serve to pick up discourse referents that are given, as follows. For examples like (19a) it may be argued that, although the existence of a bride is not entailed by the given information (after all, it could have been a gay wedding), the fact that there was a marriage makes it quite likely that there was a bride, so she counts as given.

Whatever the merits of this argument, it surely does not apply to (19b): although mock turtle soup may be served at a wedding, it is doubtful that, in general, the mere mention of a wedding will raise mock-turtle-soup expectations to any significant degree. But even such cases may be consistent with the view that definite descriptions are anaphoric in the sense that their function is to select discourse referents that are given. For it may be argued that examples like (19b) involve *accommodation* in the sense of Lewis (1979). Accommodation occurs if a speaker chooses to present new information *as if* it were given already. Here, the speaker uses the expression "the mock turtle soup" because a proper introduction of the dish would have been too much of a hassle, and he reckons that his

audience won't mind if he treats it as given. In effect, the speaker expects his audience to update their DRSs with a discourse referent for the mock turtle soup before they go on interpreting the definite description. So on this view of how examples like (19b) work, we can say that "the mock turtle soup" refers to a given entity, after all, even if the whole thing is a bit of a charade.

All this is more or less common lore. What is less widely realised is that these observations hold for anaphoric pronouns just as much as they do for definite descriptions. To begin with, the intended referent of a pronoun need not be given by the previous discourse. If Jill, who has just left the house, returns, and Jack guesses that she forgot her car keys, he may help her out by saying:

(20) They're on the kitchen table.

Similarly, if a friend of Jill's calls and Jack answers the phone, he may utter (21) without previous mention of Jill:

(21) I'm sorry, she just left for work.

In these examples, it is the situational context that furnishes the pronouns' referents, but otherwise they are no different from cases in which the intended referent is introduced by way of an indefinite NP. Nor do there seem to be any qualitative differences between the kind of pronominal reference witnessed in (20)-(21) and the "situational uses" of definite descriptions discussed earlier. The main difference is that the descriptive content of a personal pronoun is much paltrier than that of a definite description like "the chandelier" or "the moon", as a consequence of which the context is even more important in these cases (and therefore this use of personal pronouns will be rarer), but this is a quantitative difference, not one in kind (Bosch 1983).

Basically, anaphoric pronouns are just semantically attenuate definite descriptions, and not very different from non-lexical definites like "the man" or "the thing". Most importantly, for the purposes of this paper, pronouns can be interpreted by way of bridging just as full definite NPs can, though again the possibilities are somewhat reduced due to the fact that the descriptive content of a pronoun is relatively poor:

(22) a. When the doorbell rang I thought *it* was Vernon.
b. What's that shadow creeping up the wall? Could *it* be a burglar?
c. The Jones's had been happily married for six years when *he* became unemployed.

d. A car's coming up to the junction and *he* starts to turn right. (Yule 1982)
e. John bled so much *it* soaked through his bandage and stained his shirt. (Tic Douloureux 1971)
f. When Little Johnny threw up, was there any pencil-eraser in *it*? (ibid.)
g. Maxine was kidnapped but *they* didn't hurt her. (Bolinger 1977)

In each of these examples, the italicised pronoun has to be interpreted by way of bridging, since there is no antecedent expression introducing a suitable discourse referent. For example, the neuter pronoun in (22a) must be construed as referring to the person who rang the doorbell, and similarly for the other examples.

Some of the readers of an earlier version of this paper have suggested to me that the neuter pronouns in (22a) and (22b) are expletive rather than referential. Note, however, that in both cases we can exchange "it" for a demonstrative pronoun:

(23) a. When the doorbell rang I thought that was Vernon.
b. What's that shadow creeping up the wall? Could that be a burglar?

These are perhaps more marked than the original examples, but that may be attributed to the fact that "that" is marked vis-à-vis "it".

One of the stock-in-trade examples of the DRT literature is due to Barbara Partee:

(24) a. Exactly one of the ten marbles is not in the bag. *It* is under the couch.
b. Exactly nine of the ten marbles are in the bag. ?*It* is under the couch.

As observed by Kamp et al. (2004), among many others, this minimal pair shows that anaphora is contingent on more than truth-conditional content alone: the first sentence in (24a) expresses the same propositional content as its counterpart in (24b), yet anaphora is possible in the former case but not in the latter. It is often suggested that this is because the first sentence in (24a) introduces a discourse referent for the missing marble, whilst its counterpart in (24b) does not. In my view, this is a misleading, or at the very least incomplete way of describing what goes on in these examples. Note, to begin with, that it is perfectly possible to retrieve the missing marble in (24b) by means of an anaphoric expression:

(24) b′. Exactly nine of the ten marbles are in the bag. *The missing marble* is under the couch.

Furthermore, the acceptability of examples like (24b) can be improved by changing the context. Suppose that last night there was a breakout from the local prison. With 150 inmates to be accounted for, it is reported to the governor that:

(25) 138 of the prisoners are safe in their cells.

In this situation, the governor might very well exclaim, "I want them back before noon!", referring to the 12 escapees. The trick of this scenario is that it makes the missing prisoners so important that they become sufficiently salient to be retrieved by pronominal means. The same trick will work to improve (24b). Suppose that the marbles in question are pure gold and owned by the addressee. Then the first sentence of (24b) will raise the salience of the missing marble to such a degree that the speaker can reassure its owner by saying: "Don't you worry: it's under the couch."

I agree with Kamp et al. that the opening sentence of (24b) does not introduce a discourse referent for the missing marble, and that this is why the anaphor in the second sentence is odd—at first. However, the opening sentence entails that there is a missing marble, and with some help from the context this may prompt the introduction of a discourse referent, after all. The concept of bridging will reappear in Section 6, where it will figure prominently in an account of piggyback anaphora.

5. Piggyback anaphora again

All in all, the vanilla version of DRT outlined in Section 3 captures Karttunen's ideas about discourse referents quite well: an indefinite introduces a discourse referent that may be picked up by an anaphoric expression, and whose "lifespan" is bounded by any operator that has the indefinite in its scope. More accurately: if an indefinite occurs in the scope of one or more operators, the lifespan of its discourse referent is the scope of the innermost operator. This is what Karttunen had in mind, with one important exception: it doesn't give us an account of piggyback anaphora. To see what causes the problem, let us have another look at Karttunen's example (5b), repeated here as (26a):

(26) a. Harvey courts a girl$_i$ at every convention. She$_i$ always comes to the banquet with him.
b. [x: Harvey(x), [e: convention(e)]$\langle\forall$e$\rangle$[y: girl(y), x courts y at e]]

As pointed out by Karttunen, "a girl" in (26a) allows for a specific and a non-specific reading. On the specific construal, the first sentence entails that there is a particular girl who is courted by Harvey at every convention; this reading is unproblematic, so we can set it aside. On its non-specific construal, the indefinite is in the scope of the universal quantifier, which we may represent as in (26b). The problem with this reading is that the lifespan of the discourse referent y is bounded by the scope of "every convention", so there is no way the anaphoric pronoun in the second sentence could have access to it. Sometimes, this is what we want:

(27) Harvey courts a girl$_i$ at every convention. She$_i$ is called "Jackie".

Our current version of DRT predicts that this can only be read with a specific construal of "a girl", which is correct, but by the same token it prohibits a non-specific construal of the indefinite in (26a), which is not correct.

The same problem arises with other scope-bearing expressions, like modals, for example:

(28) a. Wilma may have bought a car.
b. [x: Wilma(x), $\Diamond$[y: car(y), x bought y]]

In order to account for sentences like (28a), we have to modalise the DRS language. The simplest way of doing this is by introducing the standard one-place modal operators, as illustrated by (28b). Extending the DRS semantics so as to interpret these structures is straightforward, too: we add to our models sets of worlds and accessibility relations, and interpret DRSs relative to embedding functions and worlds. The semantics is just as one would expect; the clause for $\Diamond$-conditions comes out as follows:

> f verifies $\Diamond$K at world w iff, for some world w$'$ accessible from w, there is an embedding function g such that f[K]g and g verifies K at w$'$.

This predicts that modals constrain anaphora just like negation does: a discourse referent introduced within the scope of a modal α can only be picked up by a pronoun if the pronoun occurs in the scope of α, too. As we saw in Section 1, this is only partly right:

(29) Wilma may have bought a car$_i$.
   a. *It$_i$'s a Volkswagen.
   b. It$_i$ may be a Volkswagen.

Assuming that, in this case, a specific construal of "a car" is not feasible, anaphora is possible in (29b) but not in (29a). As it stands, DRT doesn't account for this contrast, nor does any other dynamic theory of interpretation, for that matter.

Broadly speaking, there are three ways of dealing with the problems posed by piggyback pronouns. I call them the "inferential" model, the "resumption" model, and the "mixed" model. Between these, the resumption model is no doubt predominant, while the mixed model is probably the least known. In the following I will discuss each of these models in turn, and argue that the mixed model deserves to be taken more seriously than it has been thus far.

The clearest example of the inferential approach to piggyback anaphora is Roberts's (1987, 1989) work on what she calls "modal subordination", i.e. piggyback anaphora in modal environments. On Roberts's view, modal subordination is enabled by pragmatic inferences that constrain the domain of possible worlds a modal expression quantifies over. In the case of (30a), pragmatic inferences constrain the domain of the second modal, as shown in the transition from (30b) to (30c), as a result of which the neuter pronoun gets access to a suitable antecedent. (Here and in the following inferred material is marked in boldface.)

(30) a. Wilma may have bought a car$_i$. It$_i$ may be a Volkswagen.
   b. [x: Wilma(x),
      ◇[y: car(y), x bought y],
      ◇[<u>z</u>: VW(z)]]
   c. [x: Wilma(x),
      ◇[y: car(y), x bought y],
      ◇[z, **y**: **car(y)**, **x bought y**, VW(z), z = y]]

The main problem with Roberts's account is that, in a sense, it is too pragmatic: by leaving essentially everything to pragmatics, Roberts ends up with an explanation that is insufficiently constrained. For example, there is nothing in Roberts's theory that rules out sequences like the following:

(31) Wilma may have bought a car$_i$. *Fred can drive it$_i$, too.

For further discussion of Roberts's take on modal subordination, see Geurts (1995, 1999) and Frank (1997).

As observed in Section 1, the problematic anaphor in (30a) and similar examples would seem to be enabled by the circumstance that the anaphoric expression occurs in the scope of an expression that somehow extends the mode of quantification of the first sentence (whence the term "piggyback anaphors"). This at any rate is the intuitive idea underlying the resumption model, which is instantiated in a range of otherwise very different theories proposed since the mid-1990s by Kibble (1994), Geurts (1995, 1999), Krifka (1996), Frank (1997), Frank and Kamp (1997), van Rooij (2005), Asher and McCready (2007), and others. (The first exponent of this view may be a pre-publication version of Groenendijk and Stokhof's 1990 paper, discussed by Roberts (1995), which presents an analysis of piggyback anaphora in modal contexts that is very similar to later proposals by Kibble, Geurts, and others. However, this analysis didn't make it into the published article.)

Consider the discourse in (30a), and suppose that the first sentence introduces a set of possible worlds in which Wilma bought a Volkswagen. Suppose, furthermore, that the modal in the second sentence quantifies over the same worlds. Then each of the worlds at which the pronoun is evaluated contains a suitable object for "it" to refer to. The same, mutatis mutandis, for piggyback anaphors generally. Unfortunately, things aren't quite as simple as this: in order to for the resumption scheme to work, track will have to be kept of which VWs Wilma bought in which worlds: the discourse referent standing in for the neuter pronoun will have to find a discourse referent introduced by the indefinite NP in the scope of the first modal. Typically, though not invariably, this is done by pairing worlds with embedding functions. The same holds for other forms of quantification:

(32) Last week, all employees received a letter$_i$. Most of them read it$_i$ rightaway.

On a resumption theory, we have to make sure that, by the time the second sentence is being interpreted, we still know which employee received which letter; which usually is done by pairing individiuals in the domain of quantification with embedding functions. The upshot of all this is that the resumption theories of piggyback anaphora invariably run into considerable technical rigmaroles: without exception, the theories listed above are tremendously complex. For example, Krifka's (1996) definition of what he calls "parametrised sum individuals", which are pairs of ordinary individuals and assignment functions (i.e. embedding functions), takes up no fewer than eight clauses, and while other theories may be

very different from Krifka's, all of them involve technical apparatus that is at least as complex. And, crucially, none of these complexities is justified on independent grounds; they are needed solely for dealing with piggyback anaphora. It seems safe to conclude, therefore, that even if the core intuition underlying the resumption model is quite straightforward, its implementation is anything but.

Another problem with the resumption model is empirical rather than methodological; it is illustrated by the following discourse:

(33) Last week, all employees received a letter$_i$. Most of them read it$_i$ rightaway, except for Jones, who lost it$_i$.

Intuitively, the two occurrences of "it" in (33) should be interpreted analogously, but the resumption approach entails that they are fundamentally different: whereas the first occurrence of "it" is a run-of-the-mill piggyback anaphor licensed by the quantifier "most", there is no overt expression to license the second occurrence of "it", and unlike some of the cases discussed in Section 1, there doesn't seem to be a covert quantifier, either (cf. the examples in (6)). Hence, whatever the true story about the second "it" may be, the resumption view entails that the only thing the second "it" has in common with the first is that they share the same antecedent; otherwise they are bound to function in entirely different ways. Which doesn't seem very likely.

6. Piggyback anaphora and bridging

Even if these problems aren't fatal, they are serious enough not to be too happy with the resumption model. Still, I believe the basic intuition that drives this approach is sound. What causes the trouble, in my view, is an assumption that is never made explicit, let alone defended. To explain, consider the constellation in which piggyback anaphors arise:

$$\text{Q}[\ldots\ \text{Indefinite}_i\ \ldots] \quad \ldots \quad \text{Q}'[\ldots\ \text{Anaphor}_i\ \ldots]$$

(Where Q and Q′ are expressions that quantify over individuals, worlds, and so on.) The tacit assumption is that giving Q′ the right sort of interpretation is *sufficient* for enabling the anaphoric link into the scope of Q. In DRT terms, it is presupposed that once Q′ has been linked to Q, the discourse referent introduced by Indefinite$_i$ is *ipso facto* accessible to the discourse referent associated with Anaphor$_i$, and therefore the two can be equated. This is the reason why embedding functions are being incorporated into the semantic values of plurals, modals, and so on.

Once the assumption has been exposed it is evident, I trust, that there are no compelling reasons for sticking to it. So let's give it up and see what our options are. What I would like to suggest in its stead, revamping a proposal by Kamp and Reyle (1993), is that piggyback anaphora always involves a *bridging inference*, which is licensed by the interpretation of Q′ in the scheme above. In fact, once Q′ has received the right interpretation, the bridging inference is a *logical* one; it is an entailment. I call this a "mixed" model, because it supposes that piggyback anaphora requires resumption as well as an inferential element.

To show how this will work, I'm going to discuss (33) sentence by sentence:

(34) a. All employees received a letter.
b. [$\underline{X}$, Y: all(X,Y),
employees(X), Y = X ∩ x̂[x, y: letter(y), x received y]]

X and Y are discourse referents representing groups of individuals, which we will model as sets. Alternatively, we might interpret them in terms of mereological sums, for example, as long as it is understood that they don't contain embedding functions. So "all" in "all(X,Y)" is a perfectly ordinary quantifier, whose interpretation is as simple as it gets:

- f verifies all(X,Y) iff f(X) ⊆ f(Y)
- f verifies some(X,Y) iff f(X) ∩ f(Y) ≠ ∅

etc.

The condition "employees(X)" in (34b) is to mean that all individuals in X are employees, and "Y = X ∩ x̂K" defines Y as that subgroup of X whose members have the property of being an x such that K. This condition contains two complex discourse referents, which may be interpreted by extending embedding functions as follows:

For any embedding function f:

- f(X ∩ Y) = f(X) ∩ f(Y)
- f(x̂K) = {d ∈ D | there is a g such that f[K]g, g(x) = d, and g verifies K}

(It will not be hard to see that this analysis of quantification preempts the standard DRT treatment of quantified donkey sentences; we will come to that in the next section.) Finally, "all" being a quantifier, it presupposes its domain, but as (34a) just serves to get our example discourse started, we will ignore that for now.

The first conjunct of the second sentence in (33) is represented thus:

(35) a. Most of them read it rightaway.
b. [$\underline{Y'}$, Z: most(Y',Z), Z = Y' $\cap$ $\hat{z}$[z, $\underline{\underline{u}}$: z read u]]

(35b) contains two discourse referents that are presuppositional (or anaphoric, if you prefer): Y′ and u. If we merge the DRSs (34b) and (35b), bind Y′ to Y, and simplify the result, we get:

(36) [X, Y, Z:
all(X,Y), employees(X), Y = X $\cap$ $\hat{x}$[x, y: letter(y), x received y],
most(Y,Z), Z = Y $\cap$ $\hat{z}$[z, $\underline{\underline{u}}$: z read u]]

In this DRS, u does not as yet have access to a suitable antecedent. However, since any eligible value of Z is a subset of some value of Y, (36) *entails* (37), where the difference between the two DRSs is marked in boldface:

(37) [X, Y, Z:
all(X,Y), employees(X), Y = X $\cap$ $\hat{x}$[x, y: letter(y), x received y],
most(Y,Z), Z = Y $\cap$ $\hat{z}$[z, $\underline{\underline{u}}$, **y**: **letter(y)**, **z received y**, z read u]]

Now u can be bound to its intended referent, and we end up with the DRS in (38), which gives us the reading we were looking for:

(38) [X, Y, Z:
all(X,Y), employees(X), Y = X $\cap$ $\hat{x}$[x, y: letter(y), x received y],
most(Y,Z), Z = Y $\cap$ $\hat{z}$[z, y: letter(y), z received y, z read y]]

In prose: there is a set of employees X (which is presupposed and therefore determined by the context) all of whom received a letter, there is a set Y which consists of those individuals in X that read the letter in question, and Y contains most individuals in X.

The key feature of this analysis, which generalises in a straightforward way to other cases of piggyback anaphora, is that it treats the problematic pronoun in term of bridging: once the quantifier in (35a) has been interpreted, we are entitled to infer within the scope of "most" (so to speak) that there is a letter, which then may be picked up by the pronoun. The interpretation of the quantifier plays a crucial role, but it is not the whole story.

This proposal has all the virtues and none of the vices of previous accounts. First, it is very simple, both in terms of the representational machinery it requires and the procedures needed to associate a discourse with an adequate representation. Secondly, it has no need for esoteric se-

mantic types: plural discourse referents denote sets of individuals, propositional discourse referents denote sets of possibilities, and so on. Thirdly, the proposed account brings out the family resemblance between piggyback anaphora and other cases of anaphoric reference. The last sentence in (33), that is

(39) Jones lost it.

is a patent instance of bridging: assuming that Jones is one of the employees, the preceding discourse entails that he received a letter, which is the intended referent of "it". The difference with (35a) is merely that the bridging inference isn't enabled by a quantifying expression, but by an implicit premiss.

This bridging account of piggyback anaphors is based on the observation that Kamp and Reyle's (1993) treatment of their orphanage examples (see (6d), Section 1) generalises to all kinds of piggyback anapora. This type of treatment has provoked harsh criticisms, which typically fasten on the fact that information gets copied from one part of a DRS to another:

> A more principled objection to the illustrated treatment of plural anaphora concerns the power of the rules involved. The DRS construction rules are virtually unconstrained re-writing rules. […] There is nothing that would restrict the type of copying operations that are possible. (Krifka 1996: 560)

Although it has to be said that the way Kamp and Reyle present their analysis may have fostered such misgivings, they are beside the point. For, as I have been at pains to emphasise in my own presentation, the copy-and-paste "construction rules" used by this theory are in fact rules of inference that are sound, in the logical sense of the word, with respect to the semantics of the DRS language. They are, in other words, *highly* constrained.

Finally, it may be noted that Krifka's copy-and-paste objection does apply, and has been applied, to Roberts's (1987, 1989) theory of modal subordination. As we saw in the last section, Roberts holds that piggyback anaphora is chiefly dependent on pragmatic inferences, and this is precisely why it is insufficiently constrained. The mixed model of piggyback anaphora doesn't suffer from this defect.

7. Rethinking accessibility

The potential ramifications of this theory go beyond the analysis of piggyback anaphora. For, if my proposal is on the right track, there are reasons

for fearing that, in some respects at least, DRT's treatment of accessibility may not be. Consider the standard semantic clause for universal quantification, as given in Section 3:

> f verifies K⟨∀x⟩K′ iff, for all individuals d ∈ D, if there is a g such that f[K]g, g(x) = d, and g verifies K, then there is an h such that g[K′]h and h verifies K′.

On reflection, this definition is not as unproblematic as it may seem to be at first. For one thing, considering how simple the intuitive notion of universal quantification is, the standard DRT treatment is alarmingly complex. For another, this definition *stipulates* that the restrictor of the universal quantifier is accessible to its nuclear scope. Hence, the keystone of DRT's celebrated account of donkey sentences like (40) is arguably ad hoc.

(40) Every farmer who owns a donkey$_i$ beats it$_i$. (= (16a))

In other words, classical DRT suffers from the same problems that theories of piggyback anaphora run into (Section 5); the same goes for other members of the dynamic semantics family (cf. Schlenker's 2007 recent critique).

In the last section I proposed the following definition of universal quantification:

> f verifies all(X,Y) iff f(X) ⊆ f(Y)

This is very simple and intuitive and as far as I can tell it is not stipulative in any way. However, if we adopt it the standard account of donkey sentences goes by the board, because the restrictor of the universal quantifier (or any other quantifier) is no longer accessible from its nuclear scope; so the indefinite in (40) cannot bind the pronoun, as it can in DRT and other dynamic theories.

In the version of DRT presented in the last section, (40) is represented as follows (ignoring the domain presupposition associated with the quantifier):

(41) [X, Y: all(X,Y),
X = x̂[x, y: farmer(x), donkey(y), x owns y],
Y = X ∩ x̂[x, u: x beats u]]

In this representation, the anaphoric discourse referent u does not have access to its "target", y. However, using the same reasoning as before, (41) entails:

(42) [X, Y: all(X,Y),
X = x̂[x, y: farmer(x), donkey(y), x owns y],
Y = X ∩ x̂[x, u, **y**: **farmer(x)**, **donkey(y)**, **x owns y**, x beats u]]

Now u can be equated to y and we obtain:

(43) [X, Y: all(X,Y),
X = x̂[x, y: farmer(x), donkey(y), x owns y],
Y = X ∩ x̂[x, y: farmer(x), donkey(y), x owns y, x beats y]]

This says that every farmer who owns a donkey is a farmer who owns and beats a donkey; which is the reading to be accounted for. Hence, it is natural, and more economical, to explain anaphora in sentences like (40) in terms of bridging, and jettison what has become the standard account in DRT and related theories.

It is clear that the analysis I envisage for (40) will apply to any form of quantification, be it over individuals, worlds, or times; so it extends to nominal, adverbial, and modal quantification, as well as conditionals. More precisely, what I have outlined is a general explanation of anaphora from the nuclear scope of a quantifying expression into its restrictor.

As noted at the end of Section 3, there are two features that distinguish DRT from syntax-based theories of binding. One is that DRT enables indefinites sitting in the restrictor of a quantifying expression to bind anaphors in its scope. The other is that DRT makes it possible for an indefinite to bind an anaphor across "and" or a sentence boundary. If we adopt a bridging account of donkey anaphora, we lose the first feature. What about the second?

Sad to say, with the advent of dynamic theories of interpretation, the semantics of "and" has become undeservedly moot, and in this point my own views go against the dynamic current, but since I don't have much space left, the following remarks will have to be brief. First, it is not too difficult to incorporate sentential conjunction in DRT. We simply introduce conjoined DRSs of the form $K \wedge K'$, as follows:

$K \wedge K'$ is a DRS, where
– $U_{K \wedge K'} = U_K \cup U_{K'}$
– $Con_{K \wedge K'} = Con_K \cup Con_{K'}$

So (44a) is represented as (44b), which (save for the fact that the possessive pronoun gets glossed over) has the right meaning. Most importantly, for our purposes, the anaphoric link between "a Toyota" in the first conjunct and "it" in the second is nicely accounted for:

(44) a. Our vicar owns a Toyota$_i$ and it$_i$ is grey.
b. [x, y: vicar(x), car(y), x owns y] $\wedge$ [: grey(y)]

This result is obtained by inserting an old-fashioned form of conjunction into a dynamic framework. It is equivalent to treating conjunction by merging DRSs, as it is usually done (though this clear and simple idea has unfortunately been muddied, but I'm coming to that), for (44b) has the same truth conditions as:

(44) b.$'$ [x, y: vicar(x), car(y), x owns y, grey(y)]

What this analysis doesn't provide is an explanation of the contrast between (44a) and (45):

(45) *It$_i$ is grey and our vicar owns a Toyota$_i$.

It has become de rigueur, in dynamic circles, to claim that this contrast must be accounted for by hard-wiring it into the lexical entry of "and". Against this view, I have argued that the observed contrast is purely pragmatic, and that there is no good reason for doubting the conventional wisdom that "and" is commutative (Geurts 1999). So I'm going to be brazenly conservative and assume that the definition of conjunction given above is correct.

Now my point is, quite simply, that this treatment of conjunction is as frugal as it can be and that it seems most unlikely that we could come up with an alternative story in terms of bridging that explains cross-conjunct anaphora and is more economical than the conservative account given above. In brief, the second feature of the dynamic analysis of discourse referents stands unchallenged: if there is a distinctive feature of DRT and related theories, it is this.

8. Conclusion

It has long since been recognised that the interpretation of non-lexical definites, like "the president" or "the carburetor", may involve so-called bridging inferences, which anchor the newly introduced president or carburetor to a suitable object in the preceding discourse. Although, to the

best of my knowledge, it has never been denied that the same holds for anaphoric pronouns, I believe it has been widely assumed that bridging interpretations of pronouns are of marginal relevance, at best. I have argued that that this assumption is mistaken. It may be the case that pronouns are harder to interpret by way of bridging than non-lexical definites, but bridging construals of pronouns aren't hard to find, either. Once this is appreciated, there is no reason why bridging shouldn't play a key role in piggyback and donkey anaphora.

The paramount advantage of introducing bridging into the analysis pronominal anaphors is that it simplifies the dynamic treatment of anaphora, and simplifies it hugely, especially in the case of piggyback anaphora. However, there are empirical and conceptual benefits as well. On the empirical side, the bridging analysis captures the similarity between the two neuter pronouns in examples like (46):

(46) Last week, all employees received a letter$_i$. Most of them read it$_i$ rightaway, except for Jones, who lost it$_i$. (= (33))

On the conceptual side, bridging allows us to stick to a repertoire of semantic entities that are more faithful to people's semantic intuitions: instead of having to assume that modal expressions quantify over pairs of possible worlds and embedding functions, we can revert to the simpler, and intuitively more palatable, view that they quantify over worlds. Similarly, there is no need to suppose that nominal quantifiers range over pairs of individuals and embedding functions: they range over individuals, as they should. In sum, the advantages of the bridging analysis, simple as it is, are considerable.

My final remark concerns the portability of the analysis presented in the foregoing pages. One of the reasons why I have used DRT is that it is rather easy to incorporate bridging into a DRT framework, and it will be obvious that a framework that cannot accommodate bridging inferences would have been unsuitable for my purposes. Now, it is not unlikely that some dynamic theories of interpretation will run into considerable technical trouble when they try to incorporate bridging inferences, and for this reason may find it hard, if not impossible, to accommodate a bridging analysis of pronouns. But surely, it cannot be held against the proposed analysis that such theories exist.

Acknowledgements

Portions of the material in this article were presented to audiences in Amsterdam, Cambridge (the transatlantic one), Paris, and Seoul, to all of whom I am indebted for discussion. I'm especially grateful to Paul Portner for his extensive comments on an earlier version of this paper.

References

Asher, N. and E. McCready (2007). *Were*, *would*, *might* and a compositional account of counterfactuals. *Journal of semantics* 24: 93–129.

Beaver, D. (2001). *Presupposition and assertion in dynamic semantics*. Stanford: CSLI.

Bolinger, D. (1977). Pronouns and repeated nouns. Indiana University Linguistics Club.

Bosch, P. (1983). *Agreement and anaphora*. London: Academic Press.

Elbourne, P. D. (2005). *Situations and individuals*. Cambridge, Massachusetts: MIT Press.

Evans, G. (1977). Pronouns, quantifiers, and relative clauses (I). *Canadian journal of philosophy* 7: 467–536.

Frank, A. (1997). *Context dependence in modal constructions*. Doctoral dissertation, University of Stuttgart.

Frank, A. and H. Kamp (1997). On context dependence in modal constructions. In A. Lawson (Ed.), *Proceedings of SALT 7*, Ithaca, New York, pp. 151–168. Cornell University.

Gazdar, G. (1979). *Pragmatics: implicature, presupposition, and logical form*. New York: Academic Press.

Geurts, B. (1995). *Presupposing*. Doctoral dissertation, University of Stuttgart.

Geurts, B. (1999). *Presuppositions and pronouns*. Oxford: Elsevier.

Geurts, B. (2002). Donkey business. *Linguistics and philosophy* 25: 129–156.

Geurts, B. and D. Beaver (2007). Discourse representation theory. *Stanford encyclopedia of philosophy*.

Groenendijk, J. and M. Stokhof (1990). Dynamic Montague Grammar. In L. Kálmán and L. Pólos (Eds.), *Papers from the second symposium on logic and language*, pp. 3–48. Budapest: Akadémiai Kiadó.

Groenendijk, J. and M. Stokhof (1991). Dynamic predicate logic. *Linguistics and philosophy* 14: 39–100.

Haddock, N. J. (1987). Incremental interpretation and combinatory categorial grammar. In *Proceedings of IJCAI-87*.

Haviland, S. and H. H. Clark (1974). What's new? acquiring new information as a process in comprehension. *Journal of verbal learning and verbal behavior* 13: 512–521.

Heim, I. (1982). *The semantics of definite and indefinite noun phrases*. Ph.D. thesis,

University of Massachusetts at Amherst.
Heim, I. (1990). E-type pronouns and donkey anaphora. *Linguistics and philosophy* 13: 137–178.
Heim, I. (1992). Presupposition projection and the semantics of attitude verbs. *Journal of semantics* 9: 183–221.
Kamp, H. (1981). A theory of truth and semantic representation. In J. Groenendijk, T. Janssen, and M. Stokhof (Eds.), *Formal methods in the study of language*, pp. 277–322. Amsterdam: Mathematical Centre Tracts 135.
Kamp, H. and U. Reyle (1993). *From discourse to logic*. Dordrecht: Kluwer.
Kamp, H., J. van Genabith, and U. Reyle (2004). Discourse representation theory. Unpublished paper, University of Stuttgart.
Karttunen, L. (1971). Discourse referents. Indiana University Linguistics Club. Reprinted in: McCawley (**?**: 363-385).
Karttunen, L. (1973). Presuppositions of compound sentences. *Linguistic inquiry* 4: 167–193.
Kibble, R. (1994). Dynamics of epistemic modality and anaphora. In *Proceedings of the international workshop on computational semantics*. Tilburg Univeristy.
Krifka, M. (1996). Parametrized sum individuals for plural anaphora. *Linguistics and philosophy* 19: 555–598.
Lewis, D. K. (1975). Adverbs of quantification. In E. L. Keenan (Ed.), *Formal semantics of natural language*, pp. 3–15. Cambridge University Press.
Lewis, D. K. (1979). Scorekeeping in a language game. *Journal of philosophical logic* 8: 339–359.
McCawley, J. D. (1981). *Everything that linguists have always wanted to know about logic but were ashamed to ask*. Oxford: Blackwell.
Muskens, R. (1996). Combining Montague semantics and discourse representation. *Linguistics and philosophy* 19: 143–186.
Neale, S. (1990). *Descriptions*. Cambridge, Massachusetts: MIT Press.
Roberts, C. (1987). *Modal subordination, anaphora and distributivity*. Ph.D. thesis, University of Massachusetts at Amherst.
Roberts, C. (1989). Modal subordination and pronominal anaphora in discourse. *Linguistics and philosophy* 12: 683–721.
Roberts, C. (1995). Domain restriction in dynamic semantics. In E. Bach, E. Jelinek, A. Kratzer, and B. H. Partee (Eds.), *Quantification in natural languages*, Volume 2, pp. 661–700. Kluwer.
Russell, B. (1905). On denoting. *Mind* 14: 479–493.
Schlenker, P. (2007). Anti-dynamics: presupposition projection without dynamic semantics. *Journal of logic, language and information* 16: 325–356.
Sells, P. (1985). Restrictive and non-restrictive modification. Technical Report 85-28, CSLI, Stanford.
Stenning, K. (1978). Anaphora as an approach to pragmatics. In M. Halle, J. Bresnan, and G. A. Miller (Eds.), *Linguistic theory and psychological reality*, pp. 162–200. Cambridge, Massachusetts: MIT Press.

Tic Douloureux, P. (1971). A note on one's privates. In A. Zwicky, P. Salus, R. Binnick, and A. Vanek (Eds.), *Studies out in left field: defamatory essays presented to James D. McCawley on the occasion of his 33rd or 34th birthday*, pp. 45–51. Indiana University Linguistics Club.

van Rooij, R. (2005). A modal analysis of presupposition and modal subordination. *Journal of semantics* 22: 281–305.

Veltman, F. (1996). Defaults in update semantics. *Journal of philosophical logic* 25: 221–261.

Yule, G. (1982). Interpreting anaphora without identifying reference. *Journal of semantics* 1: 315–322.