Craige Roberts
Lectures at the Universidade Federal de São Carlos, Brazil
September 19, 2017; revised 6/25/2023

Anaphora in Discourse

Lecture 2:
Donkey Anaphora and Dynamic Interpretation

Outline for Lecture 2:

4. Donkey anaphora and dynamic interpretation

A real turning point in our understanding of anaphora in discourse occurred when Irene Heim and Hans Kamp published their seminal work on donkey anaphora and dynamic interpretation in the early 1980s. Each developed their work independently and their theories are superficially different; but at a more abstract level they share common features which make them much the same. This work has had ramifications well beyond the study of anaphora, and forms the basis of some of the most sophisticated formal semantic theories today. And Heim's work has been extremely influential in the study of presupposition and presupposition projection. Even though the original theory proved to have empirical problems, its central insights are still important, shedding light on the way that discourse is structured and interacts with the compositionally determined meanings of individual clauses and sentences.

We'll start with a look at the classical puzzle that Heim and Kamp tackled, and then turn to their proposals.

4.1 The problems posed by donkey anaphora[1]

Here's a classic puzzle about anaphora: the so-called donkey sentences noticed by the logician Geach (1967) (which he attributed to scholars in the Middle Ages). There are two main types of donkey sentence, exemplified in (1) and (2):

[1]The first chapter of Heim (1982) offers an excellent discussion of the problems posed by donkey anaphora, and a detailed critique of earlier attempts to account for them. I highly recommend it.

(1) Every farmer that owns a donkey beats it.
(2) If a farmer owns a donkey, he beats it.

Here I have indicated intended anaphoric relations informally with color highlighting. The first type of donkey sentence (1) has an indefinite antecedent inside a relative clause (RC) that modifies a quantificational subject NP, with the pronoun in the VP. The second type (2) has an antecedent inside the *if*-clause of a conditional, with a pronoun in the main clause. Here, for simplicity, we'll just focus on the green highlighted NPs; the same problems arise for the yellow.

Such examples pose two difficult problems for a theory of scope and anaphora:

Problem 1: **scoping out of islands**

The apparent antecedent in (1), *a donkey*, is in a RC, and, as we saw in Rodman's (1976) examples discussed in Lecture One §3.2.3, RCs are scope islands. It is clear that a QP with the determiner *every* cannot scope out of a RC island to bind a pronoun in the VP, as illustrated with (3), which is structurally identical to (1):

(3) #A farmer that owns every donkey beats it.

Note that I have captured the unacceptability of (3) not with the '*' of ungrammaticality, but with '#', indicating infelicity. None of the Principles of Binding Theory would preclude coindexing the two underlined NPs. But since *every donkey* is not referential, the two NPs cannot be coreferential. And because *every donkey* is in a scope island (§3.2.3 above), it cannot raise to take the whole sentence as its scope. But as we saw in §3.2.1 (e.g., example (76)), quantificational binding of a pronoun requires that the pronoun be in the scope of the quantifier. Thus, in (3), the pronoun *it* is free, whatever its index. The diacritic judgment '#' then indicates that the reading where we co-vary the interpretation of *it* with that of the domain of *every donkey* (informally indicated by underlining the two NPs) is not available. In other words, this is not a way of satisfying the pronoun's anaphoric presupposition. Since presupposition satisfaction is a pragmatic phenomenon, the result is infelicitous: unacceptable even though grammatical.

The same thing can be said of the antecedent of the conditional in (2); as we saw in §3.2.3, such tensed clauses are also scope islands.

Though indefinites like *a donkey* are classically treated as QPs, with the same semantics as we saw for *some* NPs in Lecture One, there is a difference between indefinites and other QPs, including *every donkey*. Unlike other QPs, the indefinites are known to sometimes be available to serve as antecedents for pronouns cross-sententially, violating ROOFING (§3.3). Again, compare:

(4) Every donkey brayed when the noon whistle went off.
#It had a loud voice.
(5) A donkey brayed when the noon whistle went off.
It had a loud voice.

Might the indefinite in examples like (5) be referential rather than quantificational? And might this then not explain how the indefinite serves as antecedent in donkey sentences like (1) and (2)? That's not a *prima facie* unreasonable hypothesis, but it won't explain the donkey anaphora either. That's because treating *a donkey* in those examples as referential would not yield their attested interpretations. For in (1) and (2) it's not that there's a particular donkey such that every farmer owns it, or that if a farmer owns that donkey there'll be certain consequences. Instead, the interpretation of *a donkey* co-varies with that of *every farmer* in (1), or with the arbitrarily chosen donkey-owning situation in (2). There is no particular donkey being referred to. So we can neither say that *a donkey* is referential, nor that it's a QP but just somehow takes wide scope out of the island.

This conclusion is reinforced by the observation that donkey indefinites cannot generally serve as antecedents for pronouns in subsequent sentences, illustrated here both with continuations of our original donkey sentences and with one other example introduced to show that there's nothing special about *every* in (1): (8) is a donkey sentence with the quantificational determiner *no*.

(6) Every farmer that owns <u>a donkey</u> beats <u>it</u>.
#It looked at me with sad eyes.

(7) If a farmer owns <u>a donkey</u>, he beats <u>it</u>.
#It suffered terribly last week.

(8) No frog that saw <u>an insect</u> ate <u>it</u>.
#It was a big fly.

Thus, in cases like (6)-(8) the indefinite under the scope of *every*, the conditional, or *no* does obey ROOFING. We conclude that donkey anaphora is neither a question of coreference, nor of quantificational binding. So what is the nature of this beast?

<u>Problem 2</u>: **wrong quantificational force**

We've ruled out a referential interpretation for *a donkey*. Suppose we give it the existential force usually associated with indefinites, like QPs headed by *some*, repeated here from §3.2 along with (9) that exemplified it:

(9) Some child smiled.
some: a relation between two sets, its domain and its scope
in (9): domain = the set of children $|child|$
scope = the set of entities that smiled $|smiled|$
the relation, in set-theoretic terms: the non-null intersection between domain and scope
in (9): $|child| \cap |smiled| \neq \varnothing$ 'the intersection of domain and scope is not the empty set', i.e. there's some entity that's in both

I call this interpretation existential because, as we see in the highlighted passage, its truth requires the existence of some entity that has both the domain and scope properties.

Assuming this interpretation, we just saw that there is no specific donkey at issue in the donkey sentences, ruling out the reading in (1′a) for (1).

(1) Every farmer that owns a donkey beats it.

(1′) (a) wide scope existential: Not Available!

formally: $\exists y(\text{donkey}'(y) \ \& \ \forall x[\text{farmer}'(x) \ \& \ \text{owns}'(x,y) \rightarrow \text{beats}'(x,y)]$

informally: 'there's a donkey such that every farmer that owns it beats it'

But we also saw that we cannot just interpret the indefinite *in situ*, since there it cannot move outside its island scope to bind the pronoun. We would get a logical form like that in (1′b):

(1′) (b) *In situ* existential: Not the attested reading!

formally: $\forall x[(\text{farmer}'(x) \ \& \ \exists y\underline{[\text{donkey}'(y) \ \& \ \text{owns}'(x,y)]}) \rightarrow \text{beats}'(x,y)]$

informally: 'for every farmer such that there's a donkey that that farmer owns, that farmer beats y'

Here, the scope of the indefinite is the underlined constituent, so that the existential doesn't bind the highlighted *y*. But our intuitions tell us that *a donkey* is clearly the antecedent of *it*, the two co-varying with the chosen farmer.

In fact, the truth conditions we want for (1) are more like the reading in (1′c):

(1′) (c) double universal reading:

formally: $\forall x\underline{\forall y}[\text{farmer}'(x) \ \& \ \underline{\text{donkey}'(y)} \ \& \ \text{owns}'(x,y) \rightarrow \text{beats}'(x,y)']$

informally: 'for every pair consisting of a farmer and a donkey that he owns, the farmer in the pair beats the donkey in the pair'

So, not only does the indefinite *a donkey* escape the RC island to act as antecedent for a pronoun that is outside the island, it suddenly acts like a universal quantifier, as we see in the underlined portions of (1′c)! This kind of behavior for an indefinite is only attested in donkey sentences, which led Kamp and Heim to suspect that it is the result not just of the meaning of the indefinite itself, but of the context in which it occurs in the donkey sentences and the role of that context in interpretation.

Since the truth conditional meaning for the conditional donkey sentence (2) seems to be the pretty nearly the same, it displays the same problems we just saw for the universal donkey sentence (1): Both *a donkey* and *a farmer* escape the antecedent of a conditional—a scope island, and then act as if they had the same force as *every*. Very strange behavior for an indefinite NP.

4.2 A first account of the donkey problem: Dynamic interpretation and unselective binding

To address the problems we just considered, Heim and Kamp proposed new theories of how interpretation proceeds in discourse. Heim's theory was called File Change Semantics (1982), or Context Change Semantics (1983). Kamp's (1981) was called Discourse Representation Theory (DRT). As we will see, they share four features that were crucial for solving those problems. Because of those shared features, we follow Heim (1990) and call them theories of dynamic interpretation. Because DRT is simpler to present in a relatively informal way than Heim's elegant and rigorous theory, I'll offer you a simplified version of the DRT account of donkey sentences. For more detail on DRT, see Kamp & Reyle (1993). For an excellent, careful introduction to both theories, with comparison, see Kadmon (2001).

We'll start by considering the simple two-sentence discourse in (10):

(10) (a) A farmer$_1$ owns a donkey$_2$.
(b) He$_3$ beats it$_4$.

The following is a Discourse Representation Structure (DRS) for (10a), constructed via a simple algorithm from the English surface sentential structure. The DRS is a kind of logical form, a translation from one language (English) to another (the language of DRSes), with the logical structure of the latter explicit and unambiguous:

DRS(10a):

x_1 x_2
farmer(x_1) donkey(x_2) owns(x_1,x_2)

The DRS has two parts separated by a horizontal line. In the bottom are a set of conditions: each noun and verb has become a predicate, and its argument is a variable whose index is given by that of the NP (for *a farmer* and *a donkey*) or by the indices of the arguments of the verb (for *owns*). In the top are the variables that are used in the conditions; Kamp calls this the Universe of the DRS, and calls the variables *reference markers*. But I'll use the term for these variables introduced by Karttunen (1976) and adopted by Heim (1982): these variables are **discourse referents**.

This is just a representation. To interpret it, we need an **embedding**. This is an assignment function *g* that maps all the variables in the universe of the DRS onto entities in the world (or a model of the world). A **truthful embedding** is one which assigns entities to the variables x_1 and x_2 in such a way as to make all the conditions in the bottom cell of the DRS true of those entities. Suppose we have an assignment function (just the sort we considered in §3.3.2) that assigns Alphie to x_1 and Giorgio to x_2. Then to see if the DRS (and hence (10a) itself) is true, we have to determine whether there is an assignment function which makes each of the three conditions true. So we check whether the value of x_i Alphie is a farmer, the value of x_2 Giorgio is a donkey, and Alphie owns Giorgio. If those conditions are true under the assignment, then DRS(10a) is true, and so is the example (10a) for which it captures the logical form.

In a simple, extensional model, we can think of such a representation as a crude first approximation to the context which results from the utterance of (10a) out of the blue. It puts a constraint on what assignment functions are licit in the discourse subsequent to that utterance. These are just those assignment functions which are constrained so as to truthfully embed DRS(10a). This filters out assignment functions which assign x_1 to non-farmers, or x_2 to aardvarks or carrots (assuming one can't be both an aardvark or a carrot *and* a donkey). In this framework, the addition of information corresponds to the addition of limitations on embedding, so that fewer assignment functions satisfy the embedding conditions.

A **discourse referent** associated with one of the variables in the universe of the DRS is technically just a way of capturing the contextually given constraints on the interpretation of any NPs coindexed with the variable: the way that DRS embedding works only permits using a given assignment function for interpreting the associated free variable (and hence, indirectly, the corresponding trace or pronoun in the surface structure of the utterance) if the assignment function respects the conditions on the coindexed discourse referent in the DRS.

In the DRS for (10a) we have two discourse referents, x_1 and x_2. (There is nothing essential about those particular referential indices; it just happens that we chose those indices for the NPs in the representation.) Functionally, the discourse referents in the universe of the DRS together with the conditions below them constrain on the permissible interpretations of utterances in subsequent discourse: only those assignments that truthfully embed (10a) may be used to interpret subsequent utterances. Then if those subsequent utterances contain any free NPs with the index *1* or *2*, the conditions they introduce to a DRS will receive an interpretation consistent with the information in (10a), as encoded in the permissible assignments. In other words, the discourse referents x_1 and x_2 and their associated conditions together tell us what assignment functions are available for interpreting free pronouns with the indices *1* or *2* in subsequent utterances. You might think of a discourse referent as a hyperlink that ties together bundles of information gathered across discourse. Another way of understanding discourse referents is that they are ways of encoding the constraints on the assignment of values to unbound variables that accumulate across discourse, so that the assignments which are felicitous are no longer arbitrary sky hooks, but really ways of keeping track of the information we've gathered, both about actual entities and (as we will soon see) about arbitrary instantiations of some quantificational domain.

Note that the true embedding of the DRS for (10a) doesn't require that the values of x_1 and x_2 in its universe be any particular farmers or donkeys; that is, they aren't referential in the usual sense. The DRS will be true just in case there's at least one such pair satisfying all the conditions; there might be many such pairs, and the truth of the DRS doesn't require picking out a particular pair. In this sense, the force of the indefinites which introduced the discourse referents is merely existential, non-specific.

So such a representation captures the truth conditions for the utterance in question: (10a) is **true** in a model *M* just in case there is a truthful embedding from (10a) into *M*. Such an embedding will have to find a farmer and a donkey who stand in the owning relation, guaranteeing an existential entailment associated with each of the discourse referents in the universe of the DRS, under the descriptions encoded in the conditions.

Now consider the DRS for (10b):

DRS(10b):

$x_3 \quad x_4$
$x_3 = ??$
$x_4 = ??$
beats(x_3,x_4)

The predicate '= ??' crudely indicates an anaphoric presupposition which must be resolved if the utterance is to be felicitous. But in (10) this string wasn't uttered out of the blue. We want the DRS in which it is represented to capture not only the conventional content of the constituent uttered, but its context of utterance as well, so as to permit us to resolve these presuppositions appropriately in *that* context. So suppose we merge DRS(10b) into the result of uttering the prior sentence, DRS(10a), taking the union of their individual universes to make the new universe, the union of their conditions for the new set of conditions:

DRS(10a+10b), step 1, unresolved presuppositions:

$x_1 \quad x_2 \quad x_3 \quad x_4$
farmer(x_1)
donkey(x_2)
owns(x_1,x_2)
$x_3 = ??$
$x_4 = ??$
beats(x_3,x_4)

Having entered the conventionally given content of (10b) into the DRS from its prior context, DRS(10a), we still have to resolve the pronominal presuppositions. But in this context there are already salient "entities"—the discourse referents x_1 and x_2. And so we can resolve these presuppositions as follows, since the pronominal gender and number of the pronouns and proposed antecedents is compatible and the resulting interpretation is a plausible and relevant follow-up to (10a) with the pronominal presuppositions so-resolved:

DRS(10a+10b), step 2, presuppositions resolved:

x_1 x_2 x_3 x_4
farmer(x_1) donkey(x_2) owns(x_1,x_2) $x_3 = x_1$ $x_4 = x_2$ beats(x_3,x_4)

Given the equivalence of some of the discourse referents stipulated in the highlighted conditions, this is truth conditionally equivalent (in terms of what worlds will yield permissible embeddings) to:

≈ DRS(10a+10b):

x_1 x_2
farmer(x_1) donkey(x_2) owns(x_1,x_2) beats(x_1,x_2)

Both the DRS in Step 2 and this equivalence can be truthfully embedded just in case there is a farmer and a donkey, such that two relations hold between them: the farmer owns the donkey and the farmer beats the donkey, precisely the intuitions native speakers report for (10).

With this sketch of how DRT works for a simple discourse, we can turn to the more interesting donkey sentence example in (1), *every farmer who owns a donkey beats it*. Here we have a quantificational determiner *every*. As we have done in previous sections, Heim and Kamp take such QPs to have the tripartite structure of quantification:

operator (**domain**, **scope**)

This will be reflected in the structure of a DRS for the example. The DRS for (1) will be a condition in the DRS representing its context. Call the pre-existing contextual DRS the matrix DRS. Then the result of adding the tripartite structure for (1) will have this general structure:

DRS Schema for tripartite quantificational structures:[2]

[universe of matrix DRS]

[conditions on embeddings for matrix DRS]

universe of domain	**Operator**	universe of scope
descriptive content of domain		descriptive content of scope

Let's start with a slightly simpler version of (1), (11):

(11) Every farmer$_1$ who owns a donkey$_2$ plows.

Assume that I utter (11) out of the blue, with no preceding linguistic context. Exemplifying the general schema to (11) we get:

DRS(11):

[universe of DRS(11)]

[Condition 1:]

x_1 x_2	$\forall$	
farmer(x_1) donkey(x_2) owns(x_1,x_2)		plows(x_1)

[2] We could do this tripartite structure in Polish notation, as well, like: Operator (Domain , Scope), with a box for each of the arguments.

In this DRS, the main universe is empty because the example is uttered out of the blue, and we have only one condition, representing the tripartite quantificational structure for (11) itself. How does such a representation bear on embedding the DRT? This DRS puts only one constraint on assignment functions, which is that they **verify** the tripartitite condition 1. Technically, we use the counterpart of Tarskian variation over variables bound by quantifiers.

> An embedding *g* **verifies** a universally quantified tripartite condition *c* just in case for every way of varying the values that the embedding assigns to elements in the universe of the domain (the first argument, or left-hand box) such that this results in a truthful embedding of the domain, there's an assignment *g′* that's just like *g* except possibly for the values of any elements in the universe of the scope (the second argument, or right-hand box) and such that *g′* truthfully embeds the scope.

Crudely paraphrasing: the universal operator tells us that **every way** of making its domain true is a way of making its scope true. The discourse referents in those arguments have no quantificational force of their own. Instead, they are unselectively bound by the universal operator over embeddings: it quantifies over ways of assigning values to those discourse referent variables.

Since there are no arguments in the scope of DRS(11), in this example verification requires of an assignment function *g* that if it truthfully embeds the domain, it also truthfully embeds the scope. For that to be the case, every donkey-owning-farmer will also be a plower, so that we replicate the truth conditions we saw for universally quantified *every* NPs in §3.2.

The requirement that we vary over any elements introduced in the scope comes into play in interpreting DRSes for examples which introduce new variables in the universe of the scope, as in (12):

(12) Every farmer$_1$ that owns a donkey$_2$ bought a plow$_3$.

DRS(12):

<table>
<tr><td>[universe of DRS(12)]</td></tr>
<tr><td>[Condition 1:]
<table>
<tr><td>x_1 x_2</td><td rowspan="2">∀</td><td>x_3</td></tr>
<tr><td>farmer(x_1)
donkey(x_2)
owns(x_1,x_2)</td><td>plow(x_3)
bought(x_1,x_3)</td></tr>
</table>
</td></tr>
</table>

The single tripartite condition in DRS(12) can only be verified if every assignment that verifies the domain, making x_1 a farmer, x_2 a donkey and the pair x_1 and x_2 stand in an owning relationship, can be modified so that it assigns x_3 to a plow that the farmer x_1 bought.

The null prior context in DRS(11) puts no constraints on assignment functions. The quantificational condition introduced by (11) just constrains what assignment functions are felicitous from this point on: all and only those such that *if* they assign values to x_1 and x_2 which are a farmer and a donkey respectively and such that the farmer owns the donkey, then the value for the farmer is also in the set of things that *plow*. This constraint on contextually felicitous assignment functions is not associated with a particular variable, as was the case with x_1 and x_2 in (10). Instead it rules out any assignments which fail the quantificational condition. Hence, DRS(11) does not yield an existence entailment—it might be only vacuously true, as when there are no farmers or donkeys who stand in the *owning* relation. If DRS(11) is truthfully embeddable and we do subsequently talk about a particular farmer, say Pierre, whom we learn has a donkey, then the context will entail that Pierre plows. So the condition does affect the contextually available information. But without an existential entailment, which would be represented by a variable in the universe of the main DRS, there is no discourse referent in the universe of the DRS, the context resulting from the utterance of (11).

Then suppose we follow up (11) with the utterance *He₃ beats it₄*. Interpretation would yield DRS(11′):

DRS(11′):

[universe of DRS(11′)]
x_3 x_4

[Condition 1:]

x_1 x_2		
farmer(x_1) donkey(x_2) owns(x_1,x_2)	$\forall$	plows(x_1)

x_3 = ??
x_4 = ??
beats(x_3,x_4)

But in the context DRS(11) we cannot resolve the highlighted presuppositions introduced by the pronouns, which would require appropriate discourse referents already in the universe of

DRS(11). The dRefs in the subordinate Condition 1 have a limited “lifespan”: They are only available antecedents for free variables under the scope of ∀. So DRS(11) offers no way of resolving the anaphora. So here the presuppositions *??* go unresolved, the resulting DRS is unembeddable as a consequence, and the discourse itself is infelicitous.

Now let’s apply what we have learned to building a representation for (1), repeated here:

(1) Every farmer$_1$ that owns a donkey$_2$ beats it$_3$.

Since, as we saw, *it* cannot be bound by its apparent antecedent *a donkey*, we have given it a new index. Then like the DRS for (11), that for (1) will introduce a new discourse referent into the universe of the scope of the universal operator:

DRS(1):

[universe of DRS(1)]

[Condition 1:]

x_1 x_2		x_3
farmer(x_1) donkey(x_2) owns(x_1,x_2)	∀	x_3 = ?? beats(x_1,x_3)

Here again we must resolve a pronominal familiarity presupposition. But in this case, something is different than in DRS(11′). In verifying the tripartite condition, any assignment *g* which truthfully embeds the domain argument (the left-hand box in the tripartite condition) will have done so by finding a farmer and a donkey owned by that farmer in the model. But in addition, there must be an assignment *g*′ which is just like *g* except possibly for the value of x_3 and which truthfully embeds the scope as well, so that the farmer, the value of x_1 (in both *g* and in *g*′) also beats something or other, the value of x_3. The value of that something or other is presupposed to be saliently entailed to exist. But at this point, the truthful embedding of the domain by *g*/*g*′ *does* guarantee such an entailment: it locally (under the scope of ∀) entails the existence of a donkey owned by the farmer. This (arbitrarily chosen) donkey of the (arbitrarily chosen) farmer is salient at that point, and so we can resolve the pronoun’s presupposition by equating x_3 with x_2, yielding the attested truth conditions:

DRS(1): Step two: anaphoric presupposition resolved

[universe of DRS(1)]

[Condition 1:]

x_1 x_2		x_3
farmer(x_1) donkey(x_2) owns(x_1,x_2)	$\forall$	$x_3 = x_2$ beats(x_1,x_3)

(1) will be true just in case there's a truthful embedding of DRS(1), which requires that the single condition be verified. Then given the equation of x_3 with x_2, this condition will be verified just in case for every way of picking out a farmer and a donkey that the farmer owns, the farmer beats the donkey. This is effectively the double-universal interpretation that we considered in (1′c) above.

We might say (using a process metaphor) that in the domain and scope of this universal operator there is *a temporarily available discourse referent* for a donkey. This value for x_2 in the scope of the operator is available by virtue of the semantics of the operator introduced by *every*, built into its verification conditions—this is what we get by requiring that the assignment verifying the scope be *just like that which verifies the domain, modulo the value of any new variables introduced in the scope*. But this way of interpreting the tripartite quantificational condition doesn't license anaphoric accessibility outside of the scope of the operator. After we verify the tripartite condition introduced by (1), the resulting DRS(1) as a whole has no discourse referents in the main universe, so no existential entailments, just as in DRS(11). So we cannot felicitously follow up with *it suffered terribly last night*, with *it* taking *a donkey* as its antecedent.

Crucially, in this approach neither the indefinite *a donkey* nor the pronoun *it* has any quantificational force of its own—in the DRS they are just treated as free variables in the universes of the sub-DRSes in which they are introduced, under the scope of the quantifier. The only force involved comes from the semantics of embedding itself, with the relationship between the assignments verifying the two arguments given by the quantificational operator. An utterance is true if **there is** an embedding for its DRS. A tripartite universal condition is verified if **every** way of embedding the domain is a way of embedding the scope as well, etc.

In the construction of these DRSes, we informally observed a distinction between indefinite and definite NPs. The indefinites introduce **novel** discourse referents, as we saw with *a donkey* in (10a) and (11) and *a plow* in (11). In contrast, definites, including pronouns, carry a **familiarity presupposition**, as we saw with *he* and *it* in (10b) and *it* in (1). This means that definites presuppose a "familiar" discourse referent, one already available in the DRS, with which they must be equated in order to resolve the presupposition. **It is this discourse referent which is technically the antecedent of the definite**. If the antecedent discourse referent is introduced by a referential expression (like a proper name, say), the definite will co-refer with it. But if the antecedent is indefinite (like *a farmer* or *a donkey* in (10), *a donkey* in (1)) or quantificational (like *every farmer* in (13)), the definite co-varies with the value of the antecedent:

(13) Every farmer$_1$ says he$_3$'s happy.

DRS(13): shown with the highlighted anaphoric presupposition of *he$_3$* resolved

[universe of DRS(13)]

[Condition 1:]

x_1		x_3
farmer(x_1)	$\forall$	say(x_1,happy(x_3)) $x_3 = x_1$

When an anaphoric presupposition is satisfied by a pre-existing discourse referent in the main universe of the DRS, as we saw in the interpretation of (10b) in the DRS (10a+10b), we say that it is **globally satisfied**. But when it is only satisfied by a discourse referent that's in the domain or scope of a quantificational operator, as in DRS(1) or DRS(13), we say that the anaphoric presupposition is **merely locally satisfied**. In the latter case, the antecedent isn't available for subsequent anaphoric antecedence after we leave the scope of the operator to which it is local. We saw this in the inaccessibility of *every farmer* and *a donkey* in (11) for subsequent anaphora with *He beats it*. In (13), we cannot follow up with *He told me so yesterday*, since this subsequent sentence is not under the scope of the universal operator in whose domain we found the merely local antecedent x_1 for *he$_3$*.

We might say that this kind of intuition underlies the approach to context sensitivity that is generally called **dynamic interpretation**. In it, context differs (or "changes") from constituent to constituent in the utterance. So in (1) or (13), the local context available for the interpretation of the operator's scope differs from the global context in which (1) was uttered. The scope of both of these operators has a "temporary" discourse referent x_1 for the farmer, and (1) also has x_2 for the donkey, neither of these globally available before or after the utterance. Informally, we might saw that the interpretation of an uttered expression is dynamic in case the context available for interpretation utterance-internal (here, for the scope of an operator) may differ from the global context in which it was uttered. In the examples like (13), the local context for the interpretation of *says he's happy* offers a discourse referent x_1 that is neither globally available in the prior context of utterance, nor globally available in the context resulting from interpretation of (13). Thus, we say that the discourse referent x_1 has a limited **life span**—the restriction and scope of the operator under which it is introduced. The context is updated in the course of interpretation, then again after completion.

Heim (1990) summaries the essential features of the dynamic DRT solution to the donkey sentence problem (including her own version in Heim's 1982) as follows:

1. The meaning of an utterance is an instruction for updating the context of utterance, a **Context Change Potential**. Technically, it tells us how to update the DRT (her File) representing the context of utterance, adding the new information conveyed in the utterance. This update may introduce new discourse referents and information about them, as in (10a) or the domain of (11). And it may include new information about pre-existing discourse referents, as in (10b) or the scope of (1). Further, it may introduce discourse referents globally or merely locally (under the scope of an operator).
2. What differentiates indefinite and definite NPs is that indefinites introduce **novel** discourse referents, while definites, including pronouns, carry a **familiarity presupposition**, a presupposition that there is an already familiar, co-indexed dRef in preceding context.
3. Definites and indefinites are treated as variables, with no quantificational force of their own.
4. **Unselective binding** (Lewis 1975) by a quantificational operator ($\forall$ in the examples above) accounts for the apparent quantificational force of an indefinite.

This approach offers an explanation for both the puzzles posed by donkey sentences.

First, it explains the scope puzzle: It isn't that *a donkey* semantically binds *it* in (1) or (2). Rather, the local context resulting from interpretation of the domain of *every* in (1) (or the *if*-clause in (2)) offers a way of giving a value to the free pronoun *it*: the constraint on admissible assignment functions introduced *merely locally* by the update of the global context with the operator's domain makes available the discourse referent corresponding to *a donkey* (x_1 in our DRS(1)) to serve as a constraint on assignment functions that interpret the operator's scope. If we equate the discourse referent for *it* with that for the familiar donkey discourse referent, we get the intended interpretation. The pronoun is free, its antecedent the locally available discourse referent.

Second, it is the unselective binding by a universal operator introduced by *every* which explains the apparent universal force of the indefinite *a donkey* that was the second part of the donkey puzzle.

We can't take the time to explain how the DRT approach captures the conditional-type donkey sentences like (2). Heim (1982) has the clearest, best semantically motivated exposition for those cases. Suffice it to say here that it involves the same kind of tripartite structure as we found in the universal donkey sentences, with the quantificational force contributed by the implicit modal force of the conditional construction itself and acting to unselectively bind both indefinites in the conditional's *if*-clause.

4.3 **The proportion problem for unselective binding**

While unselective binding gives the right truth conditions for examples where the force of the operator is universal, problems arise with donkey anaphora when the quantificational determiner is proportional, as with *most* in (14):

(14) Most farmers who own a donkey beat it.

Suppose that we construct the DRS for (14) in the same way as that for the universal donkey sentence DRS(1), replacing the universal operator with proportional *most*:

DRS(14): Step two: anaphoric presupposition resolved

[universe of DRS(14)]

[Condition 1:]

x_1 x_2		x_3
farmer(x_1) donkey(x_2) owns(x_1,x_2)	**MOST**	$x_3 = x_2$ beats(x_1,x_3)

And suppose that *most* requires (for example) that more than 60% of the ways of verifying its domain are ways of verifying its scope. Then, with unselective binding, DRS (14) is truthfully embeddable if more than 60% of the ways of picking out a pair of a farmer and a donkey that he owns also make it true that the chosen farmer beats his donkey.

Here's the problem that such proportional quantifiers pose: Consider a situation where there are a total of ten farmers. Nine of these farmers own one donkey each and treat them very kindly, never beating them. But the tenth guy owns ninety-one donkeys and beats every one of them. Then we have a total of one hundred different pairs consisting of a farmer and a donkey that he owns (ninety-one of which have the same farmer but differ in the donkey element). Since there's beating in ninety-one of these pairs (all involving the same mean farmer), (14) is predicted to be true in the circumstance described. But that does not match the intuitions of native speakers, who take it that the quantifier *most* should range not over the pairs, but over just the farmers: there are ten farmers, and for (14) to be true, six or more of those farmers would have to beat their donkeys. That is not the case, so the example should be false. Unselective binding fails to reflect the special status of the head (*farmer*) in the QP: this is what determines the size of the domain.

4.4 Existential readings of donkey sentences

Another problem for the unselective binding approach to donkey sentences comes from so-called *weak readings* of donkey sentences, as in the following example due to Pelletier & Schubert (1989):

(15) Every man who had <u>a quarter</u> in his pocket put <u>it</u> in the meter.
unselective interpretation: 'every man who had a quarter in his pocket put all the quarters in his pocket (whatever quarters he had) in the meter'
preferred reading: 'every many who had a quarter in his pocket put one of the quarters in his pocket in the meter'

These readings have been much-discussed in subsequent literature. see Barker (1993), Kang (1994), Kanazawa (1994), Chierchia (1995), Yoon (1996) Kanazawa (2001), Geurts (2002), Roberts (2005), Elbourne (2005), Brasoveanu (2006, 2007, 2008), and Champollion, Bumford & Henderson (2019), among others. The general consensus is that they cannot naturally be derived with unselective binding. But authors disagree about whether donkey sentences are ambiguous—with one weak and one strong reading—or unambiguously weak or unambiguously strong. Champollion et al. have a particularly interesting account wherein donkey sentences are underspecified rather than ambiguous, with pragmatic principles yielding particular interpretations in particular contexts, especially as a function of the Question Under Discussion (QUD).

Together, the proportion problem and the problem of existential readings have led to the abandonment of unselective binding as part of the dynamic semantics approach to donkey anaphora.

But this problem shouldn't lead us to throw out the baby with the bathwater: We will see that anaphora warrants the adoption of dynamic context update (though not necessarily a dynamic semantics) and that discourse referents generally have a role to play in these contexts, but not unselective binding.

<u>References</u>: See the course bibliography.