

Trivalence and anaphora

Matthew Mandelkern*

May 3, 2022

1 Introduction

This chapter is about the role of trivalence in theories of anaphora. The seminal puzzle here is the problem of *donkey anaphora*: how to account for the anaphoric dependency between ‘a cat’ and ‘her’/‘the cat’ in a sentence like (1):

- (1) Everyone who has a cat loves $\left\{ \begin{array}{c} \text{her} \\ \text{the cat} \end{array} \right\}$.

The idea that some kind of partiality is involved in the treatment of donkey anaphora is central to the dynamic treatment of the phenomenon (as in Kamp 1981; Heim 1982 and especially Krahmer and Muskens 1995). More recently, Rothschild (2017) has developed an explicitly trivalent treatment of anaphora in a static framework. I will start by explaining the puzzle; then I will explain the standard dynamic approach; Krahmer and Muskens’ bilateral dynamic semantics; and Rothschild’s static trivalent account. Finally, I will develop an alternative trivalent approach, based on the bivalent system in Mandelkern 2022b.

In the rest of the introduction, I will explain the problem of donkey anaphora. Readers familiar with the problem should jump to §2.

The foil here is what I’ll call the ‘classical picture’: the framework you learn in introductory logic for translating simple English sentences into sentences of predicate logic and then interpreting them. On this picture, a sentence with an indefinite noun phrase like ‘A cat sleeps’ has the logical form $\exists x(Cx \wedge Sx)$, where C stands for ‘cat’ and S for ‘sleeps’. ‘The cat sleeps’ gets a similar treatment, but with the addition of a uniqueness inference—that there is only one cat. There is controversy about how exactly this gets added: in particular, there is a divergence between Russell, on the one hand, and Frege/Strawson, on the other. Per Russell, the right schematization of this is $\exists x((Cx \wedge \forall y(Cy \rightarrow x = y)) \wedge Sx)$, where \rightarrow is the

*Department of Philosophy, NYU; mandelkern@nyu.edu. To appear in *Handbook of Three-Valued Logics*, Paul Egré and Lorenzo Rossi, eds. Thanks to seminar participants at UCL, NYU, and Oxford, and to Keny Chatain, Patrick Elliott, Matthew Gotham, Nathan Klinedinst, Daniel Rothschild, Philippe Schlenker, Benjamin Spector, and Yasu Sudo for very helpful discussion. Special thanks to Daniel Rothschild and Benjamin Spector for very helpful comments on this chapter which have greatly improved it.

material conditional. That is, ‘The cat sleeps’ says there is exactly one cat and it’s a sleeping cat. Frege/Strawson thought that uniqueness should be formulated as a presupposition, not part of the entailed content. So they would have ‘The cat sleeps’ presupposes $\exists x((Cx \wedge \forall y(Cy \rightarrow x = y)))$ and asserts $\exists x(Cx \wedge Sx)$. It seems pretty clear that Frege/Strawson were right on this count—insofar as you are going in for this kind of picture, uniqueness should be thought of as a presupposition, not an entailment. But this is basically irrelevant for our purposes, so I will schematize ‘The cat sleeps’ with $\iota x(Cx \wedge Sx)$, which you can read either as Russell’s uniqueness-entailing or Frege’s uniqueness-presupposing definite. Universal quantification is schematized with \forall and \rightarrow : ‘Every cat sleeps’ gets translated $\forall x(Cx \rightarrow Sx)$. Finally, pronouns like ‘her’ in ‘Every cat has a mother and loves her’ get translated as variables: $\forall x(Cx \rightarrow (\exists y(M(y,x)) \wedge L(x,y)))$.

Foil in hand, let’s see why this can’t be right. Consider again a sentence like (1):

(1) Everyone who has a cat loves $\left\{ \begin{array}{c} \text{her} \\ \text{the cat} \end{array} \right\}$.

Focus first on the variant with pronouns. Apply the logic-class translation rules above to ‘Everyone who has a cat loves it’, and we get (2):

(2) $\forall x((\exists y(Cy \wedge H(x,y))) \rightarrow L(x,y))$

The problem with (2) is that the y in the consequent of the material conditional is unbound. Unbound variables are allowed in the classical picture, but on the understanding that they refer to salient individuals via a contextually given variable assignment. So (2) is a reasonable parse of (1), on standard assumptions, but only in the case where ‘her’ refers to a contextually salient thing—say, a vet who you have just been discussing—*not* of the prominent, covarying reading of (1) which says that every cat-owner loves the cat she owns.

Intuitively, what we need to get the covarying reading is for the y in the consequent to be bound by a quantifier. The only way to achieve that with the tools we started with, however, is to give the existential quantifier $\exists y$ wide scope, as in (3):

(3) $\forall x \exists y((Cy \wedge H(x,y)) \rightarrow L(x,y))$

But this is, once more, a failure: these truth-conditions are absurdly weak. In fact, the mere existence of a non-cat suffices to ensure that (3) is true. In other words, this is a reasonable parse, not of (1), but rather of ‘For every person, there is something which is either not a cat, not owned by that person, or loved by that person’.

What we want, for our translation to approximate the intuitive truth-conditions of (1) (on its covarying reading), is something like (4):

(4) $\forall x \forall y((Cy \wedge H(x,y)) \rightarrow L(x,y))$

This says what (1) intuitively does: namely, that every cat owner loves the cats she owns. But notice how we have transformed the original English sentence: we have taken an indefinite noun phrase from the restrictor of a quantifier, given it scope over the nuclear scope of the quantifier, and turned it into a universal quantifier. The problem, of course, is not just

that this is a departure from the standard rules, but that indefinites can't *generally* be treated as universal quantifiers: 'A cat sleeps' doesn't mean the same thing as $\forall x(Cx \rightarrow Sx)$.

One reaction to this problem would be to try to find translation rules that tell us when indefinite noun phrases get translated as existential quantifiers and when they get translated as universal quantifiers. Maybe that is a reasonable project for some purposes, but not if our project is understand meaning in natural language, since rules of that kind would be totally non-compositional. From the point of view of that approach, donkey sentences show that the classical picture summarized above is wrong.

Things are slightly more complicated, but equally problematic, when we have definite descriptions rather than pronouns. Consider 'Everyone who has a cat loves the cat'. This has a prominent covarying reading on which it seems to again express the same thing as (4): namely, that every cat lover loves her cats. Consider the logic-class translation in (6):

$$(5) \quad \forall x(\exists y(Cy \wedge H(x,y)) \rightarrow \iota y(Cy, L(x,y)))$$

This says: for everyone, either there is nothing that is a cat owned by them; or else there is a unique cat, and they love the cat. That's obviously wrong, since it entails (or presupposes, if you prefer) that there is at most one cat, while (1) obviously doesn't communicate that.

To deal with this problem, we could either spot the definite article some covert material, or relativize its uniqueness inference to a relevant domain. Both options are pursued in the literature, and are indistinguishable for our purposes. So, if we pursued the first option, we could say that (1) has the logical form 'Everyone who has a cat loves the cat they have'. Then we'd get:

$$(6) \quad \forall x(\exists y(Cy \wedge H(x,y)) \rightarrow \iota y(Cy \wedge H(x,y), L(x,y)))$$

Here the cats covary with the universal quantifier, as desired; and (6) does not entail that there is at most one cat. So this is a lot closer to what we want. But it's still not right, because (6) entails/presupposes that every cat-owner has just one cat. Clearly, though, 'Everyone who has a cat loves the cat' can be true even when some people have more than one cat. In case you aren't immediately convinced of this, consider the variant of Heim (1982)'s sage-plant sentence in (7):

$$(7) \quad \text{Everyone who has a cat has a second cat to play with } \left\{ \begin{array}{c} \text{it} \\ \text{the cat} \end{array} \right\}.$$

It's clear enough what this says: that everyone who has one cat has at least two cats. But the sentence would be trivial (in the sense of being true only if the restrictor is empty), if the definite description carried a uniqueness inference. The basic issue is that (7) is predicted to be equivalent to (8), which is obviously trivial:

$$(8) \quad \text{Everyone who has a cat has a second cat to play with the only cat they have.}$$

In sum: when we turn from pronouns to definite articles, things are only marginally better; while it is easier to get a covarying reading than with pronouns, the results predict uniqueness inferences that are not in fact there.

Although donkey quantification illustrates the problems with the classical approach in a particularly striking way, we can illustrate the problem more simply with minimal pairs of sequences or conjunctions, modeled on Partee’s marble sentences (Heim, 1982). Compare:

- (9) a. Sue has a child. She lives at a boarding school.
b. Sue is a parent. She lives at a boarding school.

(9a) is most naturally interpreted as saying that Sue has a child at boarding school, while (9b) is most naturally interpreted as saying that Sue has a child and Sue herself is at boarding school. The classical picture does not have an obvious explanation of this contrast. We could, of course, say that what is happening here is that ‘a child’ binds ‘She’ in (9a), while this kind of binding is not possible in (9b). But it is not clear how this would be possible on the classical approach, in which quantifiers can only bind variables that are in their syntactic scope. Moreover, that response would not do anything to help explain donkey sentences.

We can characterize the contrast in (9) in general by saying that indefinites have *open scope* to their right, as Egli (1979) observed: they can, at some level of description, “bind” definites that follow them, whether the definite is in the indefinite’s syntactic scope or not. In general, where *some* stands for the indefinite, the claim is that sentences with the form $some(p) \wedge q$ and $some(p \wedge q)$ are, in some pragmatically relevant sense, equivalent.

Now, we might of course find different accounts of contrasts like those in (9), on the one hand, and donkey sentences, on the other hand. But impressionistically, these look like very much the same issue: the heart of the problem is how to account for anaphoric dependencies between definite and indefinite noun phrases which go beyond syntactic binding relations.

The pair in (9) also brings out an important point: the phenomenon we are exploring has something essentially to do with the relationship between definites and indefinites in particular. Hence ‘she’ can easily refer to Sue’s child if it’s preceded by ‘Sue has a child’, but not so easily if preceded by ‘Sue is a parent’—even though these plausibly have the same truth-conditions, and certainly do in the classical theory. We can see the same point in the context of quantifiers with minimal variants on (1), as in (10):

(1) Everyone who has a cat loves her.

(10) Every cat-owner loves her.

Even though plausibly someone is a cat-owner iff she has a cat, the prominent covarying reading of (1) seems unavailable for (10). So a different way of framing the problem of donkey anaphora is to account for anaphoric dependencies between definites and corresponding indefinites, and the unavailability of such dependencies between definites and other noun phrases which are intuitively truth-conditionally equivalent to those indefinites.

2 E-type and dynamic responses

I will briefly describe the two main existing categories of response to this problem, e-type and dynamic. Neither of these makes explicit use of trivalence, but partiality plays a central role in the dynamic approach, which also forms the central backdrop for the trivalent approaches we'll explore, so I'll explain the dynamic approach in some detail.

The e-type approach (sometimes, confusingly, called “d-type”) aims to rescue the classical treatment of (in)definites by rejecting other parts of the classical picture.¹ On this approach, pronouns are, at some level of logical form, definite descriptions. Definite descriptions, in turn, generally have unpronounced material which fill out their restrictor (in roughly the way we've sketched so far). Uniqueness disappears because definite descriptions have their uniqueness presuppositions evaluated relative to very small parts of the world, where they can end up being trivially satisfied (we can truly say that there's only one cat in existence if we limit our viewpoint to this room). Although this approach in some sense tries to save the classical semantics above, the saving doesn't go very deep, since the classical theory of quantifiers and connectives goes out the door (to be replaced with custom-made replacements, not, for instance, any familiar Kleene semantics).

The dynamic approach, which starts in the work of Kamp 1981; Heim 1982 (in turn developing ideas from Karttunen 1976) rejects the classical picture wholesale. On the dynamic approach, definite descriptions and pronouns (which I'll refer to, together, as ‘definites’) and indefinites are all essentially free variables. Indefinites are variables which must be contextually new, while definites are variables which must be contextually familiar. This approach gets coupled with a novel approach to sentential content and connectives, as I will now explain.

2.1 Dynamic (in)definites

To illustrate the basic ideas of dynamic semantics, I'll briefly sketch a slightly simplified version of Heim 1982's system. Start with the dynamic notion of a *context*: a set of pairs of variable assignments and worlds. That is, a context is a set of pairs $\langle g, w \rangle$ where w is a possible world and g is a (possibly partial) function which takes variables to individuals in the domain of w . Sentential contents, in general, are *context change potentials (CCPs)*: (possibly partial) functions which, where defined, take an input context to a new context. Where p is a sentence and c a context, we write $[p]$ for the CCP denoted by p , and we write $c[p]$ for the result of applying $[p]$ to c . $c[p]$ is itself a new context, so we can write, for instance, $(c[p])[q]$ for the result of applying $[q]$ to the result of applying $[p]$ to c . We omit parentheses since we can do so without risk of confusion.

¹E.g. Geach 1962; Evans 1977; Parsons 1978; Cooper 1979; Neale 1990; Heim 1990; Ludlow 1994; Büring 2004; Elbourne 2005; Lewis 2012, 2019; Mandelkern and Rothschild 2020; Lewerentz 2020.

The role of indefinites in this system is to *extend* the contextual variable assignments so they are defined on a new variable, and then assign that variable to something satisfying the indefinite’s argument. The role of definites, correspondingly, is to say something about a variable that has already been introduced in this way. So, for instance, ‘There is a cat’ has the CCP [$some_x(Cx)$] which takes a context c and does two things. First, it checks whether x is *novel* in c , that is, whether g is undefined on x for any g included in a pair $\langle g, w \rangle \in c$. If x is not novel, the CCP is undefined. Otherwise, the result is to extend c with all x -variants on variable assignments in c , and then keep all and only those pairs that satisfy Cx . That is, we find the set of pairs $\langle g, w \rangle$ such that there is some pair $\langle g', w' \rangle \in c$, such that $w = w'$, and g is just like g' , except that it is also defined on x (in that case we can write $g >_x g'$, and say that $\langle g, w \rangle$ x -extends $\langle g', w' \rangle$). Then keep just the pairs from that set such that Cx is true—that is, where x is assigned to a cat in the world of the pair.

More formally, it is helpful for expository purposes to uncouple the contribution of $some_x$ from that of its argument. First, we have:²

$$c[some_x] = \begin{cases} \# & \exists \langle g, w \rangle \in c : g(x) \neq \# \\ \{ \langle g, w \rangle : \exists \langle g', w' \rangle \in c : w = w' \wedge g >_x g' \} & otherwise \end{cases}$$

Then we can treat $some_x(p)$ just as successive update of the input context, first with $some_x$, then with p . In particular, where P is any n -ary predicate, and \mathcal{I} takes any world and predicate to that predicate’s extension at that world, we have:

$$c[P(x_1, x_2, \dots, x_n)] = \{ \langle g, w \rangle \in c : \langle g(x_1), g(x_2), \dots, g(x_n) \rangle \in \mathcal{I}(w, P) \}$$

Turning to definites, Heim’s proposal is that definites have a complementary definedness requirement to indefinites: the definite $\iota x(Cx, Sx)$ (‘The cat sleeps’) requires that x be *familiar*, in the sense of being defined everywhere in the context, and in particular assigned to a cat. Assuming definedness, we then update with the restrictor, so we get exactly the pairs $\langle g, w \rangle$ where $g(x)$ is a sleeping cat in w . Generally (we’ll assume the restrictor is always free in x):

$$c[\iota x(p, q)] = \begin{cases} c[q] & c[p] = \# \\ \# & otherwise \end{cases}$$

Pronouns, we can assume, simply have a tautological restrictor, so that $\iota x(\top x, q)$ simply requires that x be defined throughout the input context.

²Cases are to be read from top to bottom, so that, e.g.,

$$f(x) = \begin{cases} a & b \\ c & d \\ e & otherwise \end{cases}$$

says that $f(x)$ is a iff b holds; c iff b does not hold and d holds; and e in all other cases, i.e., iff neither b nor d holds.

The intuition behind this approach is that indefinites “open a new file card”, in Heim’s idiom, or “establish a new discourse referent”, in Karttunen’s. By extending variable assignments throughout the context, indefinites license subsequent anaphora with definites. So, for instance, once we have updated the context with $some_x(Cx)$, we can then use a definite with the form $\iota x(Cx, q)$.

More concretely, this sketch puts us in a position to see how we account for the open scope of definites. So recall (9):

- (9) a. Sue has a child. She lives at a boarding school.
 b. Sue is a parent. She lives at a boarding school.

Assume that (9a) gets parsed $some_x(Chx)$. $\iota x(\top x, Bx)$. Assume x is novel in c ; then this is guaranteed to be defined, since the update of c with the first sentence ensures that x is familiar by the time we get to the second sentence. Moreover, this update will be identical to the update we get with $some_x(Chx \wedge Bx)$: both will take c to the set of pairs $\langle g, w \rangle$ such that g assigns a child of Sue’s in w to x , that child is at boarding school in w , and $\langle g, w \rangle$ x -extends some pair in c .

This matches the intuition that (9a)’s most prominent reading says that Sue has a child at boarding school. By contrast, this reading will not be available for (9b). Assuming (9b) has the form $some_x(Sue(x) \wedge Px)$. $\iota x(\top x, Bx)$, then it will be interpreted as saying that Sue is a parent and Sue is at boarding school, matching its most prominent reading.

Conjunction is treated in dynamic systems as successive context update: that is, $c[p \wedge q] = c[p][q]$ (where defined; I’ll leave definedness caveats implicit throughout what follows). So these points apply equally to the corresponding conjunctions:

- (11) a. Sue has a child, and she lives at boarding school.
 b. Sue is a parent, and she lives at boarding school.

More generally, in the system we’ve sketched so far, indefinites have open scope to their right, in the sense that $some_x(p) \wedge \iota x(p, q)$ will always denote the same CCP as $some_x(p \wedge q)$; likewise for $some_x(p)$. $\iota x(p, q)$.

2.2 Quantifiers

We can integrate quantifiers into the system by successively updating the context with the restrictor and then the scope, so as to capture the dependencies between definites in the scope and indefinites in the restrictor.

So consider ‘Everyone who has a cat loves $\left\{ \begin{smallmatrix} \text{her} \\ \text{the cat} \end{smallmatrix} \right\}$ ’, that is, $every_x(some_y(Cy \wedge H(x, y)), \iota y(Cy, L(x, y)))$. (I focus on the definite description case, but the reasoning in the pronoun case is the same.) Essentially, we want $every_x$ to temporarily extend c with the restrictor, then check that that all the extended points in the resulting set also survive update with the scope. In more detail,

say that g extends g' , written $g \geq g'$, just in case g is defined everywhere g' is, and agrees with g' everywhere that g' is defined. Then we say:³

$$c[\text{every}_x(p, q)] = \{\langle g, w \rangle \in c : \forall g' : \langle g', w \rangle \in c[\text{some}_x][p] \rightarrow \exists g'' \geq g' : \langle g'', w \rangle \in c[\text{some}_x][p][q]\}$$

This handles donkey anaphora in the same way that dynamic semantics predicts the open scope of indefinites. Suppose the restrictor of a quantifier contains an indefinite and the scope a corresponding definite. That indefinite will free up its corresponding variable, and so the definite in the scope will be defined, and will covary with the values of the indefinite.

To see the idea, let's work through our example $\text{every}_x(\text{some}_y(Cy \wedge H(x, y)), \text{ty}(Cy, L(x, y)))$. Consider an arbitrary context c . Assume, for simplicity, that all variables are novel in c , so that all the points in c have the form $\langle g_\emptyset, w \rangle$, where g_\emptyset is the variable assignment that is nowhere defined. Let c^x stand for $c[\text{some}_x]$, i.e. the set of points which x -extend some point from c . $c[\text{every}_x(\text{some}_y(Cy \wedge H(x, y)), \text{ty}(Cy, L(x, y)))]$ will be the set of points from c such that, for any extension of that point which is in $c^x[\text{some}_y(Cy \wedge H(x, y))]$, there is a further extension of that point which is in $c^x[\text{some}_y(Cy \wedge H(x, y))][\text{ty}(Cy, L(x, y))]$. So let $\langle g, w \rangle$ be in $c^x[\text{some}_y(Cy \wedge H(x, y))]$. That means that $g(y)$ is a cat had by $g(x)$. We check whether there is a $g' \geq g$ such that $\langle g', w \rangle \in c^x[\text{some}_y(Cy \wedge H(x, y))][\text{ty}(Cy, L(x, y))]$. First note that this update is guaranteed to be defined, since the indefinite licenses the definite by ensuring that x is familiar and assigned to a cat throughout the input context. And this will hold just in case $g(x)$ loves $g(y)$. So $\langle g_\emptyset, w \rangle$ survives update with our donkey sentence just in case everything in w which owns a cat also loves that cat. We thus get the intended, covarying reading of donkey sentences (without unwanted uniqueness inferences).

2.3 Negation and disjunction

There is much to be said for and against the dynamic approach. This is not the place for a global assessment. However, I want to briefly explain a well-known problem which provides an immediate motivation for trivalent alternatives: the problem of double negation and disjunction.

To work up to the problem, we need to a semantics for negation. The obvious approach simply says that $c[\neg p] = c \setminus c[p]$. This is fine for a propositional dynamic fragment, and indeed is suggested for that purpose by Heim (1983). But it doesn't work for fragments with anaphora. To see this, consider any context c , and consider the sentence 'There is not a cat'. Given the present semantics, $c[\neg \text{some}_x(Cx)]$ would be $c \setminus c[\text{some}_x(Cx)]$. But a little reflection shows this is just c again, as long as it is defined: for, provided that x is novel in c , every variable assignment in a pair in $c[\text{some}_x(Cx)]$ will properly x -extend a pair from c , and hence will not be in c . This is obviously the wrong result.

³Things need to be more complicated than this, because of the well-known proportion problem for this kind of semantics, but this is good enough to give the general idea.

The correct result must capture the universal truth-conditions of negated indefinites: ‘There is not a cat’ says that nothing at all (in the relevant domain)⁴ is a cat. We can get this with the standard dynamic semantics for negation, which says that we keep the pairs from the input context which don’t have any extensions in the context that results from updating the input context with the negatum:

$$c[\neg p] = \{\langle g, w \rangle \in c : \nexists \langle g', w' \rangle \in c[p] : \langle g', w' \rangle \geq \langle g, w \rangle\}$$

This nicely captures the universal import of negated indefinites. On this semantics, $c[\neg \text{some}_x(Cx)]$ is the set of pairs $\langle g, w \rangle \in c$ such that no extension of the pair is in $c[\text{some}_x(Cx)]$. That is the set of pairs $\langle g, w \rangle$ such that no x -extension of g assigns x to a cat in w , which is to say, exactly those pairs $\langle g, w \rangle$ such that there there are no cats in w .

Note that this update does not license subsequent definites since variable assignments are not extended by updates with negated indefinites. (This is easy to see from the fact that $c[\neg p]$ is always a subset of c —that is, $[\neg p]$ is an *eliminative* update.) This is as desired; the ‘she’ in (12b) obviously can’t refer to Sue’s child (of course, this is overdetermined).

- (12) a. Sue has a child. She is at boarding school.
 b. Sue doesn’t have a child. She is at boarding school.

But, while things look good so far, this approach has an undesirable downstream effect. Since single negations are eliminative, double negation is as well: $c[\neg\neg p]$ is obviously always a subset of c . But that means that doubly negated indefinites will not license subsequent definites, just like singly negated indefinites.

Here’s the intuition behind this. In the classical picture, indefinites are existential quantifiers; the universal meaning of negated indefinites falls out immediately from the interaction of Boolean negation with existential quantification. By contrast, in dynamic semantics, indefinites are not existential quantifiers; they are, in essence, just free variables. This is what allows us to capture the open scope of indefinites. But that means that, to capture the universal force of negated indefinites, we need to make *negation* a quantifier. This is exactly what the standard dynamic negation does. But then when you layer two negations, you end up with two quantifiers. But quantification is not involutive: two quantifiers is not the same as no quantifiers! So negation is not involutive either: double negation elimination is not valid; in particular, doubly negated indefinites will not license subsequent definites.

More concretely, consider a doubly negated indefinite sentence, like ‘It is not the case that there is not a cat’, with the form $\neg\neg \text{some}_x(Cx)$. Given our semantics so far, $c[\neg\neg \text{some}_x(Cx)]$ will comprise exactly the pairs $\langle g, w \rangle$ from c such that w contains cats.⁵ So x will remain novel in $c[\neg\neg \text{some}_x(Cx)]$. In other words, doubly negated indefinites have the updated effect of the classical existential quantifier, rather than the dynamic indefinite.

⁴I’ll ignore domain restriction throughout.

⁵At least, provided that x is novel in c and hence our update is defined.

Again, a quick way to see that doubly negated indefinites can never license subsequent anaphora is from the fact that the dynamic negation above is eliminative: $c[\neg p]$ is always a subset of c . If x is novel in c , then, obviously, x is novel in any subset of c . So if x is novel in c , then x is novel in $c[\neg p]$, no matter what p is—and so definites will be undefined in $c[\neg p]$ if they are undefined in c .⁶

So much for the prediction; what are the data? As Karttunen (1976) discussed, doubly negated sentences do license anaphora—contrary to the present prediction. Compare:

- (13) a. It's not the case that Sue doesn't have a child. She's at boarding school.
 b. It's not the case that Sue isn't a parent. She's at boarding school.

Or, somewhat more naturally:

- (14) a. - Sue doesn't have a child. - That's not true! She's at boarding school.
 b. - Sue isn't a parent. - That's not true! She's at boarding school.

The prominent interpretation of the (a)-variants says that Sue has a child who is at boarding school, while the prominent interpretation of the (b)-variants says that Sue herself is at boarding school.

In general, it looks like definites are licensed by doubly negated indefinites.

Double negation is a somewhat exotic environment, to be sure. But the same problem turns up in other, more ordinary environments, like disjunction. As Evans (1977) and Roberts (1987) (citing Partee) observe, negated indefinites in left disjuncts license definites in right disjuncts. Thus consider (15), where the definite naturally refers to the (potential) bathroom:⁷

- (15) Either there is no bathroom, or $\left\{ \begin{smallmatrix} \text{it} \\ \text{the bathroom} \end{smallmatrix} \right\}$ is upstairs.

Intuitively, we want the negation of the left disjunct to be available when we process the right disjunct. The natural way to try to capture this intuition in dynamic systems is as follows (Beaver, 2001):

$$c[p \vee q] = c[p] \cup c[\neg p][q]$$

So consider a sentence with the form $\neg \text{some}_x(Bx) \vee \text{!}x(Bx, Ux)$, like (15). Plugging in this semantics for disjunction, we get $c[\neg \text{some}_x(Bx) \vee \text{!}x(Bx, Ux)] = c[\neg \text{some}_x(Bx)] \cup c[\neg \neg \text{some}_x(Bx)][\text{!}x(Bx, Ux)]$. If double negation elimination were valid, this would give us exactly what we wanted: this would be equivalent to $c[\neg \text{some}_x(Bx)] \cup c[\text{some}_x(Bx)][\text{!}x(Bx, Ux)]$, ensuring that the definite in the right disjunct was licensed and, in the relevant sense, 'referred' to the bathroom. But since double negation elimination is not valid, this is not what we get; in general, the prediction is that the CCP of a sentence with this form will never be defined (except in a degenerate case): the left disjunct will require that x is novel throughout the input context,

⁶Except for the degenerate case where $c[\neg p] = \emptyset$.

⁷Talk of 'reference' throughout is just meant to get at some pretheoretical intuition, and is loose at best.

while the right disjunct will require that x be familiar in some subset of the input context, which can happen only if that subset is the empty set.

The problem is deep. We do not have a lot of latitude in our general approach to negation: if we are to capture the universal import of negated indefinites, given a framework broadly like the one sketched above, we need negation to quantify over variable assignments. But then it is hard to see how two negations could cancel each other out.

I have sketched this problem, however, not to convince you to abandon dynamic semantics, but rather because it provides a particularly clear motivation for the trivalent systems which are our main topic in this chapter, and which I now turn to.

3 Bilateral dynamic semantics

I'll begin by considering the most influential response to this problem from within the framework of dynamic semantics, namely, that of Krahmer and Muskens 1995. Krahmer and Muskens solve the problems I have sketched by moving to a bilateral dynamic semantics which compositionally specifies both an extension and anti-extension for every sentence. The basic idea is that $some_x(p)$ has as its extension the standard dynamic meaning of indefinites, and as its *anti*-extension the universal meaning of negated classical indefinites. Then negation can operate simply as a switch between extension and anti-extension.

So (translating the ideas from the relational framework Krahmer and Muskens use into the functional framework we have been working with), we distinguish the extension $[p]^+$ and anti-extension $[p]^-$ of any sentence, both of which are still CCPs. Our entry for negation simply switches from extension to anti-extension, and vice versa:

$$\begin{aligned} c[\neg p]^+ &= c[p]^- \\ c[\neg p]^- &= c[p]^+ \end{aligned}$$

It is already clear from this that double negation elimination will be valid, since $c[\neg\neg p]^+ = c[\neg p]^- = c[p]^+$.

Our entry for atoms is predictable:

$$\begin{aligned} c[P(x_1, x_2, \dots, x_n)]^+ &= \{\langle g, w \rangle \in c : \langle g(x_1), g(x_2), \dots, g(x_n) \rangle \in \mathcal{I}(w, P)\} \\ c[P(x_1, x_2, \dots, x_n)]^- &= \{\langle g, w \rangle \in c : \langle g(x_1), g(x_2), \dots, g(x_n) \rangle \notin \mathcal{I}(w, P)\} \end{aligned}$$

The heart of the proposal is the following two-sided entry for *some*:⁸

$$c[some_x(p)]^+ = \{\langle g, w \rangle : \exists \langle g', w' \rangle \in c : w = w' \wedge g >_x g' \wedge \langle g, w \rangle \in c[p]^+\}$$

⁸I leave out the novelty condition for readability.

$$c[\text{some}_x(p)]^- = \{\langle g, w \rangle \in c : \nexists \langle g', w' \rangle : w = w' \wedge g' >_x g \wedge \langle g', w' \rangle \in c[p]^+\}$$

In essence, $\text{some}_x(p)$ has the extension of the dynamic update with indefinites, while it has the anti-extension of the dynamic update with negated indefinites.

This bilateral approach allows Krahmer and Muskens to navigate the poles that squeeze any theory of the indefinite. On the one hand, we want matrix indefinites to update contextual variable assignments, and hence not to quantify over assignments. On the other hand, we want a universal meaning for negated indefinites, which means they need to quantify over assignments. Achieving the latter by making negation itself a quantifier over assignments yields the problematic failure of double negation elimination we have seen. By contrast, making both of these meanings already a part of the semantics of indefinites—as an extension and anti-extension, respectively—lets us capture the two faces of indefinites, while keeping quantification out of negation, and thus validating double negation elimination.

While Krahmer and Muskens’ approach is presented in a bilateral format, their original (relational) semantics can be equivalently reformulated in trivalent terms, as Benjamin Spector has helpfully pointed out to me.⁹ Their treatment can thus be seen as a landmark in the trivalent treatment of anaphora. (Of course, partiality is already an essential part of (Heimian) dynamic semantics, so trivalence already, in a way, played a central role in that approach.)

Having said that, this system is somewhat complex. And once we start using trivalence, it is not entirely clear that we need the powerful functional/relational apparatus of dynamic semantics at all, as we will see presently.

Before going on, I should note that there are other, broadly related, broadly dynamic solutions to the double negation problem: see especially van den Berg 1996; Gotham 2019; Elliott 2020a, b. I omit discussion of these interesting systems only for reasons of space.

4 Rothschild’s trivalent approach

In the rest of the chapter, I will explore two systems which are trivalent but dispense with the functional/relational setting of dynamic semantics (though both systems borrow ideas from dynamic semantics, as will be clear throughout and as I discuss in the conclusion). The idea of accounting for anaphora in a trivalent framework *without* the apparatus of dynamic semantics is a very recent one, first explored, to my knowledge, in Rothschild 2017, which I will now summarize.

⁹It is not clear that this holds for the functional version of their semantics that I have presented here, but that doesn’t affect the basic point.

4.1 Connectives

To lay out Rothschild's system, I will work with the same language, closed under an additional one-place operator \dagger (a closure operator, on which more shortly). Pronouns are now translated into our language as variables, rather than using t .

Rothschild's treatment of the connectives is the standard strong Kleene one. That is, where g is an assignment and w a world, and $\llbracket p \rrbracket^{g,w}$ the truth-value of p at $\langle g, w \rangle$:

$$\llbracket p \wedge q \rrbracket^{g,w} = \begin{cases} 1 & \llbracket p \rrbracket^{g,w} = \llbracket q \rrbracket^{g,w} = 1 \\ 0 & \llbracket p \rrbracket^{g,w} = 0 \text{ or } \llbracket q \rrbracket^{g,w} = 0 \\ \# & \text{otherwise} \end{cases}$$

$$\llbracket p \vee q \rrbracket^{g,w} = \begin{cases} 1 & \llbracket p \rrbracket^{g,w} = 1 \text{ or } \llbracket q \rrbracket^{g,w} = 1 \\ 0 & \llbracket p \rrbracket^{g,w} = \llbracket q \rrbracket^{g,w} = 0 \\ \# & \text{otherwise} \end{cases}$$

$$\llbracket \neg p \rrbracket^{g,w} = \begin{cases} 1 & \llbracket p \rrbracket^{g,w} = 0 \\ 0 & \llbracket p \rrbracket^{g,w} = 1 \\ \# & \llbracket p \rrbracket^{g,w} = \# \end{cases}$$

If you want an intuition behind these entries, a helpful heuristic: we can (though we don't have to) think of undefinedness as classical underdetermination. A sentence is true if it's classically true on every way of filling in the underdetermination, false if classically false on every way of filling in the underdetermination, and undefined otherwise. Since these entries will be familiar to most readers, I let them stand here without discussion; see the Foundations section of this volume for more exposition.

Note that these are symmetric entries. Anaphora, of course, involves striking asymmetries: 'he' in (16b) is not naturally interpreted as being anaphoric on 'a man', while it is in (16a):

- (16) a. A man came in and he was wearing a beret.
 b. He was wearing a beret and a man came in.

Right-to-left anaphoric dependencies are, however, possible:

- (17) He wasn't satisfied, so a student of mine was disputing his grade.

While Rothschild develops a symmetric system, he notes that you could instead work with an asymmetric, "middle Kleene" system instead (I lay out such a system in §5.2 below).

4.2 Anaphora

The first innovation in the system comes in the treatment of predicational sentences. Rothschild assumes that there is a contradictory object \perp in the domain which can be the value

of variable. The idea is that a predicational sentence $P(x_1, x_2, \dots, x_n)$ is undefined at g if g assigns any of $x_1 \dots x_n$ to \perp :

$$\llbracket P(x_1, x_2, \dots, x_n) \rrbracket^{g,w} = \begin{cases} \# & g(x_1) = \perp \text{ or } g(x_2) = \perp \text{ or } \dots \text{ or } g(x_n) = \perp \\ 1 & \langle g(x_1), g(x_2), \dots, g(x_n) \rangle \in \mathcal{I}(w, P) \\ 0 & \text{otherwise} \end{cases}$$

Incidentally, while this is helpful heuristically, it should not be taken as ontologically committal; we could, as far as I can see, equally work with partial assignments, rather than with total assignments and the object \perp .

The final central move comes in treating $\text{some}_x(p)$ as having the same extension as p , but a slightly wider anti-extension than p :

$$\llbracket \text{some}_x(p) \rrbracket^{g,w} = \begin{cases} 1 & \llbracket p \rrbracket^{g,w} = 1 \\ 0 & g(x) = \perp \text{ or } \llbracket p \rrbracket^{g,w} = 0 \\ \# & \text{otherwise} \end{cases}$$

To see the basic idea, consider a sentence like $\text{some}_x(Cx \wedge Sx)$ ('A cat sleeps'). This has the same truth conditions as $Cx \wedge Sx$: it is true at $\langle g, w \rangle$ just in case $g(x)$ is a sleeping cat in w . But this sentence slices falsity and undefinedness differently from the open sentence. Both the open sentence $Cx \wedge Sx$ and $\text{some}_x(Cx \wedge Sx)$ are false when $g(x)$ is something other than \perp which is not a sleeping cat. But the open sentence is undefined if $g(x) = \perp$, whereas $\text{some}_x(Cx \wedge Sx)$ is false when $g(x) = \perp$. So some_x is, in Karttunen's terminology, a plug for undefinedness stemming from assignment of x to \perp .

The final piece of the central picture is a fairly standard pragmatic story. Rothschild adopts a combination of Heimian and Stalnakerian assumptions. Following Heim, Rothschild models a context as a set of pairs of worlds and variable assignments. Following Stalnaker, he assumes that p is only assertable in a context c if p is defined at every point in c . When p is defined throughout c , then if p is asserted and accepted, the subsequent context is the set of points in c where p is true, that is, $c \cap \llbracket p \rrbracket$.¹⁰

This puts us in a position to see how the characteristic law of dynamic semantics—the open scope of indefinites—is validated by Rothschild's system. Consider a sentence with the form

$$(18) \quad \text{some}_x(Cx) \wedge Px.$$

That is, 'There is a cat and it is purring'. The open scope of indefinites says that this sentence should be equivalent to 'There is a purring cat':

$$(19) \quad \text{some}_x(Cx \wedge Px).$$

¹⁰ $\llbracket p \rrbracket = \{ \langle g, w \rangle : \llbracket p \rrbracket^{g,w} = 1 \}$.

And indeed, in Rothschild’s system, these are equivalent. Consider any point $\langle g, w \rangle$. Given the rules we have laid out, (18) and (19) are both true whenever $g(x)$ is a purring cat in w . Both are false otherwise: if $g(x) = \perp$, then the right conjunct of (18) is undefined, but the left conjunct is false, and so by the strong Kleene rules, the conjunction is false. And, by our rule for *some*, (19) is false, too, if $g(x) = \perp$. If $g(x) \neq \perp$ but is not a purring cat in w , then both sentences are false.

This reasoning is general: $\text{some}_x(p) \wedge q$ is semantically equivalent to $\text{some}_x(p \wedge q)$. And this extends equally to sequences of sentences: $\text{some}_x(p)$. q will always lead to the same update as $\text{some}_x(p \wedge q)$.

At a high level, this approach mimics Heim’s treatment, and the novelty/familiarity intuition. An indefinite sentence $\text{some}_x(p)$ is assertable in a context which includes assignments of x to \perp , since $\text{some}_x(p)$ is simply false at any such assignment. Once $\text{some}_x(p)$ has been asserted, subsequent sentences which are free in x —that is, a sentence with a pronoun indexed to x —will be assertable, since, once we’ve updated with $\text{some}_x(p)$, the context won’t include any assignments of x to \perp . By contrast, a sentence which is free in x won’t generally be assertable without a preceding indefinite sentence, since the context then may contain assignments which take x to \perp , in which case the assertion would crash (modulo the usual points about accommodation).¹¹

Rothschild focuses on anaphora involving pronouns, but we can extend his system to definite descriptions easily enough. Extend our language so that, for any variable x and sentence p free in x , $x(p)$ is also a formula. Let the terms of our language be all the strings of the language starting with a variable (so, the variables, plus the definite descriptions). The idea is that a bare variable x is interpreted as a pronoun, while a string like $x(Cx)$ stands for ‘the cat’. We extend our semantics for predication in the usual way: where $\tau_1 \dots \tau_n$ are any terms,

$$\llbracket P(\tau_1, \tau_2, \dots, \tau_n) \rrbracket^{g,w} = \begin{cases} \# & \llbracket \tau_1 \rrbracket^{g,w} = \# \text{ or } \llbracket \tau_2 \rrbracket^{g,w} = \# \dots \text{ or } \llbracket \tau_n \rrbracket^{g,w} = \# \\ 1 & \langle \llbracket \tau_1 \rrbracket^{g,w}, \llbracket \tau_2 \rrbracket^{g,w}, \dots, \llbracket \tau_n \rrbracket^{g,w} \rangle \in \mathcal{S}(w, P) \\ 0 & \text{otherwise} \end{cases}$$

For any variable x , we have:

$$\llbracket x \rrbracket^{g,w} = \begin{cases} g(x) & g(x) \neq \perp \\ \# & g(x) = \perp \end{cases}$$

Finally, for a definite description $x(p)$, we have:

$$\llbracket x(p) \rrbracket^{g,w} = \begin{cases} g(x) & \llbracket p \rrbracket^{g,w} = 1 \\ \# & \llbracket p \rrbracket^{g,w} \neq 1 \end{cases}$$

¹¹Note that Rothschild only encodes “familiarity” (that is, his definedness condition on predicational sentences), and not novelty. But, as he notes, it’s not clear that there is any direct empirical evidence for novelty.

That is, a definite like $x(Cx)$ will denote $g(x)$ provides that $g(x)$ is a cat in w , and otherwise \perp . This will extend the reasoning about the open scope of indefinites to definite descriptions in a straightforward way. $P(x(Cx))$ ('The cat purrs') will lead to a crash in any context where x is somewhere assigned to a non-cat. But if we first update with $some_x(Cx)$, then all the points in the context will assign x to a cat, and hence $P(x(Cx))$ is guaranteed to be defined.

4.3 Negation

So far, this is supremely elegant. If we can do things with such minimal resources, we should.

The rest of the system is somewhat less compelling. Problems start when we come to negation. Consider 'There's not a sleeping cat': $\neg some_x(Cx \wedge Sx)$. Given the strong Kleene negation, this is true at $\langle g, w \rangle$ just in case $g(x)$ is a non-sleeping-cat in w or is \perp . This does not get us the strong, universal meaning we want for negated indefinites.

To deal with this problem, as we have seen, Heim (1982), in essence, makes negation a quantifier over assignments. That is the approach, anyway, of Chapter 3 of Heim 1982; a different approach comes from the second chapter of Heim 1982. That approach is to insert existential closure operators which take scope over indefinites, and, in essence, transform them from open sentences to existentially quantified sentences.

Rothschild follows this latter approach. Say that p is *definedness-sensitive* to a variable x iff there is a pair $\langle g, w \rangle$ such that $\llbracket p \rrbracket^{g[x \rightarrow \perp], w} = \#$, while $\llbracket p \rrbracket^{g[x \rightarrow o], w} \neq \#$ for all individuals o other than \perp ($g[x \rightarrow a]$ is the variable assignment which takes x to a but is otherwise just like g). In essence, p is definedness-sensitive to x iff x appears in p but is not bound by an indefinite in p . We write $g'[p]g$ to say that g' and g agree on all definedness-sensitive variables in p . This, in turn, lets us define our closure operator \dagger :

$$\llbracket \dagger p \rrbracket^{g, w} = \begin{cases} 1 & \exists g'[p]g : \llbracket p \rrbracket^{g', w} = 1 \\ 0 & \forall g'[p]g : \llbracket p \rrbracket^{g', w} = 0 \\ \# & \textit{otherwise} \end{cases}$$

So, in essence, $\dagger p$ quantifies over variable assignments which vary variables bound by indefinites in p ; so the truth conditions of $\dagger some_x(p)$ are those, not of the open sentence p , but rather of the classical existential quantification $\exists x(p)$.

The idea is that the dagger can be inserted under negation, to existentially close a sentence. There are not explicit rules about when this happens: we obviously don't want to existentially close all sentences, since this would rob the fragment of its ability to account for anaphora. So 'There's not a sleeping cat' gets parsed as $some_x(Cx \wedge Sx)$, but 'There's not a sleeping cat' gets parsed, not as $\neg some_x(Cx \wedge Sx)$, but rather as $\neg \dagger some_x(Cx \wedge Sx)$. This has the desired universal truth-conditions. $some_x(Cx \wedge Sx)$ is not definedness-sensitive to x , since taking x to \perp would make the sentence false, not undefined. Hence $\neg \dagger some_x(Cx \wedge Sx)$

is true at $\langle g, w \rangle$ just in case $\dagger some_x(Cx \wedge Sx)$ is false at $\langle g, w \rangle$ just in case, for every o in the domain, $Cx \wedge Sx$ is false at $\langle g_{[x \rightarrow o]}, w \rangle$, just in case there is no sleeping cat in w , as desired.

What about double negation elimination? Well, double negation elimination is valid for the strong Kleene connectives, as we noted. So as long as ‘It’s not true that there is not a sleeping cat’ gets parsed $\neg\neg some_x(Cx \wedge Sx)$, without any intervening daggers, it will have the desired anaphoric properties, since this is equivalent to $some_x(Cx \wedge Sx)$. So, unlike the dynamic approach, this approach is capable of accounting for the anaphoric potential of doubly negated indefinites.

Having said that, this approach does not give us an account of when we insert daggers and when we don’t, which makes it unsatisfyingly unconstrained. For instance, nothing in the account rules out the parse $\neg\neg\dagger some_x(Cx \wedge Sx)$, or indeed, $\neg\dagger\neg some_x(Cx \wedge Sx)$ or $\neg\dagger\neg\dagger some_x(Cx \wedge Sx)$. On any of these parses, doubly negated indefinites would not license subsequent anaphora. Likewise, nothing in this approach rules out the parse $\dagger some_x(Cx)$ for a matrix indefinite; that parse, of course, makes the indefinite equivalent to the classical existential quantifier, and hence will ensure that it does not license subsequent anaphora. Nor does it rule out the parse $\neg\dagger some_x(Cx)$ for a singly negated indefinite, which would predict a reading of ‘it’s not the case that there is a cat’ as saying ‘A particular thing is not a cat’. All these questions are raised, and not answered, by Rothschild’s account.

4.4 Disjunction

This treatment of double negation elimination nearly puts us in a position to deal with bathroom disjunctions, but we need one more tool. Consider (20):

(20) Either there is no bathroom, or it is upstairs.

Assume first that this gets parsed:

(21) $\neg some_x(Bx) \vee Ux$

This parse isn’t acceptable: we want the negated indefinite in the left disjunct to get its usual strong interpretation. (To see the problem, suppose there is a bathroom in the basement but not upstairs. Then intuitively (20) is false. But (21) will be true, relative to any variable assignment which takes x to something other than a bathroom.)

To get the intended strong reading of the left disjunct, we can insert a dagger:

(22) $\neg\dagger some_x(Bx) \vee Ux$

This lets us avoid the problem just sketched, but now we have a different problem: the pronoun in the right disjunct of (22) will not be licensed by the indefinite in the left. That is, (22) still puts a non-trivial requirement on the input context: namely, that it not contain any pairs $\langle g, w \rangle$ such that there is a bathroom in w , but $g(x) = \perp$. For otherwise, if w contains a bathroom but $g(x) = \perp$, then by our truth-conditions for disjunction, (22) will be

undefined at $\langle g, w \rangle$ since the left conjunct is false and the right conjunct is undefined. But intuitively, bathroom disjunctions don't require any contextual set-up—that is, they don't require a preceding indefinite to license the pronoun in the right disjunct.

So we seem stuck: we somehow need the dagger to get the intuitive truth-conditions of the left disjunct, but we can't have it if we are to capture the intuitive licensing conditions of the right disjunct. Rothschild deals with this by having it both ways. Rothschild says that a sentence with the surface form p can have the logical form q' provided that for some sentence q , p and q are classically equivalent (that is, equivalent under a classical interpretation of the language) and q' is identical to q except perhaps for the addition of some daggers. Call this mechanism of free copying *cloaking*. Since $p \vee q$ and $p \vee (\neg p \wedge q)$ are classically equivalent, we can according to this rule parse (20) as (23):

$$(23) \quad \neg \dagger some_x(Bx) \vee (some_x(Bx) \wedge Ux)$$

This has the desired meaning and (lack of) anaphoric requirements. The whole disjunction is true when there is no bathroom. In case there is a bathroom, it is true iff $g(x)$ is an upstairs bathroom in w . And it is never undefined: if $g(x) = \perp$, then the right disjunct will be false, not undefined, since *some* plugs undefinedness.

So we are in a similar situation as in the case of double negation. We can deal with Partee sentences. This requires extra machinery, however: in this case, not only a dagger, but also a cloak. It is worth noting a dialectical wrinkle at this point: cloaking could also save the dynamic account of disjunction, and perhaps also of double negation. If we are allowed to interpret (20) as $\neg some_x(Bx) \vee (some_x(Bx) \wedge Ux)$, then the standard dynamic interpretation works unproblematically. Likewise, if we can copy double-negated material across discourses, then dynamic semantics has no trouble with those cases. So, although Rothschild sells these as empirical advantages of his account, the dynamic approach could in principle also help itself to cloaking and then deal with these cases in the same way (though of course, we might still find other reasons to prefer Rothschild's account).¹²

4.5 Quantifiers

Rothschild adopts a natural extension of the strong Kleene system to the universal quantifier: roughly speaking, $every_x(p, q)$ is true if p is defined on every object, and q is true of every object that p is true of; false if there is a $p \wedge \neg q$ object; and # otherwise.¹³ So, where D is the domain (which we assume is fixed across worlds, but does not include the

¹²Rothschild also points to two nice further features of his account: its ability to deal with Stone (1992)'s disjunctive anaphora, like 'If you see a dog or a cat, give it a kiss'; and corresponding conditional cases which Rothschild points out, like 'Either its a holiday or a customer will come in. And if it's not a holiday, he'll want to be served'. The variant we develop presently can also deal with these cases in a parallel way.

¹³What I present here is slightly different from Rothschild's presentation, which I think contains an error.

contradictory object):

$$\llbracket \text{every}_x(p, q) \rrbracket^{g, w} = \begin{cases} 1 & \forall o \in D : \llbracket p \rrbracket^{g[x \rightarrow o], w} \neq \# \text{ and } \llbracket p \rrbracket^{g[x \rightarrow o], w} = 1 \rightarrow \llbracket q \rrbracket^{g[x \rightarrow o], w} = 1 \\ 0 & \exists o \in D : \llbracket p \wedge \neg q \rrbracket^{g[x \rightarrow o], w} = 1 \\ \# & \text{otherwise} \end{cases}$$

To make this work for donkey anaphora, we once more need cloaks and daggers. To see this, return to our running donkey sentence, focusing on the variant with the pronoun:

(24) Everyone who has a cat loves it.

Assume first this gets parsed as in (25):

(25) $\text{every}_x(\text{some}_y(Cy \wedge H(x, y)), L(x, y))$

This won't work. The problem is, in essence, that *some* isn't quantificational at all in this system—*modulo* definedness, indefinite sentences are equivalent to open sentences. But we need to quantify over cats to get the intended meaning. We can work around this, again, with appropriately copied material, and appropriately placed daggers: namely, by parsing (24) as (26) rather than (25):

(26) $\text{every}_x(\dagger \text{some}_y(Cy \wedge H(x, y)), \dagger(\text{some}_y(Cy \wedge H(x, y)) \wedge L(x, y)))$

The reason this works is that the \dagger turns the indefinite into, in essence, a classical existential quantifier. So consider an arbitrary object o in the domain. o is a witness to the truth of our sentence at $\langle g, w \rangle$ just in case (i) either the restrictor is false, relative to $\langle g[x \rightarrow o], w \rangle$, or else (ii) both the restrictor and scope are true, relative to $\langle g[x \rightarrow o], w \rangle$. So suppose that the restrictor is true; then there is some cat that o has. The scope is true just in case there is some cat that o has and loves. The sentence is true iff every individual in the domain witnesses its truth; in other words, iff every cat-owner loves some cat they own. We thus derive the intuitive meaning of the donkey sentence. (To be precise, we derive one of two apparent interpretations, the “weak” one. A stronger interpretation says that every cat-owner loves every cat they own. It's controversial whether both readings are semantically available, or if there is instead some kind of flexibility in interpretation.¹⁴)

5 Adding local contexts

There is a thrilling simplicity to Rothschild's system, until you confront the necessity of cloaks and daggers, which make the system both more revisionary and less constrained than the negation-free fragment promises to be. Being revisionary is not in itself an objection, but it does deprive this approach of some of its apparent advantages over dynamic semantics—especially when you see that dynamic semantics can avoid its problems with disjunction and

¹⁴See e.g. Heim 1982; Root 1986; Rooth 1987; Schubert and Pelletier 1989; Chierchia 1992; Kanazawa 1994; Chierchia 1995; Champollion et al. 2019.

double negation if it, too, invokes some kind of cloaking mechanism. Being unconstrained is a more acute problem: to have a predictive theory, we need to know when we insert daggers and when we don't. For instance, we want an explanation of why it is very hard to get a dagger-free reading of singly negated indefinites (that is, a reading on which the negated indefinite does not have universal quantificational force), and why it is, conversely, very hard to get a daggered reading of doubly negated indefinites (that is, a reading on which the doubly negated indefinite does not license anaphora).

In the rest of this chapter, I will develop an alternative picture—a different ending to Rothschild's trivalent story. The central idea about indefinites is based on a similar proposal which I develop in Mandelkern 2022b. In that paper, I develop the idea in a bivalent but two-dimensional framework, rather than a trivalent framework; in the conclusion, I return briefly to the choice between these frameworks. Both that proposal and the one I presently develop have their roots in Krahmer and Muskens (1995)'s bilateral account, as will become clear.

At a high level, here is how we deal with the problems that lead Rothschild to invoke cloaks and daggers. First, instead of using a closure operator, we derive the strong force of negated indefinites by encoding it directly in the trivalent truth-conditions: $some_x(Px)$ is true just in case Px is, but false just in case nothing is in $\mathcal{I}(w, P)$. That means that indefinites have the update effect of open sentences—hence licensing anaphoric dependencies. But singly negated indefinites have only the desired universal truth-conditions, no anaphoric effects. Together with Kleene connectives (middle, in our case, instead of strong), this gets us what we need as far as double negation goes.

Second, instead of freely copying classically redundant material, we use the apparatus of *local contexts* to define the familiarity condition on definites. This is not so different from Rothschild's approach: in fact, Schlenker (2008) influentially showed that local contexts can simply be viewed as the strongest locally redundant material in a sentence. But the result is more constrained: local contexts are recursively specified, whereas cloaking is in Rothschild's system optional but not mandatory; and the ideology of local contexts, unlike cloaking, does not require new commitments about the relationship between surface and logical form.

5.1 (In)definites

I'll work with essentially the same fragment, but we will no longer need \dagger . I'll write definites $\iota x(p)$ ('The p ') and $\iota x(\top(x))$ ('he'/'she'/'it'), treating these as terms.

Start with indefinites. The idea, again, is that indefinites have the extension of open sentences, but the anti-extension of negated existential quantifiers. That is, for any world w ,

assignment g , and context c (more on the role of contexts shortly):

$$\llbracket \text{some}_x(p) \rrbracket^{c,g,w} = \begin{cases} 1 & \llbracket p \rrbracket^{c,g,w} = 1 \\ 0 & \forall o \in D : \llbracket p \rrbracket^{c,g[x \rightarrow o],w} = 0 \\ \# & \text{otherwise} \end{cases}$$

So, for instance, $\text{some}_x(Cx)$ is true at $\langle c, g, w \rangle$ just in case $g(x)$ is a cat in w ; false just in case w is cat-less; and undefined otherwise.

Turning to definites: as in the extension of Rothschild's system above, definites have the value of the corresponding variable, provided that the variable's prejacent is true throughout the context (which is still a set of world/variable assignment pairs):

$$\llbracket \iota x(p) \rrbracket^{c,g,w} = \begin{cases} \llbracket x \rrbracket^{c,g,w} & \forall \langle g', w' \rangle \in c : \llbracket p \rrbracket^{c,g',w'} = 1 \\ \# & \text{otherwise} \end{cases}$$

Since we are assuming that pronouns are definites with tautological restrictors—that is, with the form $\iota x(\top x)$ —this amounts to the requirement for pronouns that all the variable assignments in the context are defined on the variable in question. For definite descriptions, the requirement is stronger—namely, that all the variable assignments in the context be defined on the variable in question and moreover assign it to something that makes the restrictor true. So, for instance, $\iota x(Cx)$ is defined in c only if x is assigned to a cat throughout c ; in that case, its value at $\langle g, w \rangle$ is $\llbracket x \rrbracket^{c,g,w}$.

Variables, in turn, get the value assigned to them by the variable assignment in the usual way; we assume that assignments can be partial, so $\llbracket x \rrbracket^{c,g,w}$ is $g(x)$ provided the latter is defined, and undefined otherwise. (We assume that variables are always in the scope of a co-indexed ι , some , or quantifier.)

Finally, predicational sentences again get the natural trivalent semantics:

$$\llbracket P(\tau_1, \tau_2, \dots, \tau_n) \rrbracket^{c,g,w} = \begin{cases} \# & \llbracket \tau_1 \rrbracket^{c,g,w} = \# \text{ or } \llbracket \tau_2 \rrbracket^{c,g,w} = \# \dots \text{ or } \llbracket \tau_n \rrbracket^{c,g,w} = \# \\ 1 & \langle \llbracket \tau_1 \rrbracket^{c,g,w}, \llbracket \tau_2 \rrbracket^{c,g,w}, \dots, \llbracket \tau_n \rrbracket^{c,g,w} \rangle \in \mathcal{S}(w, P) \\ 0 & \text{otherwise} \end{cases}$$

5.2 Connectives

Turning to the connectives, I'll make two changes from Rothschild's approach. First, we switch from strong to middle Kleene connectives, for reasons I'll explain below. This is not a deep change, since as Rothschild notes, his approach could be developed in a middle Kleene context as well. Second, more significantly, we incorporate local contexts into our treatment of the connectives; that is, the context we use to evaluate the part of a complex sentence is not necessarily the input context, but may be a part of the input context. There

are a lot of different ways of thinking about local contexts. A central question is whether they are part of the recursive semantic machinery, or something separable from (but determined by) the semantic machinery. Karttunen (1973, 1974) puts them in the former category; Stalnaker (1974); Schlenker (2008) in the latter. Partly for expository purposes, I'll follow the first route, specifying them recursively; see Spector 2021 for an attempt to follow the latter route in the context of a theory of anaphora. Another question is whether to implement local contexts symmetrically or asymmetrically; I will opt for an asymmetric implementation here, though, again, there is controversy about order. For any sentence p and context c , I write c^p for $\{ \langle g, w \rangle \in c : \llbracket p \rrbracket^{c,g,w} = 1 \}$. Then the middle Kleene semantics with local contexts is:

$$\llbracket p \wedge q \rrbracket^{c,g,w} = \begin{cases} 1 & \llbracket p \rrbracket^{c,g,w} = \llbracket q \rrbracket^{c^p,g,w} = 1 \\ 0 & \llbracket p \rrbracket^{c,p,w} = 0 \text{ or } (\llbracket p \rrbracket^{c,p,w} = 1 \text{ and } \llbracket q \rrbracket^{c^p,p,w} = 0) \\ \# & \textit{otherwise} \end{cases}$$

$$\llbracket p \vee q \rrbracket^{c,g,w} = \begin{cases} 1 & \llbracket p \rrbracket^{c,g,w} = 1 \text{ or } (\llbracket p \rrbracket^{c,g,w} = 0 \text{ and } \llbracket q \rrbracket^{c^{\neg p},g,w} = 1) \\ 0 & \llbracket p \rrbracket^{c,g,w} = \llbracket q \rrbracket^{c^{\neg p},g,w} = 0 \\ \# & \textit{otherwise} \end{cases}$$

$$\llbracket \neg p \rrbracket^{c,g,w} = \begin{cases} 1 & \llbracket p \rrbracket^{c,g,w} = 0 \\ 0 & \llbracket p \rrbracket^{c,g,w} = 1 \\ \# & \llbracket p \rrbracket^{c,g,w} = \# \end{cases}$$

The middle Kleene connectives are essentially order-sensitive versions of the strong Kleene connectives; if the left clause suffices to determine the truth-value of the whole sentence, whatever the truth-value of the right clause is, that suffices for the sentence to have a value; otherwise, both clauses must be valued. So for instance if p is true and q indeterminate, then $p \vee q$ is true (since, whether q is determined as true or as false, the classical value of the sentence would be true either way); while $q \vee p$ is indeterminate (since just the material to the left of q —of which there isn't any—doesn't suffice to fix the classical value of this sentence).

5.3 Pragmatics

Like Rothschild, we will borrow from Heim the notion of a context as a set of pairs of possible worlds and (possibly partial) variable assignments.¹⁵ Unlike Rothschild, however, we don't adopt the Stalnakerian update rule, which says that you can only assert p when p is defined throughout the context. Instead we just have (what seems to me) the simplest possible update rule: we update with p by keeping all and only points from the input context c where p is true relative to c .¹⁶ We thus reject "Stalnaker's Bridge", the rule that says you

¹⁵We dispense with Rothschild's contradictory object.

¹⁶This simplifies in irrelevant ways: if we think of context as representing the common ground, then further adjustments will follow, e.g., we will also update with the fact that we now accept p , and so on.

can only assert p if p is defined throughout the context. This rule plays a crucial role in Rothschild’s treatment of anaphora, and, more generally, in the trivalent approach to presupposition. But it is crucial for us that we *don’t* have this rule.

To see this, consider a sentence like $some_x(Cx)$. We want to be able to update an uninformative context c with a sentence like this, with the result licensing subsequent definites indexed to x . On my proposed update rule, we update by keeping the points $\langle g, w \rangle$ from c such that $some_x(Cx)$ is true at $\langle c, g, w \rangle$. Given our semantics, that means that the updated context c' comprises exactly the points $\langle g, w \rangle$ such that $g(x)$ is a cat in w . That means that a subsequent definite indexed to x will be licensed, since x will be assigned to a cat throughout the updated context. That is, $P(\iota x(Cx))$ (‘The cat is purring’) or $P(\iota x(\top x))$ (‘It is purring’) will be defined throughout c' , and in particular true at exactly those points $\langle g, w \rangle$ where $g(x)$ is purring. By contrast, in a context where x is not familiar, $P(\iota x(\top x))$ would be undefined at every point in that context.

If, however, we took on board Stalnaker’s Bridge, indefinites would not be able to play this essential role in setting up subsequent anaphora. In a context c where x is not familiar, $some_x(Cx)$ will be undefined at any point $\langle g, w \rangle \in c$ where $g(x)$ is undefined or where w has a cat but $g(x)$ is not a cat. So in a context where x is novel, $some_x(Cx)$ will be undefined at many points in the context. Hence if we required any assertion to be defined throughout the context, we could never use an indefinite sentence to make a novel variable familiar and hence set up subsequent anaphora.

Although Stalnaker’s Bridge is an important part of the traditional picture of semantic presupposition projection, it’s also stipulative, as many have noted (Soames, 1989; von Stechow, 2008); so it does not seem intrinsically problematic to reject it. It does mean that our approach can’t be coupled with the standard trivalent treatment of presupposition. We have an alternative, though. Given that we also have the apparatus of local contexts in our system, we can follow the tradition of Karttunen, Stalnaker, Heim and Schlenker and say that a sentence is only defined if its semantic presuppositions are satisfied throughout their local context. Once you have that, Stalnaker’s Bridge is superfluous—the quantification over context worlds implicit in Stalnaker’s Bridge is instead taken care of by way of the presupposition satisfaction rule.

5.4 Open scope

This is the core of the theory. We have seen how updating with an indefinite licenses subsequent anaphora. We are also in a position to see how this approach validates the open scope of indefinites, and how it avoids problems with negation and disjunction.

Hence, consider (27) and (28) (‘There is a cat and it is purring’ and ‘A cat is purring’):

$$(27) \quad some_x(Cx) \wedge P(\iota x(\top x))$$

$$(28) \quad some_x(Cx \wedge Px)$$

In the present system, (27) and (28) are not semantically equivalent. But that is only because their falsity conditions differ. They still have the same *truth*-conditions: each is true at $\langle c, g, w \rangle$ iff $g(w)$ is a purring cat. This entails (given our pragmatic system) that they have the same update effect; and that they have the same truth-value whenever both are defined (they are *Strawson equivalent*).

To see the central point—that these have identical truth-conditions—consider any point $\langle g, w \rangle$ and context c . Suppose first that (27) is true at that point, relative to c . Recall that indefinites are true just in case the corresponding open sentence is, which means that $g(x)$ is a cat in w . The truth of the right conjunct guarantees that $g(x)$ is also purring in w . But that means that (28) is true too, since the corresponding open sentence is. Suppose next that (28) is true; then the corresponding open sentence is, so $g(x)$ is a purring cat in w . That means that the left conjunct of (27) is true. For the conjunction to be true at $\langle c, g, w \rangle$, $P(\text{!}x(\text{T}x))$ must be true at $\langle c^{\text{some}_x(Cx)}, g, w \rangle$. For that to hold, (i) x must be defined throughout $c^{\text{some}_x(Cx)}$ and (ii) $g(x)$ must be purring. But (i) will hold, since by definition $c^{\text{some}_x(Cx)}$ is the set of points $\langle g', w' \rangle$ from c where $\text{some}_x(Cx)$ is true, which are all points where g' is defined on x . And (ii) will hold given that (28) is true.

More generally, sentences with the form $\text{some}_x(p) \wedge q$ and $\text{some}_x(p \wedge q)$ will always have the same truth-conditions. For $\text{some}_x(p) \wedge q$ is true at $\langle c, g, w \rangle$ iff $\text{some}_x(p)$ is true at $\langle c, g, w \rangle$ and q is true at $\langle c^{\text{some}_x(p)}, g, w \rangle$, iff p is true at $\langle c, g, w \rangle$ and q is true at $\langle c^{\text{some}_x(p)}, g, w \rangle$. $\text{some}_x(p \wedge q)$ is true at $\langle c, g, w \rangle$ iff $p \wedge q$ is true at $\langle c, g, w \rangle$ iff p is true at $\langle c, g, w \rangle$ and q is true at $\langle c^p, g, w \rangle$. But now note that $c^{\text{some}_x(p)} = c^p$, since the points in c where $\text{some}_x(p)$ are true are precisely those where p is true. So these truth-conditions are the same. The same reasoning extends to pairs like (27) and (28), which have the form $\text{some}_x(p) \wedge q$ and $\text{some}_x(p \wedge q')$ where q' differs from q by replacing a definite with a corresponding bare variable. As long as the prejacent of the definite is entailed by the indefinite $\text{some}_x(p)$, as in this case, this reasoning will go through unchanged.

Since truth is what matters for pragmatics in our system (that is, if two assertions have the same truth-conditions, they induce the same updates), I'm inclined to think this is enough to account for the intuitions that motivate open scope.¹⁷

5.5 Negation and disjunction

Our account, crucially, avoids the problems with negation and disjunction for dynamic semantics which we explored above. Start with negation. First, note that negated indefinites have the intuitively correct meaning. $\neg \text{some}_x(Cx)$ is true at $\langle c, g, w \rangle$ just in case there is no cat in w . Hence we capture the intuition that negated indefinites have strong, universal

¹⁷It is, of course, true in our system that sentences like (27) and (28) have different falsity-conditions, and thus that their negations have different truth-conditions. But 'It's not the case that there is a purring cat' and 'It is not the case that there is a cat and it is purring' do seem to differ slightly, though intuitions are extremely subtle.

meanings. This does not require any kind of extra closure operator, but is simply thanks to the trivalent semantics for indefinites: on our approach, indefinites have the *truth* conditions of the corresponding open sentence but the *falsity* conditions of the corresponding existential quantifier.

So consider doubly negated indefinites. $\neg\neg\text{some}_x(Cx)$ is semantically equivalent to $\text{some}_x(Cx)$: this is obvious from our Kleene semantics for negation, which switches true to false and false to true (and leaves undefinedness unchanged). So doubly negated indefinites will have the same meaning as matrix indefinites, and hence are predicted to license anaphora, as desired.

One perspective is that we have simply folded Rothschild’s dagger operator into the falsity conditions of indefinites, rather than positing it as a syntactically separate entity. That means that we don’t face the question of when the dagger operator is inserted. This also brings out clearly the precedent for the present approach in Krahmer and Muskens’ proposal, which is based on the same basic idea.

Next, consider a bathroom disjunctions like ‘Either there is no bathroom, or it is upstairs’. The obvious parse, in (29), yields the desired truth-conditions: there is no copying or closure needed.

$$(29) \quad \neg\text{some}_x(Bx) \vee U(\text{!}x(\top x))$$

Consider an arbitrary point $\langle g, w \rangle$ in context c . Suppose first that w is bathroom-free. Then the left disjunct will be true, and so (29) will be true, by our middle Kleene connectives. Suppose next that w contains a bathroom. Then, since the left disjunct is not true, the sentence is true just in case the left disjunct is false and the right disjunct is true. The left disjunct is false just in case $g(x)$ is a bathroom in w . In that case, the right disjunct will be defined, because (i) $g(x)$ is defined and (ii) the definite’s familiarity constraint is satisfied. To see the latter point, note that the right disjunct’s local context is $c^{\neg\neg\text{some}_x(Bx)}$. Since double negation elimination is valid, this is equal to $c^{\text{some}_x(Bx)}$, which will only contain pairs which assign x to a bathroom and hence which are defined on x . Given all this, the right disjunct is true just in case $g(x)$ is upstairs. In sum: (29) is true at $\langle c, g, w \rangle$ iff w is bathroom-free, or else $g(x)$ is a bathroom which is upstairs. This nicely captures the intuitive truth-conditions (modulo domain restriction, which, again, we set aside here). Things work in just the same way for a version with ‘the bathroom’ in place of ‘it’ in the right disjunct.

This helps bring out why I have adopted the middle Kleene connectives rather than weak or strong Kleene. Suppose we had weak Kleene connectives; then (29) would be undefined whenever $g(x)$ is, even if w is bathroom-free. That would mean that updating with (29) would license subsequent anaphora to x , since the sentence would only be defined at points where $g(x)$ is. But you can’t say ‘There is either no bathroom, or it is upstairs. It has a nice mirror in it.’ By contrast, given middle Kleene, (29) will be true, not undefined, when $g(x)$ is undefined, provided that there is no bathroom. Strong Kleene connectives would also go wrong here: suppose that w has only downstairs bathrooms, but $g(x)$ is something upstairs.

Then on the strong Kleene semantics, (29) would wrongly be predicted true at $\langle c, g, w \rangle$, since the right disjunct would be true.

Besides bringing out clearly the motivations for our approach to the connectives, this case helps bring out the motivation for our semantics for indefinites. The semantics looks haphazard at first. But we need something along precisely these lines if we are to capture the open scope of indefinites, while also capturing the strong meaning of negation. One way to get at the intuition behind this semantics is via a reformulation that Keny Chatain (p.c.) has suggested to me: we can equivalently think about this semantics as giving indefinites the semantics of existential quantifiers, together with a conditional definedness condition (what I call the *witness bound* in Mandelkern 2022b). The witness bound says that a sentence like $\text{some}_x(Cx)$ is only defined at $\langle g, w \rangle$ if the following conditional holds: *if* there is any cat in w , then $g(x)$ is a cat in w . That is, *if* there's a witness to the indefinite, then we make sure we keep track of all the witnesses with the context's variable assignments. Otherwise, we don't have anything to keep track of. Essentially, indefinites are still existential quantifiers—just with the added requirement that they must keep track of witnesses to their truth, if there are any.

5.6 Quantifiers

I turn, finally, to quantifiers. In the classical approach, quantifiers quantify over assignments which vary just the value of the variable they bind. Donkey sentences suggest we need more flexibility than this. Consider the covariation between, say, cats and cat-owners in 'Everyone _{x} who _{x} has a cat _{y} loves it _{y} '. In order to capture that covariation, quantifiers need to range over assignments which vary the value both of variables they bind (in this case, x) and of variables bound by indefinites in their scope (here, y). But we cannot quantify over *all* assignments, since we need to hold fixed the values of already-familiar definites. Hence, in a sentence like (30), 'him' is interpreted as being anaphoric on the cat I have, and cannot covary with 'everyone':

(30) I have a cat. Everyone loves him.

To get the right amount of variation to capture both of these facts, we will let our quantifiers range over assignments which hold fixed the values of all the variables already familiar in c . Write $g'[c]g$ just in case g' and g agree on all variables familiar in c .

A final wrinkle is that we want the quantifier's variable to count as familiar within the scope of the quantifier, so that it won't get re-shifted by embedded quantifiers. That is, when we have a sentence like 'Everyone _{x} loves everyone _{y} ', we don't want to re-shift the value of x when we arrive at the second quantifier. So we will evaluate the restrictor and scope relative to a minimal variant on the input context which makes the quantifier's variable familiar; write c_x for the context which is just like c but without any assignments where x is undefined.

With this in hand, we can extend the Kleene idea to quantifiers as follows:

$$\llbracket \text{every}_x(p, q) \rrbracket^{c, g, w} = \begin{cases} 1 & \forall o : \exists g'[c]g : \llbracket p \rrbracket^{c_x, g'_{[x \rightarrow o]}, w} \neq \# \text{ and} \\ & (\exists g'[c]g : \llbracket p \rrbracket^{c_x, g'_{[x \rightarrow o]}, w} = 1) \rightarrow \exists g''[c]g : \llbracket p \wedge q \rrbracket^{c_x, g''_{[x \rightarrow o]}, w} = 1 \\ 0 & \exists o : \exists g'[c]g : \llbracket p \wedge \neg q \rrbracket^{c_x, g'_{[x \rightarrow o]}, w} = 1 \\ \# & \text{otherwise} \end{cases}$$

The core is the classical idea that $\text{every}_x(p, q)$ is true iff every o which is p is also q , and false iff some o is p and $\neg q$. We add a natural extension of the Kleene idea to quantifiers: namely, the quantified sentence cannot be true if, for some o , it is indeterminate (on all relevant assignments) whether o is p (compare George 2008).

For quantifiers not embedding (in)definites, the result will be equivalent to the classical treatment of quantifiers. In, e.g., $\text{every}_x(Cx, Px)$ ('every cat is purring'), our semantics tells us to simply check whether every cat is purring; the sentence is true if so, false otherwise.

As for donkey quantification, the central difference to Rothschild's approach is, again, that we don't need any extra material in the logical form. Consider 'Everyone who has a cat loves it', as in (31), assessed at $\langle c, g, w \rangle$:

$$(31) \quad \text{every}_x(\text{some}_y(Cy \wedge H(x, y)), L(x, \iota y(\top y)))$$

We must consider every individual o in the domain, and check whether the restrictor $\text{some}_y(Cy \wedge H(x, y))$ is true at $\langle c_x, g'_{[x \rightarrow o]}, w \rangle$ for any $g'[c]g$. As long as y is not already familiar in c , g' can vary the value of y . Hence the restrictor will be true at some such $\langle c_x, g'_{[x \rightarrow o]}, w \rangle$ just in case o is a cat-owner. If this holds, we must then check whether the conjunction of the restrictor and scope is true at $\langle c_x, g''_{[x \rightarrow o]}, w \rangle$ for some $g''[c]g$. Once more, we get to vary y , and so this holds just in case there is some way of assigning y to a cat which is both owned and loved by o in w . The quantified sentence is thus true just in case every cat-owner in w is also a cat-lover.

Crucially, we get covariance between the definite in the scope and the indefinite in the restrictor; and the definite in the scope is guaranteed to be licensed, since its local context entails the scope, which contains an indefinite indexed to y . A definite description 'the cat' would likewise be licensed in the same place. By contrast, without a corresponding indefinite, we won't get covariance. 'Every cat-owner loves her_y' will be undefined if y is novel in the input context, thanks to the definedness constraint of definites. If y is already familiar, then the sentence gets a non-covarying reading, as desired, thanks to our definition of $g'[c]g$.

Like Rothschild, we get the weak reading of donkey sentences. Again, I'm not sure what to make of the weak/strong issue, and as far as I can tell, the present approach does not shed special light on it (though there are different, trivalent approaches which may, as discussed in Champollion et al. 2019); a close variant would get us the strong reading instead.

Donkey quantification is, of course, not restricted to universal quantifiers, as (32) illustrates:

(32) $\left\{ \begin{array}{c} \text{Most} \\ \text{Few} \end{array} \right\}$ people who have a cat love it.

Extending our treatment to these constructions is conceptually straightforward. For instance, ‘most’ can be treated as follows (ignoring definedness conditions, which will be as for *every*):

$$\llbracket \text{most}_x(p, q) \rrbracket^{c, g, w} = 1 \text{ iff } \frac{|\{o : \exists g'[c]g : \llbracket p \wedge q \rrbracket^{c_x, g'_{[x \rightarrow o]}, w} = 1\}|}{|\{o : \exists g'[c]g : \llbracket p \rrbracket^{c_x, g'_{[x \rightarrow o]}, w} = 1\}|} > \frac{1}{2}$$

6 Conclusion

In the light of Partee pairs and donkey sentences, pretty much everyone rejects the classical theory of anaphora and its interaction with quantifiers and connectives. After that, there is a lot of disagreement. E-type approaches defend the classical treatment of (in)definites, by taking on revisionary semantics of the quantifiers and connectives. Dynamic semantics departs from classical assumptions in its treatment of (in)definites as well. Trivalence is (implicitly) a central part of some dynamic systems. But in the dynamic literature, trivalence has tended to be at least somewhat in the background, overshadowed by other, more revisionary features of dynamic semantics. More recently, trivalence has come into its own as a central player in theories of anaphora.

The trivalent approach that I developed in the last section is closely influenced by dynamic semantics (especially Krahmer and Muskens (1995)’s bilateral system). It borrows from dynamic semantics a central tool, namely, that of local contexts. Indeed, depending on your preferred criterion for dynamicity, this may lead you to classify that system as itself dynamic. The system is obviously very different from standard dynamic semantics. Nonetheless, it is worth asking whether a system like the one I have presented here can be made more purely static, by eschewing entirely the apparatus of semantically specified local contexts (e.g., by using the tools developed in Schlenker 2008). This is the focus of exciting new work in Spector 2021.

A related question concerns whether trivalence is exactly the right tool for thinking about the partiality involved in anaphora. As I mentioned at the outset, the trivalent system I have developed here is based on a different system which I develop in Mandelkern 2022b. That system is similar in spirit, but has a two-dimensional, bivalent semantics. The ‘main’ dimension is essentially classical, while the witness bound is a part of the second dimension of backgrounded constraints. We can think of two-dimensionality as a kind of four-valuedness, so there is a parsimony consideration in favor of trivalence. On the other hand, the four-valued picture has an appealing division of labor between the main dimension, whose semantics—and hence logic—is classical, and the second dimension, where all the dynamic action takes place.

Another consideration concerns a point raised by Benjamin Spector for the trivalent picture I have developed in this chapter, arising from stacked negated indefinites (p.c., citing a related observation of Amir Anvari). Consider a sentence like ‘There is not a person that didn’t dance with someone’, that is, $\neg \text{some}_x(\neg \text{some}_y(\text{danced}(x,y)))$. On the trivalent approach we’ve developed, this ends up being true at $\langle c, g, w \rangle$ only if everyone danced with $g(y)$. For $\llbracket \neg \text{some}_x(\neg \text{some}_y(\text{danced}(x,y))) \rrbracket^{c,g,w} = 1$ iff $\llbracket \text{some}_x(\neg \text{some}_y(\text{danced}(x,y))) \rrbracket^{c,g,w} = 0$ iff $\forall o : \llbracket \neg \text{some}_y(\text{danced}(x,y)) \rrbracket^{c,g[x \rightarrow o],w} = 0$ iff $\forall o : \llbracket \text{some}_y(\text{danced}(x,y)) \rrbracket^{c,g[x \rightarrow o],w} = 1$ iff $\forall o : \llbracket \text{danced}(x,y) \rrbracket^{c,g[x \rightarrow o],w} = 1$! That is a serious problem, and one that the two-dimensional approach avoids. But different trivalent approaches have resources to deal with this case as well: in particular, the trivalent approach of Spector 2021 avoids this problem by using plural assignment functions. In addition to these points, the two-dimensional approach has a nicer, and more perspicuous, logic than the trivalent approach I have developed.¹⁸ These points lead me to favor a two-dimensional approach over a trivalent approach at present.

There is much more to explore here; what seems clear is that some notion of partiality will play a, and perhaps the, central role in a successful theory of anaphora.

References

- Beaver, D. (2001). Presupposition and Assertion in Dynamic Semantics. CSLI Publications: Stanford, CA.
- van den Berg, M. (1996). Some Aspects of the Internal Structure of Discourse: The Dynamics of Nominal Anaphora. PhD thesis, ILLC, Universiteit van Amsterdam.
- Büring, D. (2004). Crossover situations. Natural Language Semantics, 12:23–62.
- Champollion, L., Bumford, D., and Henderson, R. (2019). Donkeys under discussion. Semantics & Pragmatics, 12(1):1–50.
- Chatain, K. (2017). Local contexts and anaphora. Handout, MIT.
- Chierchia, G. (1992). Anaphora and dynamic binding. Linguistics and Philosophy, 15(2):111–183.
- Chierchia, G. (1995). Dynamics of Meaning: Anaphora, Presupposition, and the Theory of Grammar. University of Chicago Press, Chicago, IL.
- Cooper, R. (1979). The interpretation of pronouns. In Heny, F. and Schnelle, H. S., editors, Syntax and Semantics, volume 10, pages 61–92. Academic Press.
- Egli, U. (1979). The Stoic concept of anaphora. In Bäuerle, R., Egli, U., and von Stechow, A., editors, Semantics from Different Points of View, pages 266–283. Springer, New York.

¹⁸The basic issue for the present system is that what variables quantifiers get to vary depends on context, and the context changes over the course of the sentence, meaning there will be failures of the law of non-contradiction and excluded middle, as in many dynamic systems (e.g., the following is consistent: ‘Everyone has a donkey_y and something_y is feline, and it’s not the case that (everyone has a donkey_y and something_y is feline)’). This might be separable from other parts of the system, for instance if we added a novelty presupposition to the indefinite, or if we treated donkey anaphora using functional pronouns, following Schlenker 2011; Chatain 2017, rather than the more flexible kind of quantification we have introduced here.

- Elbourne, P. (2005). Situations and Individuals. MIT Press.
- Elliott, P. D. (2020a). Crossover and accessibility in dynamic semantics. Manuscript, MIT.
- Elliott, P. D. (2020b). Towards a principled logic of anaphora. <https://ling.auf.net/lingbuzz/005562>.
- Evans, G. (1977). Pronouns, quantifiers and relative clauses. Canadian Journal of Philosophy, 7:467–536.
- von Fintel, K. (2008). What is presupposition accommodation, again? Philosophical Perspectives, 22(1):137–170.
- Geach, P. (1962). Reference and Generality. Cornell University Press.
- George, B. R. (2008). Presupposition repairs: A static, trivalent approach to predicting projection. Master's thesis, University of California, Los Angeles.
- Gotham, M. (2019). Double negation, excluded middle and accessibility in dynamic semantics. In Schlöder, J. J., McHugh, D., and Roelofsen, F., editors, Proceedings of the 22nd Amsterdam Colloquium, volume 22, pages 142–151.
- Heim, I. (1982). The Semantics of Definite and Indefinite Noun Phrases. 2011 edition, University of Massachusetts, Amherst.
- Heim, I. (1983). On the projection problem for presuppositions. In Barlow, M., Flickinger, D. P., and Wiegand, N., editors, The West Coast Conference on Formal Linguistics (WCCFL), volume 2, pages 114–125. Stanford, Stanford University Press.
- Heim, I. (1990). E-type pronouns and donkey anaphora. Linguistics and Philosophy, 13:137–177.
- Kamp, H. (1981). A theory of truth and semantic representation. In et al., J. G., editor, Formal Methods in the Study of Language. Mathematisch Centrum.
- Kanazawa, M. (1994). Weak vs. strong readings of donkey sentences and monotonicity inference in a dynamic setting. Linguistics and Philosophy, 17(2):109–158.
- Karttunen, L. (1973). Presuppositions of compound sentences. Linguistic Inquiry, 4(2):167–193.
- Karttunen, L. (1974). Presupposition and linguistic context. Theoretical Linguistics, 1(1-3):181–194.
- Karttunen, L. (1976). Discourse referents. In Syntax and semantics: Notes from the linguistic underground, volume 7. Academic Press.
- Krahmer, E. and Muskens, R. (1995). Negation and disjunction in discourse representation theory. Journal of Semantics, 12:357–376.
- Lewerentz, L. (2020). Situation pragmatics. Manuscript, Oxford.
- Lewis, K. (2012). Discourse dynamics, pragmatics, and indefinites. Philosophical Studies, 158:313–342.
- Lewis, K. S. (2019). Descriptions, pronouns, and uniqueness. MS, Columbia University.
- Ludlow, P. (1994). Conditionals, events, and unbound pronouns. Lingua e Stile, 29:3–20.
- Mandelkern, M. (2022a). Bounds. Manuscript, NYU.
- Mandelkern, M. (2022b). Witnesses. Linguistics and Philosophy. 10.1007/s10988-021-09343-w.
- Mandelkern, M. and Rothschild, D. (2020). Definiteness projection. Natural Language

- Semantics, 28:77–109.
- Neale, S. (1990). Descriptions. MIT Press, Cambridge, MA.
- Parsons, T. (1978). Pronouns as paraphrases. MS University of Massachusetts, Amherst.
- Roberts, C. (1987). Modal Subordination, Anaphora and Distributivity. PhD thesis, University of Massachusetts Amherst.
- Root, R. (1986). The Semantics of Anaphora in Discourse. PhD thesis, University of Texas at Austin.
- Rooth, M. (1987). Noun phrase interpretation in Montague Grammar, File Change Semantics, and Situation Semantics. In Gärdenfors, P., editor, Generalized quantifiers: Linguistic and Logical Approaches, pages 237–268. Dordrecht Reidel, Netherlands.
- Rothschild, D. (2017). A trivalent approach to anaphora and presupposition. In Cremers, A., van Gessel, T., and Roelofsen, F., editors, Amsterdam Colloquium, volume 21, pages 1–13.
- Schlenker, P. (2008). Be articulate: a pragmatic theory of presupposition projection. Theoretical Linguistics, 34(3):157–212.
- Schlenker, P. (2011). On donkey anaphora. Handout of talk given at the Oxford-Paris Workshop, All Souls College, Oxford.
- Schubert, L. K. and Pelletier, F. J. (1989). Generically speaking, or, using discourse representation theory to interpret generics. In Chierchia, G., Partee, B., and Turner, R., editors, Properties, Types, and Meaning, Volume II: Semantic Issues 39, Studies in Linguistics and Philosophy, pages 193–268. Kluwer, Dordrecht, Netherlands.
- Soames, S. (1989). Presupposition. In Gabbay, D. and Guenther, F., editors, Handbook of Philosophical Logic, Volume IV: Topics in the Philosophy of Language, pages 553–616. Kluwer Academic Publishers.
- Spector, B. (2021). A non-dynamic approach to anaphora based on Trivalence and Transparency. Manuscript, Ecole Normale Supérieure.
- Stalnaker, R. (1974). Pragmatic presuppositions. In Munitz, M. K. and Unger, P., editors, Semantics and Philosophy, pages 197–213. New York University Press, New York.
- Stone, M. (1992). *Or* and anaphora. In Barker, C. and Dowty, D., editors, Semantics and Linguistic Theory (SALT), volume 2, pages 367–385.