Witnesses

Matthew Mandelkern[*]

Penultimate version; to appear in *Linguistics & Philosophy*

Abstract

The meaning of definite descriptions (like 'the King of France', 'the girl', etc.) has been a central topic in philosophy and linguistics for the past century. Indefinites ('Something is on the floor', 'A child sat down', etc.) have been relatively neglected in philosophy, under the Russellian assumption that they can be unproblematically treated as existential quantifiers. However, an important tradition, drawing from Stoic logic, has pointed to patterns which suggest that indefinites cannot be treated simply as existential quantifiers. The standard dynamic semantic treatment of those phenomena, however, has well-known problems with negation and disjunction.

In this paper I develop a new approach to (in)definites. On my theory, truth-conditions are classical. But in addition to truth-conditions, meanings comprise a second dimension of what I call *bounds*. It is at the level of bounds, not truth-conditions, that I locate the characteristically dynamic coordination between indefinites and definites. The resulting system thus has a classical logic. This approach avoids dynamic semantics' logical problems, and, more generally, yields a new perspective on the relation between truth-conditional and dynamic effects in natural language.

1 Introduction

The meaning of definite descriptions (like 'the King of France', 'the girl', etc.) has been a central topic in philosophy and linguistics for the past century. Indefinites ('Something is on the floor', 'A child sat down', etc.) have been relatively neglected by philosophers, under the Russellian assumption that they can be unproblematically treated as existential

* NYU Philosophy; mandelkern@nyu.edu. This paper was, even more than usual, a collaborative affair, and I am very grateful to audiences at UCL, the ILLC, the NYU Semantics Group, and the Virtual Language Work in Progress group; to participants in seminars at Oxford and NYU; and to Chris Barker, Kyle Blumberg, David Boylan, Simon Charlow, Cian Dorr, Patrick Elliott, Peter Fritz, Matthew Gotham, Ezra Keshet, Nathan Klinedinst, Lukas Lewerentz, Karen Lewis, Milo Phillips-Brown, Craige Roberts, Philippe Schlenker, Ginger Schultheis, Benjamin Spector, and Anna Szabolcsi for very helpful discussion. Thanks to two fantastic reviewers and to my editor Paul Dekker for extremely helpful comments. And thanks to Keny Chatain, Daniel Rothschild and Yasu Sudo for crucial discussion and help with key points throughout the project.

quantifiers. However, a different tradition, drawing from Stoic logic,[1] has pointed to patterns which suggest that indefinites are not well-modeled as existential quantifiers. To see the basic issue, compare (1-a) and (1-b):

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

b. Everyone who is a parent loves $\left\{\begin{array}{c}\text{them}\\\text{the child}\end{array}\right\}$.

(1-a) can naturally be interpreted as saying that every parent loves their child; while (1-b) can only be naturally interpreted as saying that some salient child (or person, with 'them') is loved by every parent. But if indefinites are just existential quantifiers, then 'has a child' and 'is a parent' should mean exactly the same thing: for under that assumption, y has a child iff for some x, x stands in the child-of relation to y, iff y stands in the parent-of relation to some x, iff y is a parent. But, if 'has a child' and 'is a parent' mean the same thing, it is hard to see how we could account for differences in how they embed in pairs like (1). Minimal pairs like this one form the heart of the case against the identification of indefinites with existential quantifiers.

There are two broad responses to this puzzle. *E-type* approaches argue that indefinites are existential quantifiers after all; contrasts like those in (1) are explained by the syntactic differences between the two indefinite expressions. In particular, (1-a) makes salient a predicate ('child') missing from (1-b), which can then—the thought is—be recruited to license subsequent anaphora.[2] *Dynamic* approaches instead argue that the behavior of definites and indefinites shows that meanings are more fine-grained than truth-conditions: sentential meanings are not (characteristic functions of) sets of indices, but instead functions from contexts to contexts. And, in particular, while 'x has a child' and 'x is a parent' have, in some sense, the same 'worldly' content, they have different meanings, because they update contexts differently: only the first yields a context which supports subsequent anaphora to a child.[3]

1 See Egli 1979 for some history.

2 See e.g. Geach 1962, Evans 1977, Parsons 1978, Cooper 1979, Neale 1990, Heim 1990, Ludlow 1994, Büring 2004, Elbourne 2005; see Lewis 2012, 2019, Mandelkern and Rothschild 2020, Lewerentz 2020 for more recent developments and criticism.

3 E.g. Karttunen 1976, Kamp 1981, Heim 1982, Groenendijk and Stokhof 1991, Dekker 1993, 1994, van den Berg 1996, Muskens 1996, Aloni 2001, Beaver 2001, Nouwen 2003, Brasoveanu 2007, Charlow 2014.

Both of these approaches require revisionary approaches to content: in the dynamic approach, they are, again, functions from contexts to contexts rather than sets of indices; in the e-type approach, they are sets of situations or events. And both approaches must adopt non-classical treatments of the connectives.[4] The proper treatment of patterns like those in (1) thus raises central questions, not just about the meaning of (in)definites, but also about the nature of content and the logic of natural language.

In this paper I develop a new approach, which answers these questions differently from either dynamic or e-type theories. On my *bounded theory*, contents are sets of indices, connectives are classical, and indefinites have the truth-conditions of existential quantifiers; in this respect, my theory is very classical. But, on the bounded theory, indefinites also have a secondary dimension of meaning: they have what I call a *witness bound* which requires that, if the indefinite is true, then a witness to its truth is assigned to the indefinite's variable. Definites have as their truth-conditional content simply the conjunction of the content of their scope and restrictor. But they, too, have a bound which requires that the content of their scope be familiar—that is, true throughout the definite's local context. This approach builds on the insights of existing theories, especially dynamic semantics. But by locating the coordination of indefinites and definites in a dimension of meaning separate from truth-conditions—namely, the dimension of bounds—I avoid well-known problems that arise from the revisionary logical framework of dynamic semantics. More generally, the resulting system yields a new perspective on the source and location of dynamic effects in natural language.

2 Problems for the classical picture

I will begin by summarizing the motivation for departing from the classical approach to (in)definites, and explaining the dynamic response and some problems with it, before developing my own system.

On the classical treatment, again (due to Frege, Russell, Quine, and Strawson) an indefinite sentence like ⌜Some F is G⌝ is equivalent to $\exists x(Fx \& Gx)$, where $\exists$ is the

4 I will explain this presently vis-à-vis dynamic semantics; in the e-type literature, the connectives have not been much discussed, but in Mandelkern and Rothschild 2020, we show that empirically adequate e-type approaches must likewise adopt non-classical connectives.

classical existential quantifier.[5] Definite descriptions have the same meaning plus a uniqueness inference: so ⌜The F is G⌝ says that something is both F and G, and also says (or presupposes) that there is exactly one (relevant) F-thing. Pronouns, which also fall under the broad heading of 'definite', are treated as variables.

Given these assumptions, and given the intuitive meaning of 'parent' and 'child', 'Sue is a parent' and 'Sue has a child' will mean exactly the same thing. But they seem to pattern differently in terms of their interaction with subsequent definites (descriptions and pronouns). We have already seen this in the context of quantifiers in (1) (so-called *donkey sentences*). We can also see this in an even simpler way by looking at minimal pairs like (2):

(2) a. Sue has a child. She is at boarding school.
  b. Sue is a parent. She is at boarding school.

Only (2-a) has a natural interpretation which says that Sue has a child at boarding school; (2-b) is naturally interpreted as saying that Sue herself is at boarding school. (This is not to say (2-b) is impossible to interpret in the same way as (2-a); the observation is rather that there is a striking *contrast* in the availability of these interpretations between (2-a) and (2-b) which needs to be accounted for.) This pair is modeled on Partee's famous marble sentence from Heim 1982, so I'll call pairs like this *Partee pairs*. For another example, compare:

(3) a. Sue has a twin. She lives in Dubuque.
  b. Sue is a twin. She lives in Dubuque.

Again, having a twin and being a twin are presumably the same property. Yet these indefinites set up very different anaphoric potential: (3-a) is naturally heard as saying that Sue has a twin who lives in Dubuque, while (3-b) is naturally heard as saying that Sue is a twin who lives in Dubuque.

Note that Partee pairs cannot be accounted for simply by saying that the indefinite in these cases takes wide-scope over the two sentences and binds 'she'. First, the idea that indefinites can take scope over whole discourses—sequences of sentences—

5 (Corner) quotes are omitted around expressions of formal language.

would already require a huge revision to the classical approach to logic. Second, more importantly, this would be a local solution to a very global problem, which would do nothing to help explain corresponding contrasts in other embedded environments, like the donkey sentences in (1), repeated here, or the corresponding examples with 'is a twin'/'has a twin':

(1) a. Everyone who has a child loves {them / the child}.
b. Everyone who is a parent loves {them / the child}.

(4) a. Everyone who has a twin loves {them / the twin}.
b. Everyone who is a twin loves {them / the twin}.

To see the problem, focus on the variant of (1-a) with 'them'. The classical assumptions above would yield (5) as the gloss for (1-a) (with $\forall$ the classical universal quantifier and $\rightarrow$ the material conditional):

(5) $\forall x((\exists y(\textit{child-of}(y,x))) \rightarrow \textit{loves}(x,y))$

The problem with (5) is that the variable y in the consequent is unbound, so we don't get the intended covariation between 'a child' and 'them'. A natural thought is that we could give the existential quantifier wide scope over the material conditional, but that doesn't help: while y would end up bound, we would get absurdly weak truth-conditions for (1-a) (which would end up being true provided that no one is everyone's parent). So it is not at all obvious how to derive the intended meaning of (1-a) given the classical assumptions above.

Donkey sentences with definite descriptions rather than pronouns raise slightly different, but equally serious, issues, as Heim (1982) showed. Consider the variant on (1-a) with a definite description:

(6) Everyone who has a child loves the child.

Given the classical picture above, (6) will say that every parent loves their child, but will also assert (or presuppose, depending which version of the classical picture you adopt) that every parent has exactly one child. Clearly, though, (6) does not lead to a uniqueness inference like that. This is brought out clearly by a variant on Heim's sage-plant sentence, in (7):

(7) Everyone who bought a sage plant bought seven others along with the sage plant.

(7) clearly does not license the incoherent inference that everyone who bought a sage plant bought exactly one sage plant.

3 Dynamic semantics

Dynamic semantics responds to these problems by proposing that predicates like 'has a child' and 'is a parent' have different meanings, and adopting a corresponding revisionary treatment of definites, and, more generally, logical connectives and contents. To explain the basic ideas behind this response—which, in turn, form the background for my own account—I will informally sketch a simplified version of Heim 1982's dynamic system. There are many dynamic systems in the literature; the one I will present gives a reasonably representative sense of the basic ideas.[6]

On this approach, sentence meanings are functions from contexts to contexts. A context, in turn, is a set of pairs of (possibly partial) variable assignments and worlds. The worlds represent the possibilities treated as live in the conversation (the conversation's *context set*, in the sense of Stalnaker 1974). The variable assignments track anaphoric relations between indefinites and definites.

The role of indefinites is to extend the contextual variable assignments so they are defined on a new variable; the role of definites is to pick up on a variable that has been introduced this way. So, as long as x is novel in c—that is, nowhere defined in c—'There is a_x $\text{cat}(x)$' denotes the function that takes c and extends every variable assignment in c to an assignment which assigns x to a cat. More carefully, the resulting

6 I will not extensively discuss or criticize e-type approaches here. This is largely because to my knowledge, e-type approaches have not been able to deal with the basic data that I take to be our explananda, namely Partee pairs. The issue is that in a pair like (3), the relationship of being a twin is equally salient, linguistically and cognitively, in both constructions, but only 'has a twin', not 'is a twin', licenses subsequent anaphora to the twin. Of course there are dialectical moves available, but dynamic semantics has been developed much more explicitly, and, I think, more successfully, and so makes a more effective foil.

context will be the set of pairs $\langle g, w\rangle$ such that $g(x)$ is a cat in w, and for some pair $\langle g', w'\rangle \in c$, $w = w'$, and g' and g agree everywhere except on x.

This captures the idea that indefinites 'open a file card', in Heim's metaphor, or a 'discourse referent', in Karttunen's: an indefinite indexed to x extends a context c so that every variable assignment is defined on x, and, in particular assigns x to a witness of the indefinite.

That, in turn, makes possible subsequent anaphora with co-indexed definites. Definites are used to talk about variables that have already been assigned to a witness throughout the context. In the Heimian metaphor: they are used to add information to 'file cards' that have already been opened (that are already "familiar"). So, for instance, 'The$_x$ cat(x) is sleeping(x)' presupposes that x is defined throughout c and moreover assigned to a cat through c. Where this presupposition is satisfied, 'The$_x$ cat(x) is sleeping(x)' updates c so that the resulting context includes exactly the pairs $\langle g, w\rangle \in c$ where $g(x)$ is a sleeping cat in w. Pronouns are treated analogously, but with only the requirement that x is defined throughout the input context: so, for instance, 'He$_x$ is sleeping(x)' presupposes that x is defined throughout c, and, where that presupposition is satisfied, takes c to the set of pairs $\langle g, w\rangle \in c$ such that $g(x)$ is sleeping in w.

Updating with an indefinite thus sets the stage for subsequent updates with definites: indefinites open file cards that definites can then add to. So suppose we have updated our context with 'There is a$_x$ cat(x)'. This guarantees that, at every point $\langle g, w\rangle$ in the updated context, $g(x)$ is a cat in w. So the familiarity presupposition of a definite like 'The$_x$ cat(x)' is satisfied, and so subsequent updates with 'The$_x$ cat(x) is sleeping(x)' will be defined, and will take us to the context comprising just those pairs $\langle g, w\rangle$ where $g(x)$ is a sleeping cat. By contrast, a definite out of the blue will not have its familiarity presupposition satisfied, since (we assume) default starting contexts include, for any variable x, assignments where x is undefined.

With this in hand, let's return to the Partee pairs we saw above. Focus in particular on (1). Consider any context c. Updating with either 'Sue has a$_x$ child(x)' or 'Sue is a$_x$ parent(x)' will have the same *worldly* effect on c: a world in some pair in c will be in some pair in the updated context just in case Sue is a parent in that world. But these updates have very different effects on variable assignments. Updating c with 'Sue has

a_x child(x)' will result in a context comprising exactly the pairs $\langle g, w \rangle$ such that $g(x)$ is a child of Sue's in w and for some $\langle g', w' \rangle \in c$, $w = w'$ and g' agrees with g except (possibly) on x. By contrast, updating with 'Sue is a_x parent(x)' will instead result in a context which comprises exactly the pairs $\langle g, w \rangle$ such that $g(x)$ is Sue, Sue is a parent in w, and for some $\langle g', w' \rangle \in c$, $w = w'$ and g' agrees with g except (possibly) on x. So the first update opens a file on Sue's child (indexed to x). The second instead opens a file on Sue.

That means that only the first update licenses subsequent anaphora to Sue's child; the second, by contrast, licenses anaphora only to Sue. A context that has been updated with 'Sue has a_x child(x)' can then be updated with 'She$_x$ is at boarding school(x)'; the resulting context will comprise exactly the pairs from the intermediate context which take x to a child of Sue's who is at boarding school in the world of the pair. By contrast, while a context that has been updated with 'Sue is a_x parent(x)' can be updated with 'She$_x$ is at boarding school(x)', the 'she' here will be felt to refer to Sue, not her child: the subsequent context will comprise exactly the pairs from the intermediate context such that Sue is is at boarding school in the world of the pair. This is the heart of the dynamic account of Partee pairs.

This story extends to conjunctive versions of Partee pairs, like those in (8):

(8) a. Sue has a child and she is at boarding school.
    b. Sue is a parent and she is at boarding school.

Conjunction, like sequential assertion, is treated by dynamic semantics as successive context update, first with the meaning of the left conjunct and then the right. That is: where $[p]$ is the dynamic meaning of any sentence p, and $c[p]$ is the application of $[p]$ to context c, the dynamic conjunction says $c[p \& q] = (c[p])[q]$. So, 'Sue has a_x child(x) and she$_x$ is at boarding school(x)' first takes c to a context comprising just pairs $\langle g, w \rangle$ where $g(x)$ is a child of Sue's, then further updates this context by keeping just those pairs $\langle g, w \rangle$ where $g(x)$—that is, Sue's child—is at boarding school. By contrast, 'Sue is a_x parent(x) and she$_x$ is at boarding school(x)' takes c to a context comprising pairs $\langle g, w \rangle$ where $g(x)$ is a parent and is Sue; then further updates this context by keeping just those points $\langle g, w \rangle$ where $g(x)$—that is, Sue—is at boarding school.

A general way of characterizing the dynamic account of Partee pairs is to say that, in dynamic semantics, indefinites have *open scope* to their right (Egli 1979): they can "bind" co-indexed definites, whether or not the definite is in their syntactic scope. Schematically:

- *Open scope of indefinites*: The following are equivalent:

(9) a. Some F is G and H.
b. Some F is G, and $\left\{ \substack{\text{it} \\ \text{the } F} \right\}$ is H.
c. Some F is G. $\left\{ \substack{\text{It} \\ \text{The } F} \right\}$ is H

So, for instance, 'Sue has a child, and she is at boarding school' will be equivalent to 'Sue has a child who is at boarding school'; while 'Sue is a parent, and she is at boarding school' will be equivalent instead to 'Sue is a parent who is at boarding school'.

This explanation of Partee pairs extends naturally to the corresponding contrasts in donkey sentences. For reasons of space I won't spell this out in detail, but the intuition is the same: predicates like 'has a child' and 'is a parent' update contexts differently as far as assignments go, making available different possibilities for subsequent anaphora.

4 Logical problems

Dynamic semantics thus has a compelling account of Partee pairs. It also has well-known problems involving negation (Karttunen 1976) and disjunction (Heim 1982), which I will explain in this section.

The problem, abstractly, is that the logic of dynamic semantics is non-classical in ways that do not match intuition. In classical logic, $\neg\neg p$ and p are equivalent. Likewise, $\neg p \vee q$ is equivalent to $\neg p \vee (p \& q)$. These equivalences fail, in problematic ways, in dynamic semantics.

To work up to the problem, let's start by thinking about how to extend our dynamic system to negation. A natural first thought is that $c[\neg p] = c \setminus c[p]$. This doesn't work. Consider a negated sentence like (10):

(10) It's not the case that Sue has a_x child(x).

Negated indefinites have strong truth-conditions: (10) intuitively communicates that Sue is childless. Thus what we want, when we update c with (10), is to keep just those pairs $\langle g,w\rangle$ in c such that Sue has no children in w. But the current proposal gives us something much weaker: in fact, since 'Sue has a_x child(x)' is (wherever defined) a pointwise extension of pairs in c, no pair in c will survive, unchanged, after update with [Sue has a_x child(x)], meaning that c itself would remain unchanged after update with (10).

A more general way of putting the point is that, in dynamic semantics, indefinites, once they extend the relevant variable assignments, are equivalent to open sentences. If indefinites are existential quantifiers, as in the classical approach, the universal meaning of negated indefinites follows from their interaction with Boolean negation. But given a non-quantificational approach to indefinites, if we are to capture the strong meaning of negated indefinites, we need negation itself to quantify over assignments.

Hence the standard approach to negation in dynamic semantics says $c[\neg p]$ is the set of pairs from c which *can't be extended* in any way to be a part of $c[p]$. That is, $c[\neg p] = \{\langle g,w\rangle \in c : \neg\exists g' \geq g : \langle g',w\rangle \in c[p]\}$ (where $g' \geq g$ iff g' and g agree everywhere that g is defined). Given this treatment of negation, when we update c with (10), we keep just those pairs $\langle g,w\rangle$ in c such that no extension of g assigns a child of Sue's in w to x. Provided that x is novel (which is required for the whole update to be defined), that means that we keep exactly those pairs $\langle g,w\rangle$ from c such that Sue has no children in w, as desired.

This approach captures the intuitive truth-conditions of (10). But it has a problematic upshot: double negation elimination is not valid in this system.[7] Because negation quantifies over assignments, 'Not (Not (Sue has a_x child(x)))' doesn't have any effect on assignments in the resulting context. A quick way to see this is that negation is eliminative—$c[\neg p]$ is always a subset of c, where defined—and so, if x is novel in c, it will also be novel in $c[\neg\neg p]$.

7 There are different ways to define validity in dynamic semantics, but double negation elimination is not valid on any of them. The most common approach says that the inference from p to q is valid just in case for any context c, $c[p] = (c[p])[q]$ (where both are defined). $\neg\neg p$ does not entail p in this sense, which would require $c[\neg\neg p]$ to equal $(c[\neg\neg p])[p]$ whenever defined. In particular, when p is 'Sue has a_x child(x)' and x is novel in c, $c[\neg\neg p]$ will be the set of points $\langle g,w\rangle \in c$ where Sue has a child in w, while $(c[\neg\neg p])[p]$ will be the set of points $\langle g,w\rangle$ such that $g(x)$ is a child of Sue's in w and for some $\langle g',w'\rangle \in c : w = w'$ and g agrees with g' except on x.

So updating c with 'Not (Not (Sue has a_x child(x)))' will yield a context containing exactly the pairs $\langle g, w\rangle$ from c where Sue has a child in w. Since this update puts no constraints on variable assignments, it does not set up subsequent anaphora dependencies.

How big of a problem is this? While stacked negations are strange in natural language, there still seems to be a striking contrast in pairs like (11):

(11) 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.

The doubly-negated indefinite in (11-a), like the non-negated indefinite 'Sue has a child', seems to license subsequent anaphora to Sue's child, while the doubly-negated indefinite in (11-b) only seems to naturally license subsequent anaphora to Sue. This is brought out more naturally in exchanges like those in (12):

(12) 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.

So the fact that doubly-negated indefinites do not license anaphora in standard dynamic systems does not seem to match observations.

Furthermore, this problem with negation infects other environments, in particular disjunction. As Partee observed, negated indefinites in left disjuncts license definites in right disjuncts. Compare:

(13) a. Either Sue doesn't have a child, or she's at boarding school.
b. Either Sue isn't a parent, or she's at boarding school.

Only in (13-a) can 'she' be naturally interpreted as referring to Sue's child. But this is not captured by dynamic semantics. The standard dynamic disjunction says $c[p \vee q] = c[p] \cup c[\neg p][q]$. Thus $c[\neg p \vee q] = c[\neg p] \cup c[\neg\neg p][q]$. What we want is for this to come out equivalent to $c[\neg p] \cup c[p][q]$. Then indefinites in p would be accessible to definites in q. But, because double negation elimination is not valid, this equivalence does not hold, and so we won't be able to predict that the pronoun in the right disjunct of (13-a) is licensed by the negated indefinite in the left disjunct.

Schematically, to account for the contrast in (13), we want $\neg p \vee q$ to be equivalent to $\neg p \vee (p \& q)$. Then (13-a) would be equivalent to 'Either Sue doesn't have a child, or she has a child and she is at boarding school'. But, while this equivalence holds in classical logic, it doesn't hold in dynamic systems. Thus negated indefinites in left disjuncts will not in general license definites in right disjuncts, contrary to observation.

5 The bounded theory

These problems with negation and disjunction are worrying enough that I think we should take a skeptical second look at the foundations of dynamic semantics. If we want to follow dynamic semantics in holding that indefinites have open scope to their right—as I think we should—then we need *something* non-classical in our system. But the present problems suggest that dynamic semantics overdoes non-classicality.

In the rest of the paper I will develop a new theory which, like dynamic semantics, predicts the open scope of indefinites—and thus accounts for the key contrasts that motivate dynamic semantics—but which is more conservative, logically and foundationally, than dynamic semantics. I start with the classical treatment of indefinites as existential quantifiers. Then I propose that indefinites with the form ⌜Some$_x$ F is G⌝ have a *witness bound*, operating in a separate dimension from truth-conditions, which requires at $\langle g, w \rangle$ that, if *anything* is F and G in w, then $g(x)$ is. The witness bound ensures that indefinites license subsequent definites, which require that the content in their scope be familiar. Anaphora is thus coordinated in the dimension of bounds. At the level of truth-conditions, however, everything is classical, which means that my approach avoids the problems just sketched for dynamic semantics.

5.1 Relation to the existing literature

Before presenting my theory, let me briefly situate it in the existing literature. First, at a high level: my system crosscuts the standard boundaries between dynamic vs. static systems. At the level of truth-conditions, everything is purely static. But, as we will see, the projection of bounds is calculated by way of recursively specified local contexts, which, while not formulated in the traditional functional or relational architecture of dynamic semantics, is still a characteristically dynamic tool. There are various precise criteria of dynamicness (see van Benthem 1996, Rothschild and Yalcin 2015, 2016), as well as various broader heuristics for what counts as dynamic (cf. Chierchia 1995). An interesting question, which I leave for further work, is where my system falls vis-à-vis those criteria, and more generally whether we should think of it as a dynamic theory or not.

My system builds on a variety of proposals in the literature. Krahmer and Muskens (1995), van den Berg (1996), Rothschild (2017), Heim (2018), Elliott (2020) all propose solutions to the double negation problem which likewise exploit partiality or multidimensionality (see also related discussion in Onea 2013). Rothschild 2017 and Krahmer and Muskens 1995 are particularly direct inspirations for my approach. The idea that witnesses play a central role in the semantics of indefinites has an important precedent in the system of Dekker 1994 and following.[8] Witnesses also play a central role in the class of approaches which use Hilbert's epsilon calculus to analyze indefinites (e.g. Egli and von Heusinger 1995; compare also the arbitrary object approach of Fine 1986).[9] Finally, unpublished work in Schlenker 2011, Chatain 2017 develops systems which exploit statically-conceived local contexts to deal with anaphora, as in my system.

My approach is deeply indebted to all these, but also differs from all of them. I will not go through those differences in detail, nor will I argue that my approach is superior to these. My goal is, rather, to concisely lay out and explore a new system which I think deserves serious consideration.

8 See Dekker 2000, 2002, 2004a,b, 2008, 2012 for developments. The most salient difference between Dekker's system and mine is that Dekker's system is developed with a conjunction which is capable of extending the input sequence, and (thus) with a negation which quantifies over possible extensions of the input sequence; in the resulting system, double negation elimination is logically invalid, as in dynamic systems. Dekker (2001, 2015), however, develops the system so that anaphora resolution is sensitive to information structure, in a manner that could be used to deal with double negation and disjunction.

9 See Avigad and Zach 2020 for a general overview of the epsilon calculus. Superficially, in our system, indefinites and definites are sentence operators, not terms as in the epsilon calculus. But deep differences remain even if we translate our approach into something syntactically more like the epsilon calculus. For instance, we don't validate the characteristic epsilon axiom $\exists xPx \equiv P(\varepsilon xPx)$. Of course, $\exists xPx \equiv P(\text{ɜ}xPx)$ is not well-formed for us anyway, but we could make $\text{ɜ}xPx$ a term by letting $\text{ɜ}xPx = g(x)$ iff $g(x) \in \mathfrak{I}(P,w)$, $\bot$ (an object of which everything is false) if $\mathfrak{I}(P,w) = \varnothing$, and # otherwise. By contrast, in the epsilon calculus, εxPx is an *assignment-insensitive* thing in $\mathfrak{I}(P,w)$, if $\mathfrak{I}(P,w) \neq \varnothing$, and otherwise some arbitrary individual. So, for instance, $P(\varepsilon xPx)$ will be true whenever $\mathfrak{I}(P,w) \neq \varnothing$, while $P(\text{ɜ}xPx)$, on the proposed interpretation, can be undefined, invalidating the characteristic axiom of the epsilon calculus.

5.2 Truth and falsity

I now turn to presenting my theory. On my theory, again, meanings have two-dimensions: a dimension of truth-conditions and a dimension of bounds. I start, in this subsection, by introducing the truth-conditions for my target language.

I work with a standard predicative language closed under the two-place definite operator ιx ('the') and two-place indefinite operator $\mathfrak{z}x$ for any variable x (I reserve $\exists$ for the classical existential quantifier). $\iota x(p,q)$ and $\mathfrak{z}x(p,q)$ are well-formed only if p and q are free in x. $\mathfrak{z}x(p,q)$ stands for ⌜Some p is q⌝. $\iota x(p,q)$ stands for ⌜The p is q⌝. Pronouns are treated as definites with tautological restrictors, so e.g. 'He is p' is translated $\iota x(\top x, p)$. The sentences of the language, as usual, are just those formulae without any variables that are free (in the classical syntactic sense).

The truth-conditional semantic values of our language are those of classical logic, with indefinites getting the truth-conditions of existential quantifiers and definites the truth-conditions of the conjunction of their restrictor and scope. That is, for any (possibly partial) assignment g, world w, and atomic valuation $\mathfrak{I}$, where $g_{[x\to a]}$ is the assignment which takes x to a and otherwise is just like g, we have:

- *Atoms*: $[\![A(x_1,x_2,\ldots x_n)]\!]^{g,w} = 1$ iff $\langle g(x_1), g(x_2), \ldots g(x_n)\rangle \in \mathfrak{I}(A,w)$
- *Conjunction*: $[\![p\&q]\!]^{g,w} = 1$ iff $[\![p]\!]^{g,w}=[\![q]\!]^{g,w}= 1$
- *Disjunction*: $[\![p\vee q]\!]^{g,w} = 1$ iff $[\![p]\!]^{g,w}= 1$ or $[\![q]\!]^{g,w}= 1$
- *Negation*: $[\![\neg p]\!]^{g,w} = 1$ iff $[\![p]\!]^{g,w}= 0$
- *Indefinites*: $[\![\mathfrak{z}x(p,q)]\!]^{g,w} = 1$ iff $\exists a : [\![p\&q]\!]^{g_{[x\to a]},w} = 1$
- *Definites*: $[\![\iota x(p,q)]\!]^{g,w} = [\![p\&q]\!]^{g,w}$

I assume bivalence: if a formula is not true at $\langle g,w\rangle$, it is false there.[10]

10 That means, then, that $A(x_1,x_2,\ldots x_n)$ is false at $\langle g,w\rangle$ if g is undefined at any $x_i : i \leq n$; this is somewhat unusual, but is a natural way to get a bivalent system (which I want) with partial assignments (which I also want).

5.3 The witness bound

The second dimension of meaning in my system is the dimension of bounds. It is in this dimension that I aim to capture the characteristic interaction of indefinites and definites.

Treating bounds as a separate dimension of content from truth-conditions builds on multi-dimensional approaches developed in the literature on presupposition[11] and anaphora (in particular Dekker 2008). Bounds are akin to presupposition, but, as we will see, the pragmatics of bounds are different from the standard pragmatics of presupposition: in contrast to the standard treatment of presupposition, we do not require that bounds be required to be satisfied throughout the antecedent context. We might still choose to identify them as presuppositions (and revise the standard pragmatics of presupposition, as I have argued we need to do anyway (Mandelkern 2016)), or as something else (for instance, Brasoveanu (2013)'s post-suppositions), or as a sui generis kind of content; while I am hesitant to introduce a new term, I do so because I want to remain neutral for the present on the question of how to think about the relation between bounds and other kinds of not-at-issue meaning. While other implementations of the basic idea are possible, the two-dimensional approach provides a simple and illuminating version which brings out particularly clearly the logical points which are central to my interests here.[12]

The role of bounds will become clearer as we proceed; but let me note that bounds never affect truth or falsity, which is crucial in preserving a classical logic for my system.

I propose that indefinites in particular have a *witness bound* which says that if their scope is true relative to *any* assignment, then their scope is true (and has its bounds satisfied) relative to the starting assignment. For brevity, I will use 'satt' for 'has its bounds satisfied'; so, where c is a context (on which more in a moment), g an assignment, and w a world, the witness bound says:[13]

11 See Herzberger 1973, who credits Buridan, and Karttunen and Peters 1979.

12 An alternative would be to develop this system in a one-dimensional but trivalent, rather than two-dimensional but bivalent setting. I discuss this route in Mandelkern fc. Spector (2021) points out a serious problem for that implementation, and proposes a different trivalent system with plural assignments. Detailed comparisons are required but are left for future work.

13 Many thanks to Keny Chatain for suggesting this formulation of the witness bound, based on a much more tortuous earlier version.

- *Witness bound*:
  $\text{з}x(p,q)$ is satt at $\langle c,g,w\rangle$ only if, if $\exists a : [\![p\&q]\!]^{g[x\to a],w} = 1$, then $p\&q$ is true and satt at $\langle c,g,w\rangle$.

We can equivalently reformulate the witness bound in terms of the existential quantifier:

- *Witness bound*:
  $\text{з}x(p,q)$ is satt at $\langle c,g,w\rangle$ only if, if $[\![\exists x(p\&q)]\!]^{g,w} = 1$, then $p\&q$ is true and satt at $\langle c,g,w\rangle$.

So, for instance, an indefinite with the form $\text{з}x(cat(x), sleeping(x))$ is *true* at $\langle g,w\rangle$ just in case there is some sleeping cat; and it is satt only if, if there is some sleeping cat in w, then $g(x)$ is a sleeping cat. So the witness bound coordinates the variable assignment with a witness whenever there is one. This is our way of ensuring that indefinites 'open up a file' on x: by ensuring that x is assigned to a witness of the corresponding existential quantifier.

5.4 The familiarity bound

Definites have a corresponding bound that their scope is 'familiar', as in Heim's system. In other words, in Heim's metaphor, while indefinites *open* files, definites presuppose that a file has *already* been opened, and that whatever information is in their scope is already contained in that file.

To implement this idea, we exploit a context parameter. Contexts will be modeled like Heimian contexts, as sets of pairs of (possibly partial) variable assignments and worlds. As in the Heimian model, worlds represent the context set, and assignments are used to track discourse referents. In the Heimian model, the notion of a null context was a context which pairs each world with the empty assignment; assignments are then gradually extended as conversation proceeds. By contrast, in our system, a null context is a context which pairs each world with every possible assignment; pairs are then eliminated as conversation proceeds.

The role of the context parameter is to provide the domain throughout which we require that the definite's scope is true (and satt). That is, the definite's familiarity bound runs as follows:

- *Familiarity*:
  $\iota x(p,q)$ is satt at $\langle c,g,w\rangle$ only if $\forall\langle g',w'\rangle \in c : p$ is true and satt at $\langle c,g',w'\rangle$.

So a definite description like 'The$_x$ sleeping cat(x) is tabby(x)' is only satt in a context c if, at every point $\langle g,w\rangle \in c$, $g(x)$ is a sleeping cat in w. Pronouns, again, can be treated as definites with tautological restrictors, so a sentence like 'She$_x$ is tabby', parsed $\iota x(\top x, tabby(x))$, is satt in c only if, at every point $\langle g,w\rangle \in c$, $g(x)$ is defined. So pronouns require that their variable be familiar, in the sense of being defined throughout their context; definite descriptions require that the variable be familiar in this sense and also that their restrictor be true at every point in the context.[14]

5.5 Updating

Although there is more to do in laying out the system, let me pause here to explain roughly how things will fit together over the course of updates.

First, let me explain the update rule I will assume. In Heimian dynamic semantics, the update rule is simply functional application: c updated with p is $c[p]$. In my system, updating is instead much closer to the standard eliminative update rule. In particular, updating a context c with a sentence p results in a subsequent context which comprises exactly the points from c where p is true and has its bounds satisfied relative to c: that is, the points $\{\langle g,w\rangle \in c : p$ is true and satt at $\langle c,g,w\rangle\}$.

While this is structurally similar to the standard rule for assertion in semantically static systems, from Stalnaker 1978, this rule incorporates bounds differently from the way that presuppositions are standardly incorporated into that rule. The standard account, from Stalnaker 1974 and incorporated into dynamic semantics in Heim 1983, says that, if we are update c with p, then p has to have its presuppositions satisfied

14 Plausibly, there is also a uniqueness-based definite. Schwarz (2009) has influentially argued that these are two different lexical items. A uniqueness-based definite can of course straightforwardly be incorporated into our system. However, I will continue to focus on the familiarity-definite throughout.

throughout c; if its presuppositions fail to be satisfied at any point in c, then the update will fail to be defined. If we generalized this Stalnakerian assumption to bounds, then updating with indefinites would lead to constant crashes. Consider a null context, and suppose we update it with $\exists x(Fx, Gx)$. The desired result is the set of points $\langle g, w \rangle$ such that $g(x)$ is in $F_w \cap G_w$ (F_w is the extension of F at w). The Stalnakerian assumption would lead to a crash, however, since the null context contains every kind of point, including ones where $F_w \cap G_w$ is non-empty but $g(x)$ is not in $F_w \cap G_w$. This is why I assume that when we update c with p, we simply keep all the points from c where p is true and satt.

This raises the important question of how to integrate theories of semantic presupposition into my framework. I will not take up this question here, partly because I don't know the answer, and partly because there is no need to take a stand on this as far as my central goals here are concerned.[15]

With this in hand, suppose we update a context c with the indefinite $\exists x(\mathit{cat}(x), \mathit{exists}(x))$. Consider an arbitrary point $\langle g, w \rangle \in c$. Recall that $\exists$ has the *truth* conditions of the existential quantifier, so $\exists x(\mathit{cat}(x), \mathit{exists}(x))$ is *true* at $\langle g, w \rangle$ iff there is a cat in w. Suppose first that w has no cats; then our sentence is false at $\langle g, w \rangle$ and so we eliminate this point. Suppose next that there is a cat in w. Then our sentence is true at $\langle g, w \rangle$. But this isn't enough for this point to survive; we must also check whether our sentence's witness bound is satisfied. That bound says that, if w has a cat, then $g(x)$ is a cat in w. Since w does have a cat, by hypothesis, the witness bound is thus satisfied iff $g(x)$ is a cat in w. So, $\langle g, w \rangle$ survives update with $\exists x(\mathit{cat}(x), \mathit{exists}(x))$ iff $g(x)$ is a cat in w.

This, in turn, is the key to the subsequent licensing of anaphora. Let c' be the context that results from updating c with $\exists x(\mathit{cat}(x), \mathit{exists}(x))$. Note that, in c', not only does every world contain cats, but also every variable assignment assigns x in particular to something that is a cat in its paired world. Suppose that someone now says 'The cat is tabby', $\iota x(\mathit{cat}(x), \mathit{tabby}(x))$. Consider an arbitrary point $\langle g, w \rangle$ in c'. $\iota x(\mathit{cat}(x), \mathit{tabby}(x))$ has its familiarity bound satisfied at $\langle g, w \rangle$ iff, for every point $\langle g', w' \rangle \in c'$, $g'(x)$ is a cat in w'. But this is guaranteed to hold because of our prior

15 A reviewer for this journal suggests an interesting idea: perhaps the familiarity requirement is a presupposition, not a bound, and presuppositions just are those bounds that must be satisfied throughout a context, rather than at a given point in a context.

update with the corresponding indefinite. $\iota x(cat(x), tabby(x))$ is true at $\langle g, w\rangle$ iff the corresponding open formula $cat(x)\&tabby(x)$ is true, iff $g(x)$ is a tabby cat in w.

Putting these two updates together, a point $\langle g, w\rangle$ in c survives update with $\exists x(cat(x), exists(x))$ and then $\iota x(cat(x), tabby(x))$ just in case $g(x)$ is a tabby cat in w. Things work in essentially the same way for pronouns: updating with 'There is a cat. It is tabby' has exactly the same effect as updating with 'There is a cat. The cat is tabby'. Both take us to a context which comprises exactly those points $\langle g, w\rangle$ from the initial context where $g(x)$ is a tabby cat in w. That's because the familiarity bound of $\iota x(\top x, tabby(x))$ requires that, for all points $\langle g, w\rangle$ in c', $g(x)$ is defined. But this will hold, thanks to the preceding update with the indefinite.

Note something important in the calculation of familiarity bounds: since the familiarity bound quantifies universally over points in the context, it will hold at either all or none of them. This is what accounts for the infelicity that results from asserting a definite without a corresponding indefinite.[16] This is very different from the witness bounds of indefinites, which, crucially, can be satisfied at some points in a context and not at others.

There is a deep similarity between this system and Heim's file card system. In our system, asserting, for instance, $\exists x(Fx, Gx)$ 'opens a file' on x by making x defined at every point in the updated context, and adds to this file the information that x is F and G. Asserting a definite $\iota x(Fx, Hx)$ is then licensed, in the sense that its familiarity bound is satisfied throughout the context, because a file has been opened on x that contains the information that F. The definite adds to the file the further information that x is H. So, abstractly, the overall system is very similar to Heim's; but the mechanics are very different.

5.6 Projection

Now I will turn from cross-sentential anaphora to intra-sentential anaphora, where I will argue my system is a clear improvement over standard dynamic systems.

16 Of course, definites can also be accommodated, as Heim 1982 and others discuss. As Heim discusses, there is a spectrum of difficulty in accommodation from pronouns (hardest) to definite descriptions to possessives (easiest). Heim speculates that this is because of the increasing amount of descriptive material across this spectrum, which serves as an aide to accommodation.

In order to spell this out, I need to add a final piece to our theory. Since our system contains bounds, I need to say how they project out of complex sentences. The basic idea extends the approach to the corresponding question about presuppositions of Schlenker 2009, 2010, who, building on Karttunen 1973, 1974, Stalnaker 1974, develops a theory of local contexts to account for presupposition projection in a broadly static framework. So, for instance, the local context for a right conjunct will be the set of points from the global context where the left conjunct is both true and satt (relative to the global context); a conjunction is satt, at a point $\langle g, w \rangle$ and a global context c, iff both conjuncts are satt at $\langle g, w \rangle$ relative to their local contexts. In the appendix, I spell out this generalization recursively. The base case is atomic sentences: $A(x_1, x_2, \ldots x_n)$ is satt at $\langle g, w \rangle$ iff g is defined on all of $x_1, x_2, \ldots x_n$.

There are two questions about projection that I want to flag here. The first concerns order. Traditional dynamic systems are asymmetric: the local context for a left 'junct is the global context, while the local context for the right 'junct is the global context updated in a suitable way with the content of the left junct. This is motivated by contrasts like those in (14):

(14) a. A man came in and $\left\{ \begin{smallmatrix} \text{he} \\ \text{the man} \end{smallmatrix} \right\}$ sat down.
b. $\left\{ \begin{smallmatrix} \text{He} \\ \text{The man} \end{smallmatrix} \right\}$ sat down and a man came in.

It is very difficult to interpret the 'he' as anaphoric on 'a man' in (14-b).

So order seems to matter for anaphora. But the data are very subtle. While there is an indisputable contrast in (14), in the case of bathroom disjunctions, both orders are clearly interpretable, even if there is a preference for the standard order:

(15) a. Either there isn't a bathroom, or $\left\{ \begin{smallmatrix} \text{it} \\ \text{the bathroom} \end{smallmatrix} \right\}$ is upsatirs
b. Either $\left\{ \begin{smallmatrix} \text{it} \\ \text{the bathroom} \end{smallmatrix} \right\}$ is upstairs, or there isn't a bathroom.

To see that the indefinite is doing work here, we can look at our running minimal pair:

(16) a. Either $\left\{ \begin{smallmatrix} \text{she} \\ \text{the child} \end{smallmatrix} \right\}$ is at boarding school, or Susie doesn't have a child.
b. Either $\left\{ \begin{smallmatrix} \text{she} \\ \text{the child} \end{smallmatrix} \right\}$ is at boarding school, or Susie isn't a parent.

The negated indefinite in the right disjunct clearly facilitates an anaphoric reading in (16-a) which is much less available in (16-b).

Moreover, as many have noted, right-to-left anaphora is possible across conjunction:

(17)
a. It took him a long time, but a student of mine wrote a really good paper.
b. I told him off after a friend of mine told me he had just adopted eight cats.
c. He complained all the way to the fair, and then one of my kids just disappeared.
d. Everyone who received it late gave a student's paper an F.

So it's unclear what to make of order asymmetries.

In fact, the question of order is modular in my system, in the sense that it is straightforward to implement both a standard, asymmetric version of my system, and a symmetric version, depending how we define local contexts. So, for instance, we could either say $p\&q$ is satt at $\langle c,g,w\rangle$ iff p is satt at $\langle c^q,g,w\rangle$ and q is satt at $\langle c^p,g,w\rangle$, for a symmetric entry; or we can say that $p\&q$ is satt at $\langle c,g,w\rangle$ iff p is satt at $\langle c,g,w\rangle$ and q is satt at $\langle c^p,g,w\rangle$, for an asymmetric entry. We could also, following Schlenker, posit that both systems are present, with a preference for the asymmetric one. I do not know what to make of the question of order. For reasons having to do with modality, I am sympathetic to a symmetric approach to local contexts in general (see **?**), and so I give a symmetric system in the conclusion; but, again, it is easy to transform that into an asymmetric system instead.

The second question concerns how bounds should project out of (in)definites. Consider first indefinites. Here is my proposal about how indefinites project the bounds of their complements:

- *Indefinites*: $\mathrm{E}x(p,q)$ is satt at $\langle c,g,w\rangle$ only if $\exists g'$: $p\&q$ is satt at $\langle c,g',w\rangle$.

The motivation is this. We don't want to require that p itself is satt at $\langle c,g,w\rangle$, since then updating with, say, $\neg\mathrm{E}x(Fx,Gx)$ would end up making x familiar (since $Fx\&Gx$ is satt at g only if $g(x)\neq\#$); but negated indefinites do not generally license subsequent anaphora (*modulo* the possible, but dispreferred, wide-scope interpretation, and some further points about modal subordination which I address in the conclusion):

(18) We don't have a cat. # She is a tabby.

On the other hand, we don't want $\exists$ to have trivial projection conditions, since that would mean that negated indefinites could contain non-familiar definites in their scope, which is also not possible: (19), for instance, is only possible when 'the cat' is familiar (or uniquely salient):

(19) We don't have the cat.

A middle ground is to require that the indefinite's prejacent is satt relative to *some* assignment, holding fixed the local context. That means that a sentence like $\neg\exists x(Fx, Gx)$ will be satt even if $g(x) = \#$; but a negated indefinite taking scope over a novel definite will not be satt, since varying the assignment parameter will not generally change whether a definite is satt.

With this in hand, we can look at the update effect of a negated indefinite, like 'There is not a cat', $\neg\exists x(cat(x), exists(x))$. Consider any point $\langle g, w\rangle \in c$. $\neg\exists x(cat(x), exists(x))$ is true at $\langle g, w\rangle$ iff its negatum is false; since the negatum has the truth-conditions of the existential quantifier, that holds iff there is no cat in w. Suppose that there is no cat at w. $\neg\exists x(cat(x), exists(x))$ is satt iff (i) its scope is satt at some g—which it is; and (ii) its witness bound (which projects through negation) is satt, which holds just in case, if there is a cat in w, then $g(x)$ is a cat in w. But, by assumption, there's no cat in w; so the witness bound is (trivially) satisfied. So $\langle g, w\rangle$ survives update with $\neg\exists x(cat(x), exists(x))$ just in case w has no cats.

In general, whenever an indefinite sentence is false, its witness bound is trivially satisfied, since the witness bound is a material conditional whose antecedent asserts the truth of the indefinite. So *negated indefinites* are always satt, meaning that negated indefinites have exactly the same update effect as the corresponding negated existential quantifiers (*modulo* any familiarity bounds of their scope). This means that my system captures the strong, universal meaning of negated indefinites, as well as their inability to license subsequent co-indexed definites.

Consider next the question of how definites pass on the bounds of their prejacent. I propose that they do so in a conditionalized fashion as follows:

- *Definites*: $\iota x(p,q)$ is satt at $\langle c,g,w\rangle$ only if, if p is true and satt at $\langle c,g,w\rangle$, then q is satt at $\langle c,g,w\rangle$.

We don't want $\iota x(p,q)$ to require that q be satt unconditionally. Consider a bathroom sentence with the form, say, $\neg\exists x(Fx,Gx) \vee \iota x(\top x,Hx)$ (more on these in a moment). We don't want a bathroom sentence like this to make x familiar, since updating with something like 'Either there isn't a bathroom, or it is upstairs' doesn't license subsequent anaphora (#'It has horrible wall paper'). On the other hand, we can't simply ignore the bounds of q, since any witness bounds of indefinites in q should be satisfied once we update with $\iota x(p,q)$ ('The bathroom has a sink' licenses subsequent anaphora to the sink). So I propose to require that q be satt whenever p is both satt and true—roughly, q has to be satt provided we're in a case where we are keeping a point because of how it values p.[17]

5.7 Open scope of indefinites

With this in hand, we can look at some of the characteristic features of intrasentential anaphora in my system.

Recall that one standard way to formulate the key generalization about indefinites which motivates dynamic semantics is to say that indefinites have open scope to their right: ⌜Some F is (G and H)⌝, for instance, is in some sense equivalent to ⌜Some F is G, and $\left\{\begin{smallmatrix}\text{it}\\ \text{the } F\end{smallmatrix}\right\}$ is H⌝. Thus, for instance, we saw above that 'Sue has a child, and she is at boarding school' is intuitively equivalent to 'Sue has a child at boarding school'; while 'Sue is a parent, and she is at boarding school' is intuitively equivalent to 'Sue is a parent at boarding school'.

In our language, the claim that indefinites have open scope to their right can be formulated as the claim that the three variants in (20) are equivalent:[18]

17 One corollary is that a disjunct can contain locally novel definites without that making the sentence unsatt. Hence consider $\neg\exists x(Fx,Gx) \vee \iota x(\top x,\iota y(\top y,R(x,y)))$. This will be satt at $\langle c,g,w\rangle$ when $g(x)=\#$ and nothing is F and G in w, even if y is novel in c, since the restrictor of ιx (that is, $\top x$) is not true and satt here. Still, if y is novel, then this sentence will be equivalent to $\neg\exists x(Fx,Gx)$ whenever satt, which will suffice to account for its infelicity via standard theories of redundancy.

18 I focus on conjunction here, but the same points go for sequences, and for any substitution instances of Fx, Gx, and Hx.

(20) a. $\exists x(Fx, Gx \& Hx)$ *Some F is G and H*

b. $\exists x(Fx, Gx) \& \iota x(Fx, Hx)$ *Some F is G and the F is H*

c. $\exists x(Fx, Gx) \& \iota x(\top x, Hx)$ *Some F is G, and it is H*

Of course, (20-a)–(20-c) are not *logically* equivalent in my system, where p and q are logically equivalent whenever, for any index i in any intended model, p is true at i iff q is true at i. For instance, consider a point $\langle g, w \rangle$ where $F_w \cap G_w \cap H_w$ is non-empty, and $g(x) \notin H_w$. Then (20-a) will be true at $\langle g, w \rangle$, since (20-a) is truth-conditionally equivalent to the corresponding existentially quantified sentence, which is true. By contrast, (20-b) and (20-c) will both be false at $\langle g, w \rangle$, since Hx is false there; and $\iota x(Fx, Hx)$ and $\iota x(\top x, Hx)$ are both true only if Hx is true.

But you will have noticed that, while (20-a) is is true at $\langle g, w \rangle$, its witness bound is not satisfied—there is something in $F_w \cap G_w \cap H_w$, but $g(x)$ is not in $F_w \cap G_w \cap H_w$. This points the way towards the sense in which (20-a)–(20-c) are equivalent: they have the same truth-value *provided their bounds are satisfied*. To characterize this kind of equivalence, we can generalize von Fintel (1999)'s notion of Strawson-entailment (which he applies to presupposition fragments) to bounds as follows:[19] *p bound-entails q* iff, for any index (context/assignment/world-triple) i in any intended model, if p and q are both satt at i and p is true at i, then q is true at i (I'll write $p \models q$ for logical entailment, and $p \models_b q$ for bound-entailment). p and q are bound-equivalent iff each bound-entails the other. (20-a)–(20-c) are not pairwise logically equivalent, but they are pairwise bound-equivalent.

The reasoning behind this is simple. For any point $i = \langle c, g, w \rangle$, suppose (20-a) is satt and is true at i. As we have seen, this holds just in case $g(i) \in F_w \cap G_w \cap H_w$. Whenever (20-b) or (20-c) are satt, they are also true just in case $g(i) \in F_w \cap G_w \cap H_w$. Focusing on (20-b): given its witness bound, the left conjunct of (20-b) is true at $\langle g, w \rangle$ iff $g(x) \in F_w \cap G_w$. The right conjunct is true iff $Fx \wedge Hx$ is true, that is, iff $g(x)$ is also in $F_w \cap H_w$; together, this is equivalent to the requirement that $g(x) \in F_w \cap G_w \cap H_w$. The reasoning for (20-c) is similar. All three sentences are both satt and true iff $g(x) \in F_w \cap G_w \cap H_w$.

19 Thanks to a reviewer for suggesting this helpful terminology.

In fact, there is another sense in which (20-a)–(20-c) are equivalent: it is easy to see that, if any one of (20-a)–(20-c) is satt and true, then all three are satt and true.

Since it is the combination of bounds plus truth-conditions that matters for pragmatics in the system I have developed, this account of the open scope of indefinites thus carries appropriate pragmatic weight. It is crucial, however, that I do *not* predict these to be logically equivalent. By making these bound-equivalent but not logically-equivalent, we can validate the open scope of indefinites (in a pragmatically relevant sense) while still retaining a classical logic.

5.8 Classicality

Recall the two closely related problems for dynamic semantics discussed in §4: in dynamic semantics, $\neg\neg p$ and p are not always equivalent; nor are $\neg p \vee q$ and $\neg p \vee (p\&q)$. By contrast, since our connectives are classical, the logic of our system is the logic of classical predicate logic (under the obvious translation schema between the two slightly different languages). And so it follows from this more general fact that the two classical rules $\neg\neg p \dashv\models p$ and $\neg p \vee q \dashv\models \neg p \vee (p\&q)$ are valid in my system. Moreover, the bounded logic of any system is always a superset of the system's logic: that is, if $p \models q$, then $p \models_b q$. This is for the obvious reason that, if $p \models q$, then q is true at any point in any model where p is, and thus a fortiori q is true at any point where p and q are satt and p is true. So we also have $\neg\neg p \dashv\models_b p$, and likewise $\neg p \vee q \dashv\models_b \neg p \vee (p\&q)$.

The basic reasoning behind all this, again, is very simple: our connectives are, at the level of truth and falsity, just the classical connectives.

More concretely, that means that the bounded system avoids the empirical problems that these logical failures lead to in dynamic semantics. First, doubly negated indefinites will thus license subsequent definites, as desired. Consider 'It's not the case that Susie doesn't have a child', $\neg\neg(\exists x(\mathit{child}(x), \mathit{Susie's}(x)))$. This will be semantically equivalent to $\exists x(\mathit{child}(x), \mathit{Susie's}(x))$, since our truth-conditions for negation are classical and since bounds project out of negation; and thus this will license subsequent definites like '$\left\{\begin{smallmatrix}\text{She}\\ \text{The child}\end{smallmatrix}\right\}$ is at boarding school', $\iota x(\mathit{child}(x), \mathit{at\text{-}boarding\text{-}school}(x))$.

Similar reasoning applies to disjunctions like ‘Either Susie doesn’t have a child, or $\left\{ {\text{she} \atop \text{the child}} \right\}$ is at boarding school’ $(\neg \exists x(\mathit{child}(x), \mathit{Susie's}(x))) \vee \iota x(\mathit{child}(x), \mathit{at\text{-}boarding\text{-}school}(x))$. The local context for the right disjunct will only include points where the negation of the left disjunct is true and satt, which holds at a point iff the indefinite $\neg\neg\exists x(\mathit{child}(x), \mathit{Susie's}(x))$ is true and satt there, which holds, thanks to the validity of double negation elimination, iff $\exists x(\mathit{child}(x), \mathit{Susie's}(x))$ is true and satt there. Thus the local context for the right disjunct will only contain pairs $\langle g, w \rangle$ where $g(x)$ is a child of Susie’s in w. That means that the familiarity bound of the definite description (or a pronoun) will be satisfied. The whole sentence will thus be true and satt at $\langle c, g, w \rangle$ iff either Susie is childless in w; or (i) $g(x)$ is Susie’s child in w (this follows from the indefinite’s witness bound, which projects to the whole sentence) and (ii) $g(x)$ is at boarding school in w.

Our system thus avoids the problem that negation and disjunction pose for dynamic systems. Importantly, this is not because of a local fix, but rather follows from the much more general fact that the underlying architecture of the bounded system is fully classical. Because indefinites are still, at the level of truth-conditions, existential quantifiers, we do not need to make negation a quantifier over assignments. Instead, a classical treatment of negation—and of the other connectives—works perfectly well, which means not just that the resulting system has a certain elegant simplicity to it, but also that it avoids the serious problems that dynamic semantics runs into because of its non-classical treatment of negation.

6 Quantifiers

This concludes the exposition of my basic system. Before concluding, I want to briefly discuss how to integrate generalized quantifiers like ‘every’ and ‘most’. This is important given how central a role donkey sentences have played in the literature on anaphora, though the extension is fairly straightforward, and my discussion of this complicated area is necessarily brief.

Recall the core data we are trying to capture in the interaction between quantification and anaphora, namely the co-variation between indefinites like ‘a child’ and definites like '$\left\{\substack{\text{them}\\\text{the child}}\right\}$' in sentences like (21-a), and the unavailability of a co-varying reading in a minimal variant like (21-b):

(21) a. Everyone who has a child loves $\left\{\substack{\text{them}\\\text{the child}}\right\}$.
b. Every parent loves $\left\{\substack{\text{them}\\\text{the child}}\right\}$.

I assume generalized quantifiers take three arguments: an unpronounced domain variable δ, a restrictor, and a scope. Instead of treating the domain variable as corresponding to a set of individuals, I treat it as corresponding to a non-empty set of pairs of individuals and variable assignments. Then we proceed in the natural way: for instance, 'every' and 'most' get the following truth-conditions:

- $[\![\text{EVERY}x_\delta(p,q)]\!]^{g,w}=1$ iff $\forall\langle a,g'\rangle \in g(\delta) : [\![p]\!]^{g'_{[x\to a]},w} = 1 \to [\![p\&q]\!]^{g'_{[x\to a]},w} = 1$
- $[\![\text{MOST}x_\delta(p,q)]\!]^{g,w}=1$ iff for most $\langle a,g'\rangle \in g(\delta)$ s.t. $[\![p]\!]^{g'_{[x\to a]},w} = 1, [\![p\&q]\!]^{g'_{[x\to a]},w} = 1$

I also assume that quantifiers have bounds, specifically about the domain parameter. The basic idea is that the domain parameter should not contain duplicates (each individual should be in a pair in the domain only once); it should contain only pairs that make the restrictor and scope satt; and, finally, it should only contain assignments which agree with the starting assignment on any variables that are familiar in c. So $Qx_\delta(p,q)$ is satt at $\langle c,g,w\rangle$ iff:

- $\forall a : \exists! \langle a',g'\rangle \in g(\delta) : a' = a$. In other words, each individual a in the domain is included in exactly one pair in $g(\delta)$ (this is crucial for ensuring both that we quantify over every individual in the domain, and that we avoid the 'proportion problem' which arises for some versions of dynamic semantics).
- $\forall\langle a,g'\rangle \in g(\delta) : p\&q$ is satt at $\left\langle c, g'_{[x\to a]}, w\right\rangle$. This ensures that (i) definites in p and q are satt (if they weren't, then $g(\delta)$ would be empty, contrary to assumption); and (ii) indefinites have their witness bounds satisfied relative to the variable assignments in $g(\delta)$.
- $\forall\langle a,g'\rangle \in g(\delta) : g' \sim_c g$, where $g' \sim_c g$ iff g' agrees with g on the values of all variables which are familiar in c—that is, which are defined at every point in

c. This ensures that all the assignments in the domain agree with our starting assignment on any variables which are already familiar in c.

For sentences without (in)definites, the variable assignments in $g(\delta)$ don't do any interesting work. So, e.g., $\text{EVERY}x_\delta(\mathit{farmer}(x), \mathit{tall}(x))$ is true and satt at $\langle g, w\rangle$ just in case every individual in the domain who is a farmer in w is tall in w.

The more interesting case is that of a donkey sentence like (21-a), repeated here, with the parse in (22-b):

(22) a. Everyone who has a child loves it.

b. $\text{EVERY}x_\delta(\underbrace{\mathrm{ɜ}y(\mathit{child}(y), \mathit{of}(y,x))}_{p}, \underbrace{\iota y(\top y, \mathit{loves}(x,y))}_{q})$

Suppose (22-b) is satt in $\langle c, g, w\rangle$. (22-b) is true in $\langle c, g, w\rangle$ iff for every pair $\langle a, g'\rangle \in g(\delta)$, if p is true at $\left\langle c, g'_{[x\to a]}, w\right\rangle$, then so is $p\&q$. Consider an arbitrary pair $\langle a, g'\rangle \in g(\delta)$. Given that p is satt at $\left\langle c, g'_{[x\to a]}, w\right\rangle$, it is true at $\left\langle c, g'_{[x\to a]}, w\right\rangle$ iff $g'(y)$ is a child of a, false iff a is childless. If false, then $\langle a, g'\rangle$ doesn't count against the truth of (22-b). If true, then $p\&q$ must also be true at $\left\langle c, g'_{[x\to a]}, w\right\rangle$ in order for (22-b) to be. Given that $p\&q$ is satt at $\left\langle c, g'_{[x\to a]}, w\right\rangle$, it is true there iff a loves their child $g'(y)$. So, (22-b) is true and satt at $\langle c, g, w\rangle$ iff, for every a in the domain, if a has a child in w, then a loves their paired child in w. We thus derive a covarying interpretation of donkey sentences.[20]

By contrast, a co-varying reading will not be available for a sentence like (21-b):

(21-b) Every parent$_x$ loves $\left\{\begin{matrix}\text{it}\\ \text{the child}\end{matrix}\right\}_y$.

This is simply because there is no indefinite corresponding to the definite in (21-b), so the definite won't have its familiarity bound satisfied unless it has already been

20 The interpretation is slightly non-standard. Standardly two "readings" of donkey sentences are distinguished: one which says that everyone who has a child loves some child of theirs; the other of which says that everyone who has a child loves every child they have. It is famously difficult to distinguish these two readings, and there is controversy about whether these are really two readings or two pragmatic interpretations see e.g. Heim 1982, Root 1986, Rooth 1987, Schubert and Pelletier 1989, Chierchia 1992, Kanazawa 1994, Chierchia 1995, Champollion et al. 2019). This is a complicated issue, which I don't have anything in particular to add to here, but it is interesting to note that our truth-conditions cross-cut these two readings in an interesting, assignment-sensitive way.

introduced in the global context. In that case, (21-b) will be interpreted as saying that some particular child is loved by every parent.

There is, again, much more to explore, but this brief discussion shows the basic contours of how the bounded approach can make sense of donkey sentences.

7 Conclusion

There is a difference between indefinites like ‘has a child’ and ‘is a parent’. This poses a challenge to the classical analysis of indefinites as existential quantifiers. Both dynamic semantics (which I have focused on here) and e-type theories captures this difference by rejecting, in different ways, classical notions of meaning and corresponding classical treatments of connectives.

The bounded system I have presented here captures the contrast between pairs like this, and, more generally, the coordination between indefinites and corresponding definites, with very different tools from existing theories. I have argued that we should separate the two characteristic contributions of indefinites into two dimensions of contents: their existential import is captured by their quantificational truth-conditions; their ability to license subsequent anaphora (to open a file, to start a new discourse referent) is captured by their witness bound. This system avoids the specific problems for dynamic systems involving negation and disjunction explored above. But it also, more importantly, shows that we can pull apart many of the insights of dynamic semantics from its revisionary approach to content and connectives. This classicality of the bounded system distinguishes it from dynamic semantics, which departs in deep ways from classical logic, invalidating not just double negation elimination—our focus here—but also (in many systems) classical laws like non-contradiction and excluded middle (see van Benthem 1996, Mandelkern 2020). In this sense, my system is very conservative. All the *dynamic* action in the system comes via bounds; it is (only) in the dimension of bounds that the logic extends classical logic, in particular predicting that indefinites have open scope to their right as a matter of bounded (but not logical) validity.

There is obviously much more work to do in exploring the bounded system. We should look, for instance, at extensions of the system to other domains, like modality, attitude reports, quantificational subordination, adverbial quantifiers, conditionals, and plural anaphora. I take up these topics in Mandelkern 2021b. A particularly important extension, given the main focus of this paper, is to modal subordination, which interacts with negation in complex ways (Roberts 1987, Kibble 1994, Frank 1996, Brasoveanu 2007, 2010, Hofmann 2019), as in (23), due to an anonymous referee:

(23) Sue doesn't have a child. You would know them by now.

Cases like this pose a *prima facie* challenge to the standard generalization that singly negated indefinites do not license subsequent definites (though not, obviously, to the more limited generalization that singly negated indefinites do not license subsequent *matrix* definites).

While I don't have space for a full discussion, I want to note that I think a natural extension of our approach can deal with cases like this, by combining our system with ideas recently developed (in a standard dynamic framework) by Hofmann (2019). The idea is to have the value of a variable, relative to an assignment, be a partial individual concept rather than an individual (following Stone 1999, Aloni 2000). Then we can generalize the witness bound across worlds: we say that $\mathrm{ɜ}x(p,q)$ is satt, at $\langle g,w\rangle$, only if, *for any world* w', if $\exists x(p\&q)$ is true at $\langle g,w'\rangle$, then $p\&q$ is true at $\langle g,w'\rangle$. That means that the witness bound is not inert in the case of negated indefinites: $\neg\mathrm{ɜ}x(p,q)$ will license subsequent anaphora to *x just in case we are within the scope of an operator whose local context entails* $\exists x(p\&q)$. That means that singly negated indefinites will still not license subsequent matrix definites, but they will license definites in a sequence like (23), given any theory of modal subordination on which the local context for the indefinite under 'would' entails that Sue has a child.

This is a conservative extension of the system sketched above, in the sense that all the points I have made so far also apply to this extension. This is no more than a sketch, which, again, I develop in Mandelkern 2021b; but I think this suggests a promising route for dealing with the interaction of anaphora with intensional operators.[21]

21 A harder case comes from what Kibble (1994), citing Paul Dekker, calls *negative subordination*—cases like (24):

(24) John doesn't have a car so he doesn't wash $\left\{\begin{smallmatrix}\text{it}\\ \text{the car}\end{smallmatrix}\right\}$.

We should also explore questions of order. As we saw above, the empirical situation is complicated, and local contexts can just as easily be specified in a symmetric or asymmetric fashion. This means that we have more flexibility than standard dynamic systems in accounting for order symmetries, but how we should use that flexibility remains a mystery to me.

We should compare the bounded systems in more detail to other theories of anaphora, especially those which, like mine, avoid the double-negation and disjunction problems in different ways. In addition to those mentioned as precedents above, we might look at type-theoretic approaches like Bekki (2013) or Gotham (2019)'s intuitionistically inspired treatment; as well as trivalent systems (something I explore in Mandelkern fc in a trivalent implementation of the bounded system with local contexts; which Spector (2021) takes up in a trivalent system without local contexts, and which Elliott (2020) explores in a trivalent dynamic system).

We should explore further foundational questions about the bounded system: questions about its internal foundations (like those helpfully posed in Lewis 2012, Chatain 2017, and work in progress by Keny Chatain); and questions about the relationship between anaphora, presupposition, and modality, where bounded theory contributes to a developing research program which aims to capture the insights of dynamic semantics in systems that are more conservative, foundationally and logically (Schlenker 2008, 2009, Dorr and Hawthorne 2013). Indeed, the system I have presented here is architecturally very similar to the theory of modals I argue for in Mandelkern 2019 and the theory of conditionals I develop in Mandelkern 2021a; in Mandelkern 2021b I bring these strands together, arguing that there are common arguments across these domains for the virtues of separating the dynamic effects of modals, conditionals, and anaphora into a local-context-sensitive strand distinct from a classical truth-conditional core.

Let me close with a high-level comment on the structure of bounded theory. In an illuminating discussion, Cumming (2015) identifies what he calls *the dilemma*

I am not sure what to make of cases like this, which are not straightforwardly accounted for in our account, even with the cross-world witness bound. One possibility is to appeal to something like Heim (1982, 1983)'s notion of local accommodation, or, relatedly, to analyze this as a kind of meta-linguistic negation. Obviously this needs further exploration; while I don't think my view shines any special light on this issue, I think it is amenable to whatever general solution one prefers. See Lewis (2020) for interesting further problems involving negation.

of indefinites. On the one hand, they seem to have existential import: whether an indefinite sentence is true or false apparently depends just on the truth or falsity of the corresponding existential quantifier. Intuitively 'Sue has a child' is true just in case Sue is a parent, false otherwise, whether or not the speaker has a particular child in mind. If Sue is a parent, but John thinks she is the parent of Latif when in fact she is the parent of Arden, then 'Sue has a child' is as true when John says it with Latif in mind, as when I say it with Arden in mind. On the other hand, indefinites license subsequent anaphora in ways not predicted by a purely existential account: 'Sue is a parent' and 'Sue has a child' seem inequivalent when we look at how they contribute to environments like sequences of sentences, conjunctions, or quantifiers. Crudely speaking, the two main approaches to indefinites in the literature aim to generalize to one of these two faces. On e-type approaches, indefinites are, after all, just existential quantifiers; their ability to license subsequent anaphora is explained by appeal to pragmatic and/or syntactic reconstruction that they make available. On dynamic approaches, by contrast, indefinites are fundamentally variables; their existential import is explained by appeal to more complicated notions of context and truth, and a quantificational treatment of negation.

The bounded theory suggests a synthesis: both faces of indefinites are present, but in different dimensions of content. At the level of truth-conditions, indefinites are existential quantifiers, accounting for the existential import of indefinites and explaining the validity of classical inference patterns. But at the level of bounds, indefinites do more: they require the presence of a witness to their truth, a witness that enables subsequent coreference with definites. These bounds help us keep track of anaphoric relations, and thus follow the twists and turns of conversation.

A Semantics

For ease of reference, I summarize the semantics given in the text. We recursively specify the truth-conditions for our language, relative to pairs of a world and a variable assignment: where g is a variable assignment, w a world, and sentence p, $[\![p]\!]^{g,w}$ is the bivalent truth-value of p at $\langle g, w\rangle$. We also specify the bounds of sentences relative to a context (set of assignment-world pairs), an assignment, and a world. Bounds never influence truth-conditional content, so contexts play no role in determining the latter. For any context c and sentence p, c^p is the set of world-assignment pairs where p is true and satt relative to c. $\mathfrak{I}$ is an atomic valuation taking n-ary atoms and worlds to sets of n-tuples.

- *Atoms*:
  $[\![A(x_1,x_2,\ldots x_n)]\!]^{g,w}=1$ iff $g(x_1),\ldots g(x_n)$ are all defined and $\langle g(x_1),\ldots g(x_n)\rangle \in \mathfrak{I}(A,w)$, 0 otherwise
  $A(x_1,x_2,\ldots x_n)$ is satt at $\langle c,g,w\rangle$ iff $g(x_1),\ldots g(x_n)$ are all defined

- *Conjunction*:
  $[\![p\&q]\!]^{g,w}=1$ iff $[\![p]\!]^{g,w}=[\![q]\!]^{g,w}=1$
  $p\&q$ is satt at $\langle c,g,w\rangle$ iff p is satt at $\langle c^q,g,w\rangle$ and q is satt at $\langle c^p,g,w\rangle$

- *Disjunction*:
  $[\![p\vee q]\!]^{g,w}=1$ iff $[\![p]\!]^{g,w}=1$ or $[\![q]\!]^{g,w}=1$
  $p\vee q$ is satt at $\langle c,g,w\rangle$ iff p is satt at $\langle c^{\neg q},g,w\rangle$ and q is satt at $\langle c^{\neg p},g,w\rangle$

- *Negation:*
  $[\![\neg p]\!]^{g,w}=1$ iff $[\![p]\!]^{g,w}=0$
  $\neg p$ is satt at $\langle c,g,w\rangle$ iff p is satt at $\langle c,g,w\rangle$

- *Indefinites*:
  $[\![\text{ɜ}x(p,q)]\!]^{g,w}=1$ iff $[\![\exists x(p\&q)]\!]^{g,w}=1$ iff $\exists a:[\![p\&q]\!]^{g_{[x\to a]},w}=1$
  $\text{ɜ}x(p,q)$ is satt at $\langle c,g,w\rangle$ iff $\exists g':p\&q$ is satt at $\langle c,g',w\rangle$ and $[\![\exists x(p\&q)]\!]^{g,w}=1\to p\&q$ is true and satt at $\langle c,g,w\rangle$

- *Definites*:
  $[\![\iota x(p,q)]\!]^{g,w}=[\![p\&q]\!]^{g,w}$
  $\iota x(p,q)$ is satt at $\langle c,g,w\rangle$ iff $\forall\langle g',w'\rangle\in c:p$ is true and satt at $\langle c,g',w'\rangle$ and, if p is true and satt at $\langle c,g,w\rangle$, then q is satt at $\langle c,g,w\rangle$

- *Quantifiers*:
  $[\![\text{EVERY}x_\delta(p,q)]\!]^{g,w}=1$ iff $\forall\langle a,g'\rangle\in g(\delta):[\![p]\!]^{g'_{[x\to a]},w}=1\to[\![p\&q]\!]^{g'_{[x\to a]},w}=1$
  $\text{EVERY}x_\delta(p,q)$ is satt at $\langle c,g,w\rangle$ iff

  $$\forall a:\exists!\langle a',g'\rangle\in g(\delta):a'=a;$$
  $$\forall\langle a,g'\rangle\in g(\delta):p\&q \text{ is satt at } \left\langle c,g'_{[x\to a]},w\right\rangle\text{; and}$$
  $$\forall\langle a,g'\rangle\in g(\delta):g'\sim_c g.$$

  $[\![\text{MOST}x_\delta(p,q)]\!]^{g,w}=1$ iff $\frac{|\{\langle a,g'\rangle\in g(\delta):[\![p\&q]\!]^{g'_{[x\to a]},w}=1\}|}{|\{\langle a,g'\rangle\in g(\delta):[\![p]\!]^{g'_{[x\to a]},w}=1\}|}\geq .5$
  $\text{MOST}x_\delta(p,q)$ is satt at $\langle c,g,w\rangle$ iff

  $$\forall a:\exists!\langle a',g'\rangle\in g(\delta):a'=a;$$
  $$\forall\langle a,g'\rangle\in g(\delta):p\&q \text{ is satt at } \left\langle c,g'_{[x\to a]},w\right\rangle\text{; and}$$
  $$\forall\langle a,g'\rangle\in g(\delta):g'\sim_c g.$$

References

Aloni, M. (2000). Conceptual covers in dynamic semantics. In Cavedon, L., Blackburn, P., Braisby, N., and Shimojima, A., editors, *Logic, Language and Computation*, volume III.

Aloni, M. D. (2001). *Quantification Under Conceptual Covers*. PhD thesis, University of Amsterdam, Amsterdam.

Avigad, J. and Zach, R. (2020). The epsilon calculus. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, fall 2020 edition.

Beaver, D. (2001). *Presupposition and Assertion in Dynamic Semantics*. CSLI Publications: Stanford, CA.

Bekki, D. (2013). A type-theoretic approach to double negation elimination in anaphora. In *Proceedings of Logic and Engineering of Natural Language Semantics (LENLS)*, volume 10, Tokyo.

van Benthem, J. (1996). *Exploring Logical Dynamics*. Center for the Study of Language and Information, 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.

Brasoveanu, A. (2007). *Structured nominal and modal reference*. PhD thesis, Rutgers.

Brasoveanu, A. (2010). Decomposing modal quantification. *Journal of Semantics*, 27:437–527.

Brasoveanu, A. (2013). Modified numerals as post-suppositions. *Journal of Semantics*, 30(2):155–209.

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.

Charlow, S. (2014). *On the Semantics of Exceptional Scope*. PhD thesis, NYU.

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.

Cumming, S. (2015). The dilemma of indefinites. In Bianchi, A., editor, *On Reference*. Oxford University Press.

Dekker, P. (1993). *Transsentential Meditations: Ups and Downs in Dynamic Semantics*. PhD thesis, University of Amsterdam.

Dekker, P. (1994). Predicate logic with anaphora. In Harvey, M. and Santelmann, L., editors, *Semantics and Linguistic Theory (SALT)*, volume IV, pages 79–95.

Dekker, P. (2000). The semantics of dynamic conjunction. CSLI Publications: Stanford, CA.

Dekker, P. (2001). On *If and Only*. In Hastings, R., Jackson, B., and Zvolenszky, Z., editors, *Proceedings of Semantics and Linguistic Theory*, volume XI, pages 114–133. Cornell University.

Dekker, P. (2002). Meaning and use of indefinite expressions. *Journal of Logic, Language and Information*, 11:141–194.

Dekker, P. (2004a). Grounding dynamic semantics. In Reimer, M. and Bezuidenhout, A., editors, *Descriptions and Beyond*, pages 484–502. Oxford University Press.

Dekker, P. (2004b). The pragmatic dimension of indefinites. *Research on Language and Computation*, 2:365–399.

Dekker, P. (2008). A multi-dimensional treatment of quantification in extraordinary English. *Linguistics and Philosophy*, 31(1):101–127.

Dekker, P. (2012). *Dynamic Semantics*. Springer.

Dekker, P. (2015). Not *Only* Barbara. *Journal of Logic, Language and Information*, 24(2):95–129.

Dorr, C. and Hawthorne, J. (2013). Embedding epistemic modals. *Mind*, 122(488):867–913.

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.

Egli, U. and von Heusinger, K. (1995). The epsilon operator and e-type pronouns. In et al., U. E., editor, *Lexical Knowledge in the Organization of Language*, pages 121–141. Benjamins.

Elbourne, P. (2005). *Situations and Individuals*. MIT Press.

Elliott, P. D. (2020). Crossover and accessibility in dynamic semantics. Manuscript, MIT.

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

Fine, K. (1986). *Reasoning with arbitrary objects*. Blackwell: Aristotelian Society Series.

von Fintel, K. (1999). NPI licensing, Strawson entailment, and context dependency. *Journal of Semantics*, 16(2):97–148.

Frank, A. (1996). *Context Dependence in Modal Constructions*. PhD thesis, Universität Stuttgart.

Geach, P. (1962). *Reference and Generality*. Cornell University Press.

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.

Groenendijk, J. and Stokhof, M. (1991). Dynamic predicate logic. *Linguistics and Philosophy*, 14(1):39–100.

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.

Heim, I. (2018). Do we need "plural information states"? Handout of talk given at Dynamic Semantics Conference, All Souls College, Oxford.

Herzberger, H. G. (1973). Dimensions of truth. *Journal of Philosophical Logic*, 2(4):535–556.

Hofmann, L. (2019). The anaphoric potential of indefinites under negation and disjunction. In Schlöder, J. J., McHugh, D., and Roelofsen, F., editors, *Proceedings of the 22nd Amsterdam Colloquium*, pages 181–190.

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.

Karttunen, L. and Peters, S. (1979). Conventional implicatures in Montague grammar. In *Syntax and Semantics 11: Presupposition*. Academic Press, New York.

Kibble, R. (1994). Dynamics of epistemic modality and anaphora. In Bunt, H., Muskens, R., and Rentier, G., editors, *International Workshop on Computational Semantics*, Tilburg.

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. (2014). Do we need dynamic semantics? In Burgess, A. and Sherman, B., editors, *Metasemantics: New Essays on the Foundations of Meaning*, chapter 9, pages 231–258. Oxford University Press, Oxford.

Lewis, K. (2020). Negation and anaphora. *Philosophical Studies*.

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. (2016). Dissatisfaction theory. In Moroney, M., Little, C.-R., Collard, J., and Burgdorf, D., editors, *Semantics and Linguistic Theory (SALT)*, volume 26, pages 391–416.

Mandelkern, M. (2019). Bounded modality. *The Philosophical Review*, 128(1):1–61.

Mandelkern, M. (2020). Dynamic non-classicality. *Australasian Journal of Philosophy*, 98(2):382–392.

Mandelkern, M. (2021a). If p, then p! *The Journal of Philosophy*.

Mandelkern, M. (2021b). Pseudo-dynamics. Manuscript, NYU.

Mandelkern, M. (fc.). Trivalence and quantification. In Egré, P. and Rossi, L., editors, *Handbook of Three-Valued Logics*. MIT Press.

Mandelkern, M. and Rothschild, D. (2020). Definiteness projection. *Natural Language Semantics*, 28:77–109.

Muskens, R. (1996). Combining Montague semantics and discourse representation. *Linguistics and Philosophy*, 19(2):143–186.

Neale, S. (1990). *Descriptions*. MIT Press, Cambridge, MA.

Nouwen, R. (2003). *Plural Pronominal Anaphora in Context*. PhD Thesis, Utrecht University, Netherlands Graduate School of Linguistics Dissertations 84. LOT.

Onea, E. (2013). Indefinite donkeys on islands. In *Proceedings of SALT 23*, pages 493–513.

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.

Rothschild, D. and Yalcin, S. (2015). On the dynamics of conversation. *Noûs*, 51(1):24–48.

Rothschild, D. and Yalcin, S. (2016). Three notions of dynamicness in language. *Linguistics and Philosophy*, 39(4):333–355.

Schlenker, P. (2008). Be articulate: a pragmatic theory of presupposition projection. *Theoretical Linguistics*, 34(3):157–212.

Schlenker, P. (2009). Local contexts. *Semantics and Pragmatics*, 2(3):1–78.

Schlenker, P. (2010). Presuppositions and local contexts. *Mind*, 119(474):377–391.

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.

Schwarz, F. (2009). *Two Types of Definites in Natural Language*. PhD thesis, University of Massachusetts Amherst.

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.

Stalnaker, R. (1978). Assertion. In Cole, P., editor, *Syntax and semantics*, volume 9, pages 315–322. New York: Academic Press.

Stone, M. (1999). Reference to possible worlds. RuCCS Report 49, Rutgers University.