Philosophical Studies (2021) 178:1403–1440
https://doi.org/10.1007/s11098-020-01489-w

Anaphora and negation

Karen S. Lewis[1]

Published online: 23 July 2020

Abstract
One of the central questions of discourse dynamics is when an anaphoric pronoun is licensed. This paper addresses this question as it pertains to the complex data involving anaphora and negation. It is commonly held that negation blocks anaphoric potential, for example, we cannot say "Bill doesn't have a car. It is black". However, there are many exceptions to this generalization. This paper examines a variety of types of discourses in which anaphora on indefinites under the scope of negation is felicitous. These cases are not just of intrinsic interest, but I argue present serious problems for the dynamic semantic framework, which builds the licensing facts into the semantics. I argue in favor of adopting a dynamic pragmatics, a theory that explains context change through general Gricean principles, and combining it with a static, d-type theory of anaphora, in which pronouns go proxy for definite descriptions.

Keywords Anaphora · Negation · Dynamic semantics · Dynamic pragmatics · Descriptions · Pronouns · Discourse

1 Introduction

The task of an account of discourse dynamics is to explain how information flows through a discourse, that is, how the context is affected by content and how content is affected by context. One specific area of interest is pronouns that are anaphoric on indefinite descriptions, as in (1):

(1) a. A woman walked in.
b. She sat down.

'She' in some sense co-refers with 'a woman', but since the indefinite 'a woman' is not a referring term, we cannot simply say that they refer to the same thing. Nor can

✉ Karen S. Lewis
klewis@barnard.edu

[1] Department of Philosophy, Barnard College, Columbia University, Milbank Hall, 3009 Broadway, New York, NY 10027, USA

we say that 'a woman' syntactically binds the pronoun, since it is beyond its syntactic scope. This gives rise to two questions: 1) what sort of semantic value do pronouns like 'she' get? (call this *the semantic question*) and 2) when are pronouns like 'she' acceptable? (call this *the licensing question*). Unlike referring terms, indefinites do not always allow for anaphoric reference with subsequent pronouns. For example, it is uncontroversial that in general, indefinite descriptions in ordinary, simple, affirmative sentences like (2a) license anaphoric pronouns, whereas those under the scope of negation, as in (3a) do not:

(2) a. Bill has a car.
b. It is black.

(3) a. Bill doesn't have a car.
b. #It is black.[1]

Such data has motivated many theorists to think of semantics dynamically, taking discourse and context change as central.[2] In a dynamic semantics, *context change potential* (CCP) rather than truth conditions is the basic semantic notion. Contexts are information states. They keep track of the state of the discourse at any given point: what information has been accepted, what objects or questions are under discussion, etc. In dynamic semantics, indefinites introduce novel *discourse referents* into the context. Discourse referents are representations of objects under discussion, that is, they represent that according to the conversation, there is some individual that bears certain properties (in the case of (1) *being a woman*, *walking in*, *sitting down*). In this way, they can serve as vehicles for acting as though there is a specific object under discussion and for recording information about it without being committed to discourses like (1) being about a particular woman. Furthermore, discourse referents need not be satisfied by any object in the world (we could easily substitute 'unicorn' for 'woman' in (1) and still have a felicitous discourse). Anaphoric pronouns pick up on existing discourse referents. Traditional dynamic semantics answers the licensing question and the semantic question with the same tools: the discourse referent provides the value for the pronoun, but also explains what licenses it—since pronouns presuppose that there is a familiar discourse referent in the context to update, the fact that (1) introduces a discourse referent explains why the pronoun is felicitous. Operators like 'not' act like plugs; as part of their context change potential, they block the anaphoric potential of indefinites within their scope, disallowing the introduction of a novel discourse referent into the context (this will be further discussed in § 2 below). Hence it is predicted that pronouns like that in (3b) are infelicitous.

A popular static account of discourses like (1) is to treat pronouns as *d-type*, which means they go proxy for definite descriptions like 'the woman who walked in'.[3]

[1] Examples from Karttunen (1976, 4).

[2] E.g. See Kamp (1981) and Kamp and Reyle (1993) on Discourse Representation Theory, Heim (1982) on File Change Semantics, Groenendijk and Stokhof (1991) on Dynamic Predicate Logic, and Muskens (1991), Chierchia (1995), Groenendijk et al. (1996), Beaver (2001) for different versions of update semantics.

[3] E.g. See Evans (1977), Cooper (1979), Davies (1981), Heim (1990), Neale (1990), Elbourne (2005, 2013). D-type pronouns are also sometimes called e-type pronouns.

Definite descriptions are then treated in a classical way, either in a Russellian or Fregean manner.[4] D-type theory, therefore, is primarily an answer to the semantic question. D-type theory, since it doesn't focus on discourse, has historically been less interested in the licensing question as a central focus (e.g. Elbourne (2005, 2013) does not address it at all). Nevertheless, d-type theorists do generally address the question to some extent. Unlike dynamic semantics, the answer has nothing to do with the licensing or blocking powers of operators. For example, Neale (1990) gives a pragmatic answer to the licensing question. On his view, pronouns go proxy for Russellian definite descriptions, and so have no presuppositions at all. Rather, according to him, a discourse like (3) is always straightforwardly contradictory, since the first sentence asserts that Bill doesn't have a car, and the second sentence asserts that there is a unique car that Bill has that is black. For Neale, the infelicity of the discourse is due to the fact that "it would make no practical sense to use these sentences" (p. 232). I will discuss the different ways in which d-type theorists address the licensing question in § 3.

Negation is not the only operator of interest when it comes to anaphora. There is interesting, complex data involving quantifiers, modals, conditionals, and disjunction. Nevertheless, the present paper will focus on the relationship between anaphora and negation. The licensing question is the main focus of this paper, though the semantic question will of course come up along the way.

When it comes to anaphora licensing, the negation data is more complicated than the simple contrast between (2) and (3). It isn't true that negation always blocks felicitous anaphoric reference, as the following examples show. Call these examples of *external felicity*.

In the first class of cases, the pronoun is under the scope of a modal, which sometimes, but not always, makes anaphora felicitous:

(4) a. There wasn't a thief here.
    b. He would have had to have been magical (to break in without leaving a trace).

(5) a. I don't have a microwave oven.
    b. I wouldn't know what to do with it.[5]

An intuitive, pre-theoretic gloss on these cases is that we understand them as though there is an implicit counterfactual antecedent in the second sentence like "if there was a thief" or "if I did have a microwave".

The second class of cases involves a second negation taking scope over the anaphoric pronoun, which also sometimes, but not always, makes anaphora felicitous:

(6) a. Mary doesn't have a car.
    b. So she doesn't have to park it.[6]

[4] But see also King (1987, 1991, 1994) for another static option, Context Dependent Quantifier (CDQ) theory.

[5] Geurts (1999, 188).

[6] This kind of example was pointed out to me by Maria Bittner (p.c.).

This sort of example is felicitous in a context in which someone has said something about Mary parking her car. For example, imagine the following discourse:

(7) a. A: Where's Mary? She's never so late.
b. B: Maybe she's having trouble finding parking.
c. A: That can't be it. Mary doesn't have a car, so she doesn't have to park it.

In a way, A is playing along with B's presupposition that Mary has a car in her denial of that very fact.

A third class of cases is when there is reason to think there was a relevant object in the past, even though there is no such object now.

(8) a. Bryan doesn't have an apartment in Paris (anymore).
b. He gave it up years ago.[7]

Discourse (8) with the 'anymore' presupposes that Bryan used to have an apartment in Paris, and without the 'anymore' strongly suggests it.

Related to this last type, there are also cases in which what is being denied is not the existence of a certain entity (in this case, a restaurant that Moses's daughter owns), but a certain property of that entity (it being a sushi restaurant):

(9) a. Moses doesn't have a daughter who owns a sushi restaurant.
b. It's a pizza restaurant.

This category is related to the last because one way of glossing the last category is that what is being denied is not the existence of Bryan's apartment in Paris, but its property of it *presently* being Bryan's Paris apartment.

The external felicity cases are varied: they are not all triggered by a uniform type of linguistic expression. Some of them involve modals or negation, other involve presupposition triggers like 'anymore', but still others involve no remarkable linguistic expressions at all (as in (8) without the 'anymore' or (9)).

Neither of the major theories of anaphora currently account for the full range of licensing data involving anaphora and negation. A careful examination of ways in which each attempts to—or might attempt to—address the data suggests that the licensing question is more amenable to a primarily pragmatic than a primarily semantic answer. The next section examines and rejects the ways in which dynamic semantics has dealt with the external felicity cases. § 3 looks at the ways in which d-type theory falls short, arguing it is incomplete as an answer to the licensing question. § 4 turns to my positive view. I argue in favor of combining d-type theory with a *dynamic pragmatics*, a theory that takes context change as central, but explains context change in pragmatic rather than semantic terms. The dynamic pragmatics accounts for both the basic cases like (1) as well as the felicitous and infelicitous negation cases by appealing to broadly Gricean considerations, speakers' discourse plans, and hearers' discourse expectations. § 5 turns to the question of sentence-internal anaphoric links

[7] Examples of this kind are found in Chierchia (1995) and Elbourne (2005).

under the scope of negation, arguing that these do not require appealing to discourse dynamics at all.

2 Dynamic semantics and the licensing question

2.1 Negation in dynamic semantics

To get a sense of how dynamic semantic negation works, I'll begin by considering a simple update semantics in which the context (i.e. information state) only tracks truth-conditional information. Here we can think of a context as a set of worlds, namely the set of worlds that are possible given the information already established in the conversation (the *context set* as defined by Stalnaker (1978)). Where 'C' stands for context and [ ] indicates update on C, negation is defined as follows:

$$C[\neg\phi] = C - C[\phi]^{8}$$

That is, the negation update works as follows. The input context C is updated with the material under the scope of the negation, and the result of this update is subtracted via set subtraction from the input context. If we are considering a sentence like 'It is not the case that Stacey wants candy', then C is updated with ⟦Stacey wants candy⟧, thus yielding the subset of worlds from C in which Stacey wants candy. This set is then subtracted from the original input context, yielding an output context in which there are no worlds in which Stacey wants candy.

For our purposes, the context must also track discourse referents. Contexts are here defined as sets of assignment function/world pairs. The (partial) assignment functions are functions from indices (discourse referents) to entities, namely all those possible entities that satisfy the properties associated with a discourse referent. Speaking generally and glossing over the differences between theories, negation encodes an instruction to update the context with *the anaphoric closure* of the material under the negation. That is, it encodes an update just like the one in the simple system just described, except that any discourse referents introduced during the calculation of ϕ are prevented from having an effect outside the scope of the negation. Consider the following example:

(10) It is not the case that a woman walked in.

The clause for negation tells us to temporarily update with the sentence in the scope of the negation, i.e., *a woman walked in*. As explained in the previous section, this adds a discourse referent for a woman who walked in to the context. All assignment function/world pairs are changed such that the assignment functions assign a new discourse referent index (say *1*) to a woman who walked in (in the paired world), for each such possible assignment/world pair. But the new discourse referent does

[8] See for example, Yalcin (2012), Veltman (1996), Heim (1983), Beaver (2001), ch. 4.

not survive. Negation tells us that the assignment functions in the input and output context are identical; they cannot be extended with new indices. The informational content is still subtracted from the input context, and so we are left with just the subset of assignment function/world pairs in which no woman walked in. The input assignment functions remain unchanged; most crucially, they have not been extended with the index *1* in the output context. In this way, negation acts as a plug, blocking anaphoric reference outside its scope, predicting external *infelicity* of pronouns that attempt to be anaphoric on indefinites under the scope of negation. In other words, negation is externally static (it disallows anaphoric connections beyond its scope) though internally dynamic (it allows anaphoric connections within its scope). Another way of stating this is that negation allows anaphoric connections within the local context under negation, but not in the global context. A global context is a context for a whole discourse, whereas a local context is a temporary context created in the calculation of the CCP of a sentence.[9]

Dynamic semantics builds these kinds of constraints into all quantifiers and operators, for example 'every' is also generally treated as externally static and internally dynamic, predicting the felicity of (13a) and the infelicity of (13b), and disjunction is often treated as both internally and externally static.

[9] It is helpful to look at some specific examples of clauses for negation in dynamic semantics. For example, in Beaver (2001)'s system ABLE (ch.7), the clause for negation is as follows:

(11) $\quad [\![\neg\phi]\!] = \lambda I \lambda J [\exists K\, I \downarrow [\![\phi]\!] K \wedge J = I \backslash K]$

In ABLE, an information state tracks assignment function/world pairs. Negation is a function from an input information state I to an output information state J, such that there is some intermediate information state K, where K is the output of updating I with the anaphoric closure of $[\![\phi]\!]$. Updating I with the anaphoric closure of $[\![\phi]\!]$ involves there being some further state L which is the result of updating I with $[\![\phi]\!]$, and the output, in this case K, is the subset of the assignment function/world pairs in I which have extensions in L, that is, where an extension of an assignment function g is one that is the same as g except that it has a larger domain (i.e. it assigns values to more indices). In this sense, we see here that anaphoric connections are allowed within ϕ, in particular in the update of I to L, the intermediate state in calculating the anaphoric closure of ϕ, but disallowed outside of it (once we get to the anaphoric closure).
Groenendijk and Stokhof (1991)'s Dynamic Predicate Logic (DPL) is only concerned with tracking anaphoric information, so contexts only track assignment functions. In DPL, the clause for negation is:

(12) $\quad [\![\neg\phi]\!] = \{\langle g, h\rangle |\ \ h = g\ \ \&\ \ \neg\exists k : \langle h, k\rangle \in [\![\phi]\!]\}$

The criteria that the input assignment function g has to be identical to the output assignment function h is what makes negation externally static. In other words, no changes can be made to assignment functions outside the scope of negation, which is what represents that no discourse referents are introduced into the global context. Thus external infelicity is predicted. To calculate a negation, the material in the scope of the negation, however, is processed as a whole, including any anaphoric relations within it. (Heim (1982)'s definition of negation works similarly).
Chierchia (1995)'s Dynamic Intensional Logic (DIL) also predicts external infelicity and internal felicity, since the clause for negation is: $\underline{\neg}\mathrm{A} = \uparrow \neg \downarrow \mathrm{A}$ (where the underlined negation sign is dynamic negation and the other is regular static negation). The down arrow is here an assertion operator—it takes a CCP to a static proposition—and the up arrow is the opposite—it maps a static proposition to its corresponding CCP. So A is calculated normally (with all the dynamic CCPs it may contain inside), but the negation blocks its context change potential by taking only its static content, negating it, and turning the result into a test (its corresponding CCP) on the context.

(13) a. Every farmer who owns a donkey vaccinates it.
b. #It lives in a stable.

In this way, answering the licensing question is a key component of (many) dynamic semantic theories.

2.2 External felicity and dynamic semantics

The existence of external felicity examples on their own does not mean that dynamic semantics is wrong. First of all, the dynamic framework itself does not necessitate a particular treatment of negation—other definitions of negation are possible, as are more complex treatments of the semantics more generally, which have been adopted by second-generation dynamic semantic theories to account for quantificational subordination,[10] plural anaphora, and other phenomena not accounted for by first-generation dynamic semantics.[11] Second of all, just because some of the data (i.e. the external infelicity cases) is accounted for by the semantics, it doesn't mean that other cases (i.e. the external felicity cases) can't be accounted for by something else, e.g. perhaps supplementary pragmatic considerations, especially if such considerations are independently motivated. However, I will argue that the existing tools used by dynamic semantics to explain (or to potentially explain) these cases are all problematic. Insofar as an important feature of dynamic semantics is to answer the licensing question, this is a serious problem for the framework. There are four categories of proposed solutions in the literature that I know of; these appear to exhaust the viable options for solutions. They are: an ambiguity in pronouns, an ambiguity in operators and quantifiers, modal subordination, and accommodation/pragmatic repair. I will examine each one in turn.

Chierchia (1995) argues that some pronouns should be treated as dynamically bound variables, while others are d-type pronouns (while still others are ambiguous between both interpretations). Examples like (8) are one of the motivations for adopting this view; he treats the pronoun in (8b) as a d-type pronoun. As mentioned previously (in footnote 9), Chierchia's semantics, a dynamic version of Montague's intensional logic, builds constraints on anaphoric connections into the meaning of operators like negation. Negation turns a context change potential into a test, i.e., it checks the input context for a particular property, but does not have the power to change the context. Hence, nothing in the scope of a negation establishes a discourse referent for future anaphoric connection. This nicely explains the many cases of external infelicity. But Chierchia also recognizes that there are many cases of external felicity. He thinks these are rightfully treated separately since they are "highly sensitive to various aspects of the context—what is known or presupposed by the speaker, the specific properties of the lexical items involved in interaction with what the extralinguistic facts are, and so

[10] Quantificational subordination is the phenomena by which anaphoric relationships hold across two quantificational sentences; that is, the indefinite antecedent is under the scope of one quantifier and the pronoun under the scope of another. For example, "Harvey courts a woman at every convention. She always comes to the banquet with him". (Example from Brasoveanu (2010)).

[11] Classic examples of first generation dynamic semantics includes Kamp (1981), Heim (1982), Groenendijk and Stokhof (1991). Second generation dynamic semantics and beyond is too prolific to include in its entirety, but some relevant references are van den Berg (1996), Brasoveanu (2010), and Keshet (2018).

on" (p. 9) while other cases, like donkey anaphora are "no more affected by contextual and pragmatic factors than ordinary c-command [i.e. bound] anaphora" (p. 10). Thus Chierchia thinks it is plausible to study them (and presumably account for them) independently. The sort of pragmatic, contextual factors that Chierchia is alluding to are exhibited by examples like (8), where background facts are such that it is known (or communicated by the second sentence) that Bryan used to have an apartment in Paris. Of course, I'm sympathetic to the claim that felicitous anaphora is sensitive to various contextual features—this is in fact what I will argue—but Chierchia does not give an account of when a d-type pronoun is licensed. Furthermore, though Chierchia doesn't think that his mixed theory amounts to an ambiguity theory, since he treats both kinds of pronouns as variables at the level of syntax, they are still interpreted differently by the semantics, which is in the end an ambiguity theory: some are dynamically bound and others get assigned a value by a salient function. Since there doesn't seem to be any evidence of an overt ambiguity in pronouns in any language, positing an ambiguity theory should be a last resort.[12] So at best, Chierchia's theory is incomplete in two ways. One, it is not really a defense of dynamic semantics in terms of an answer to the licensing question. On Chierchia's view, dynamic semantics applies where it applies and when it doesn't apply, d-type theory does. Two, it needs to be supplemented with a pragmatic account of when a d-type pronoun is licensed. At worst, Chierchia's theory is simply wrong, positing an ambiguity where there is no evidence of one.

The second proposed solution in the literature is ambiguity in negation instead of ambiguity in the pronouns. In other words, the hypothesis is that *sometimes* negation is externally static and sometimes it is externally dynamic. This view is tentatively proposed by Groenendijk and Stokhof (1990). When it comes to negation, they are motivated by cases in which pronouns anaphoric on an indefinite embedded under double negation are felicitous. In dynamic semantics, double negation is not equivalent to the affirmative — a second negation does not allow anaphoric connections across sentences. Rather than a plug unplugged, double negation is a double plug. This sometimes makes the right predictions, as evidenced by (14):

(14) a. It is not the case that no man walks in the park.
b. # He whistles.

But this is not always right, as Groenendijk and Stokhof observe:

(15) a. It is not the case that John doesn't own a car.
b. It is red and it is parked in front of his house. (Groenendijk and Stokhof 1990, 27)

So they posit an externally dynamic negation in addition to the externally static negation operator. The view is not exactly an ambiguity theory, since the static and dynamic interpretations of the operators are not independent—the static interpretation is reached by a closure operation on the dynamic interpretation. By their own admission, they

[12] Kurafuji (1998, 1999) argues that there is evidence for a dynamically bound variable/d-type distinction in third-person Japanese pronouns. Elbourne (2005) (pgs. 26–31) argues persuasively that the data does not support this ambiguity.

don't have a general theory that predicts when a particular instance of negation should be interpreted as dynamic or as static.[13] But suppose such an account could be worked out—perhaps something like the pragmatic account I ultimately support—dynamic semantics still loses its predictive power regarding anaphora licensing facts, since for any given instance, for all the formal semantics tells us, we could have an externally dynamic interpretation of the operator or quantifier, predicting that anaphora is licensed, or an externally static interpretation, predicting that anaphora is blocked. It would be the pragmatic story (that presumably determines which version of the operator is in play) that accounted for these facts. Perhaps this is a good way to go for those who think that dynamic semantics is required for other reasons, e.g. that we need dynamic binding in the compositional semantics to account for donkey anaphora or something like that. But this solution simply abandons dynamic semantics as an answer to the licensing question. Furthermore, even though the two interpretations of the operators and quantifiers are related, it is on the face of it a less parsimonious theory than one that doesn't posit multiple interpretations for each operator and quantifier.

It has been defended by Roberts (1996) and Frank and Kamp (1997), among others, that cases like (4) and (5) are cases of modal subordination. Modal subordination is the phenomenon by which the domain of a modal in a discourse is interpreted relative to a set of possibilities introduced by an earlier sentence. Pronouns anaphoric on indefinites under the scope of a modal are licensed so long as the discourse continues to be modal (in a way that is understood to be subordinate to the initial possibility introduced), as in:

(16) a. A wolf might come to our campsite.
b. It would eat you first.

For the negation/modal cases such as (4) and (5), the idea roughly goes, that the modal is restricted to counterfactual worlds introduced by the negated sentence, i.e., worlds in which there is a thief or the speaker does have a microwave oven.[14]

Note that modal subordination, unlike the previous two proposed solutions, only applies to cases that contain modals, so it is not a complete solution. This is not an argument against it—it is entirely possible that the different kinds of cases require different solutions. As I glossed the cases above, the natural interpretation of (4) and (5) can be paraphrased as "if there was a thief…" and "if I had a microwave…". A modal subordination account takes this gloss literally. The problem with the solution is that the negation data is different from ordinary modal subordination data in a way that suggests they shouldn't be assimilated. First of all, pronouns in an ordinary modal subordination case are equally felicitous whether the two sentences succeed one another, separated by a period, as in (16), or are conjoined as in (17):

[13] They do propose that whatever the rules are, they obey a monotonicity constraint, namely that $\neg\phi$ is interpreted dynamically only if ϕ is downward monotonic (so that every step in a discourse is upward monotonic in the sense that we never lose truth conditional information as updates occur). But even if this is right, this only provides a necessary condition.

[14] Since it doesn't matter for present purposes, I am glossing over the details of the accounts, including important differences between them, such as whether modal subordination is accounted for by antecedent accommodation (Roberts) or anaphorically (Frank & Kamp). Technically, on Roberts's view, the pronoun is not anaphoric on the original indefinite, but on the indefinite in the accommodated material, but again, nothing rests on such details for the purposes of my argument.

(17) A wolf might come to our campsite and it would eat you first.

This contrasts with the case of negation. Though (4) is perfectly felicitous, its conjunction counterpart is significantly marked, if not downright infelicitous:

(18) ??There wasn't a thief here and he would have to have been magical to get in without leaving a trace.[15]

Second, as Geurts points out, the modal subordination account predicts that examples like (19) should be felicitous (because the second sentence should be interpreted as though it read 'if someone had turned up…'), though they are not:

(19) ?Nobody turned up at Betty's party. He would have seen her wearing a gorgeous evening dress.[16]

There are plenty of other cases that also support the claim that modal subordination doesn't generally occur with the negation/modal pattern. Consider:

(20) ?There wasn't a visitor today. He would have been happy to be here.

(21) ?My sister doesn't have a dog. It might have been a poodle.

Note that the modal versions of these sentences are perfectly fine:

(22) There could have been a visitor today. He would have been happy to be here.

(23) My sister might have had a dog. It might have been a poodle.

Whatever the phenomenon of external felicity in negation cases is, it is a lot more limited than modal subordination. Theories that appeal to modal subordination to explain the felicity of these cases cannot account for these limitations; they predict a widespread phenomenon. This is not to say that in the end, one couldn't use the tools of modal subordination in the formal account of the modal cases, but it would need to include something else to explain when the cases are felicitous and when they are not; the presence of the modal and a modal subordination treatment alone doesn't do that.

Probably the most popular way to address these cases (or at least some of the cases) is to appeal to accommodation. Accommodation is the phenomenon by which the context is tacitly adjusted so that something that would otherwise be infelicitous is felicitous. Accommodation is often appealed to in dynamic semantics to account for informative presuppositions; most relevantly here, there is a history of appealing to accommodation to satisfy the familiarity presupposition on definites. For example,

[15] The same goes for (5). Note that (18) is fine with 'because' instead of 'and'.

[16] (Geurts 1999, 189). Geurts has other examples, but they are problematic because the modal sentence is in the indicative mood, which Frank and Kamp (1997)'s theory rules out.

one can say (24) discourse initially, even if the hearer is not aware that the speaker has a tree in her backyard:

(24) I'm really sad. The apple tree in my backyard got a disease and needs to be cut down.

It is commonly thought that a discourse referent for an apple tree in the speaker's backyard is accommodated to satisfy the familiarity presupposition. In fact, van der Sandt and Maier (2003) appeal to accommodation to account for an example similar to (6), except that it has a possessive definite rather than a pronoun:

(25) Sally has no dog, so it was not Sally's dog that bit her neighbour. (van der Sandt and Maier (2003, 10))

Since the presupposition that Sally has a dog is not satisfied by the context, it is accommodated. Normally global accommodation is preferred,[17] but since accommodating the information that Sally has a dog conflicts with information accepted in the global context (that Sally has no dog), it is accommodated locally, under the scope of the negation in the second part of the sentence. We can imagine a similar strategy applied to the other cases, where local accommodation is employed where the accommodated discourse referent conflicts with the accepted information (as in the modal and negation cases), but globally accommodated in the other cases (as in (8) and (9)).

There are some issues, I think, with employing widespread accommodation in a dynamic semantic framework; dynamic semantics is supposed to account for discourse phenomena that are thought to be too systematic to be amenable to a pragmatic account. Accommodation as pragmatic repair repeatedly comes in just when the phenomena are not as systematic as they seemed, and the predictions of the account fail. However, I will not pursue this line of criticism further here.[18] Even if appealing to accommodation is generally a completely legitimate strategy for dynamic semantics in the face of recalcitrant data, the problem is that it is not available in this case. There is a long history in dynamic semantics of viewing the anaphoric relationship between antecedent and pronoun as more stringent than other discourse relationships—the thought is that there has to be a formal, linguistic link between indefinite antecedent and anaphoric pronoun. This is often cited as a central reason in favor of dynamic semantics.[19] For example, van der Sandt (1992) argues that one of the central differences between presupposition and anaphora is that "*unlike pronouns* they [presuppositions] contain descriptive content which enables them to accommodate an antecedent in case discourse does not provide one" (van der Sandt 1992, 341, emphasis mine). This "more descriptive content" difference is also often invoked to explain why there doesn't have to be the same sort of formal link between definite descriptions and their antecedents as there does with pronouns:

[17] Global accommodation is when the necessary adjustment (such as the introduction of a discourse referent) is added to the global (rather than local) context.

[18] See Gauker (2008) for more arguments against accommodation.

[19] See e.g. Heim (1982, 1990), Elbourne (2005).

(26) a. Sam is married. His wife is a lawyer.
b. Sam is married. #She is a lawyer.
(27) a. If John is married, his wife is very patient.
b. # If John is married, she is very patient.
(28) a. Every married man is sitting next to his wife.
b. # Every married man is sitting next to her.

Another example of the formal link between antecedent and pronoun is that, infamously, being common ground that a certain object exists is not sufficient for licensing an anaphoric pronoun. In fact, these sorts of cases motivate dynamic semantic treatments:

(29) a. Bryan is a bicycle-owner.
b. #It is grey.
(30) a. Bryan owns a bicycle.
b. It is grey.

Of course, one can accept a dynamic semantics and give up the notion that pronouns don't allow for accommodation, perhaps even taking examples like the external felicity ones as evidence for this point. However, doing so both gives up dynamic semantics as an explanation of licensing and the formal link. The contrasts in the above examples, central to motivating dynamic semantics, would have to be explained pragmatically, by restrictions on accommodation. Accounting for the formal link between pronoun and antecedent is generally thought to be one of the strengths of dynamic semantics (and lack of explanation the weakness of its static rival, d-type theory).

In fact, all the proposed solutions canvassed above encounter a similar problem. Provided other problems with the individual accounts can be overcome, they all require adopting a pragmatic account that would play the primary role in answering the licensing question. The pragmatic account would explain when a pronoun was d-type or dynamic (if we adopted Chierchia's solution), or whether negation was externally dynamic or static (if we adopted Groenendijk and Stokhof's solution), or whether modal subordination or accommodation was licensed in a particular instance or not (if we adopted one of the latter two solutions). This is a serious problem for dynamic semantics as an answer to the licensing question.

3 D-type pronouns and the licensing question

D-type theories give a very different style of answer (when they give an answer at all) to the licensing question. D-type theories assign a classic static semantics in the tradition of Frege or Russell, so pronouns have no familiarity presuppositions, and therefore don't place constraints on the context in terms of requiring the presence of a discourse referent. Contexts don't track discourse referents, which don't play a role in the theory. Negation is treated in its ordinary truth-functional way, and so

the answer to the licensing question does not involve the blocking power of negation (or any other operator or quantifier). In part, whether a d-type pronoun is licensed depends on how the theory accounts for the recovery of the description in the first place. A theory that requires there to be, e.g. syntactic material from which to recover the description is less permissive of anaphoric pronouns than a theory that requires only that a function is made salient. Regardless, whether a d-type theory requires there to be prior linguistic material from which to recover the description or simply recovers a description from one made contextually salient, both strategies are available in cases in which the pronoun attempts to be anaphoric on an indefinite under the scope of negation, as in (3). So something else needs to be said to explain why such pronouns are not licensed.

Different d-type theorists have explained the infelicity of discourses like (3) in different ways. Evans (1977) stipulated two well-formedness conditions for d-type descriptions:

a. The antecedent determiner must be existential in force.
b. The antecedent-containing sentence must be affirmatively embedded relative to the minimal sentence that contains the pronoun.[20]

So, for example, the determiner 'no' is not existential in force, and so there is no well-formed d-type pronoun anaphoric on sentences like 'No woman walked in', and 'a woman' is not affirmatively embedded under negation relative to 'she sat down' in a subsequent sentence, hence 'It is not the case that a woman walked in' also doesn't license cross-sentential anaphora.[21] Heim (1990) treats descriptions in the Fregean tradition, which means they presuppose uniqueness of referent. Heim therefore attributes the infelicity to presupposition failure; for example, if the first sentence of (3a) is true, Bill has no car, so there's nothing that satisfies the uniqueness presupposition of the definite "Bill's car".[22] As I mentioned in the introduction, Neale (1990) argues that discourses like (3) are infelicitous because they are straightforwardly contradictory on his Russellian treatment of definite descriptions, since they are of the form "Bill does not have a car. There is a unique car that belongs to Bill that is black". I will assume that it is best to avoid stipulating conditions as Evans does, and we'd prefer an answer to the licensing question that falls out of semantics and/or pragmatic considerations, so I will focus on the style of answers given by Heim and Neale.

[20] The definition of affirmative embedding is a little hard to apply here, because (1) Evans is only thinking about sentences rather than discourses when he defines affirmative embedding and (2) he is thinking of pronouns as referring expressions, the reference being fixed by a definite description. The definition is as follows: Let $\Sigma(\sigma,\sigma')$ be a sentence embedding an existential sentence, σ, and a sentence σ' that contains a pronoun anaphoric on the indefinite in σ. σ is affirmatively embedded in Σ relative to σ' iff when the truth of Σ turns on the truth or falsity of σ', σ is true. Intuitively, the idea is that whenever a sentence's truth turns on the truth or falsity of the pronominal sentence contained in it, there is something that the pronoun refers to. We can extend this idea to discourses by thinking of the discourses as the conjunction of the sentences within them, and we could tweak the definition to better suit d-type theories so that the requirement is that there is a unique object that satisfies the description.

[21] King (1994) claims something similar with respect to his CDQ theory, "that an occurrence of a quantifier in a sentence must be existentially positive to support subsequent (simple) anaphora in another sentence"(p. 229), where a quantifier being existentially positive means it is not non-existence entailing, i.e. (Dx:Fx)Gx does not entail that the intersection of F and G is empty.

[22] Elbourne (2005, 2013) is silent on these cases; his theory runs into the same problems, but he could appeal to the same considerations as Heim.

Both Heim and Neale encounter difficulties in accounting for the full range of data. Both predict that anaphora on antecedents embedded under double negation should be just fine, since the doubly negated sentence results in neither presupposition failure nor contradiction with the pronominal sentence. But as we've already seen in § 2, examples like (14) are infelicitous. So whereas dynamic semantics has trouble accounting for double negations that *allow* for anaphora, d-type theory has trouble with double negations that *block* anaphora.

Furthermore, Neale's contradiction story predicts that all cases in which the anaphoric pronoun is under the scope of a negation or a modal should be good. There is no contradiction between asserting there does not exist an F and there does not exist/there might/would exist a unique F that is G. This prediction is wrong; some modal cases are bad, as I argued in § 2 (see examples (19)-(21)). The same goes for negation cases. (31) is not a contradiction (on one of its readings), but the pronoun is still infelicitous.

(31) Bill doesn't have a car.
#It is not black.

Since presupposition projects out of modal and negated environments, unlike Neale, Heim predicts that the cases in which a pronoun is scoped under a modal or negation should be bad, since the presupposition is not satisfied. While this accounts for the infelicitous cases, something else needs to be said to account for the external felicity cases as in (4)–(6).

Dynamic semantics attempts to answer the licensing question by appealing to the semantics of the linguistic expressions involved. In the face of recalcitrant data, it seems that at best, dynamic semantics has to be supplemented with a pragmatic account. On the other hand, d-type theorists give a mostly pragmatic answer to the licensing question (though it does involve the semantics of the pronouns). D-type theories have tended to be less focused on these questions, and as a result, the account is incomplete. I also think d-type theorists focus their explanations in the wrong place; solely in the content or presupposition of the pronoun rather than something to do with the connection between (potential) antecedent and pronoun. That is, true to many static accounts, they ignore the dynamics of discourse almost completely. I propose to solve this problem by combining a static d-type theory with a dynamic pragmatics.

4 Static semantics and dynamic pragmatics

In previous work, I have argued that we should adopt a static semantics with a dynamic pragmatics.[23] Dynamic pragmatics on my view takes context change seriously, and also accepts that the introduction and updating of discourse referents is part of what the context tracks, but this is an essentially pragmatic rather than semantic phenomenon. That is, whether a discourse referent is introduced or not is not about novelty being semantically encoded in indefinites, or blocking powers encoded in negation, but a matter of rational agents reasoning in the context of a communicative activity. This sec-

[23] See Lewis (2012, 2014).

tion extends my previous work in several ways: I provide a more detailed explanation of the background mechanisms in thinking of pragmatics in terms of plan recognition, I show how dynamic pragmatics can answer the licensing question in terms of accounting for both the external infelicity and felicity cases, and I show how it can combine with d-type theory to give a complete answer to the licensing and semantic questions.

As I argued in Lewis (2012), there are some good reasons for thinking that the introduction of a novel discourse referent is best treated pragmatically rather than semantically. If an indefinite semantically encodes an instruction to add a new discourse referent, then discourses like (32) should be very easy to interpret, and should have a salient reading equivalent to 'a woman walked in; another woman sat down'.

(32) A woman walked in. A woman sat down.

As a matter of fact, such discourses are quite confusing as to their intended meaning (especially absent a richer context), which is exactly the prediction of an account such as mine that takes the perspicuity of a speaker's plan as central to the introduction of a novel discourse referent.

Another point in favor of a pragmatic account is that the introduction of a novel discourse referent is both explicitly and contextually cancellable, in the sense of Grice (1989). (33) is an example of explicit cancellation of the introduction of a novel discourse referent and (34) an example of contextual cancellation. Suppose we've been talking for some time about a student of mine, Jane. Then I say:

(33) Sorry, I have to go now. I have a meeting with a student. In fact, it is Jane I'm meeting with.

(34) a. I passed my final exam for Chemistry.

b. Can you believe it? I passed a Chemistry exam!

In Lewis (2012), I called the latter *summary uses* of indefinites, which come in a variety of linguistic environments and never involve the introduction of a discourse referent.

I'm thinking of pragmatics in a broadly Gricean sense, in that pragmatic explanations appeal to the fact that rational agents are engaged in a basically co-operative communicative activity. I am not committed to the letter of Grice's explanations, nor do I think that the introduction of a discourse referent is a matter of implicating a *proposition* that the speaker wants to introduce a novel object under discussion or anything of the like. Rather it is the update to the context itself that is explained by reasoning about speakers' intentions as rational agents in a co-operative activity. I am indifferent to whether this should be properly called an implicature or not. I take some insights from Thomason (1990) as a starting point: if the general Gricean picture of pragmatic phenomena is right, then "it should be possible to single out certain important types of reasoning mechanisms and data structures that figure in communication among intelligent agents, and that work together to make implicature possible. These features should be independently motivated by linguistic and philosophical considerations, and should be theoretically central". (p. 330) Thomason goes on to argue that the applicable reasoning mechanism is plan recognition, and the data structure is

the conversational context. Though I start with these basic insights from Thomason, I develop them in very different ways from him.

A context contains, for our purposes, the context set, a set of questions under discussion, and a set of discourse referents. The context set is the set of worlds that are still open given the mutual presuppositions of the conversation at any given point. The set of questions under discussion record what question(s) is (are) currently being addressed in the discourse, and they are either introduced explicitly (by asking a question) or implicitly.[24] Discourse referents are indices that represent objects under discussion.[25] The context is updated not by rules encoded in the semantics, as in dynamic semantics, but as a result of the pragmatics. The context set is updated with the informational content of an assertion by eliminating all worlds that conflict with that content (for the reasons described by Stalnaker (1978)). I take this to be uncontroversial—it is generally accepted that a static semantics and dynamic pragmatics works like this. What is different about my view is that the system also includes discourse referents that are added and updated for pragmatic reasons. There is no reason why discourse referents can't be part of the context when combined with a static semantics.[26] Many people (not just in a dynamic semantic framework) accept that questions under discussion are an important part of the context. This is information about the discourse itself, rather than information about the world, just like the sort of information discourse referents track. Similar to tracking questions under discussion, I think that the context also tracks the objects under discussion.

On my view, the set of discourse referents is a formal representation of the psychologically real fact that conversational participants keep track of the objects under discussion, or things that are considered single objects according to the conversation, where objects under discussion may or may not be identified with specific objects in the world, and include hypothetical objects. This takes very literally the idea that "[i]ntuitively, a discourse referent is the address for a maximal cluster of information assumed by the interlocutors to bear on a single individual".[27] Conversational participants often make *as if* a particular object is under discussion, e.g. in (1), there is a sense in which we treat the conversation as being about a specific woman even though the truth of what the speaker has said does not depend on a particular woman. Truth conditionally, (1) is very different from (35):

(35) a. Michelle Obama walked in.
b. She sat down.

(1a) has existential truth conditions while (35a) has object-dependent truth conditions (its truth depends on Michelle Obama), but the tracking of objects under discussion

[24] See Roberts (2004, 2012) for more on questions under discussion. The set of questions under discussion likely has a hierarchical structure, but that is not important for present purposes, so here I treat it as a simple set.

[25] Formally the discourse referents are modeled as indices and the information associated with them as partial assignment functions; see "Appendix A" for more details.

[26] See Lewis (2012, 2014, 2017) for detailed arguments in support of this point.

[27] Roberts (2003), 294–5.

goes in exactly the same way. The assertion of (1a) introduces a discourse referent, say 1, which has the properties that 1 is a woman and 1 walked in. The assertion of (35a) introduces a discourse referent, say 2, which has the properties that 2 is identical to Michelle Obama and 2 walked in. While (35a) introduced a specific individual, we can think of sentences like (1a) as introducing a pseudo-specific individual.

In very informal terms, the basic idea of the dynamic pragmatic account is that new discourse referents are added to the context based on recognizing a speaker's plan to talk about an F in saying something of the form 'An F is G'. This supports discourse expectations that the same speaker or someone else in the conversation will go on to say more about an F. Since conversational contexts track the objects under discussion, a discourse referent is added. The presence of this discourse referent is what licenses subsequent anaphoric pronouns, which are used to talk about the salient objects under discussion. So much for the basic idea; on to the details.

4.1 Plans and plan recognition

Conversations are joint, co-operative activities. Following Bratman (1999) (ch.6), we can think of shared cooperative activities as being a matter of having a shared intention that we (i.e. the participants of the activity) do the activity in question. In other words, having a conversation involves a shared intention that we have a conversation. Shared cooperative activities have the following structure:

We intend to J if and only if:

1. (a) I intend that we J and (b) you intend that we J.
2. I intend that we J in accordance with and because of 1a, 1b, and meshing subplans of 1a and 1b; you intend that we J in accordance with and because of 1a, 1b, and meshing subplans of 1a and 1b.
3. 1 and 2 are common knowledge between us. (p. 121)

In this case, the J'ing in question is having a conversation. I am *not* arguing that the sub-plans in the conversation, i.e. the specific communicative and other discourse intentions that are part of the conversation are shared intentions.[28] Rather the shared intention is to engage in conversation together. Subplans *mesh* "just in case there is some way we could J that would not violate either of our subplans but would, rather, involve the successful execution of those subplans" (p. 120). As Bratman notes, this does not mean that we need to have all our subplans filled in when we set out to J. This is unrealistic for conversation, but also for many other activities. Nor does it require that our subplans *match*. Rather the commitment to a joint activity J involves the intention to J by way of subplans that do not conflict with each other, however they may be filled in along the way. One of Bratman's examples is a plan to paint a house together. If I have a subplan to paint the house entirely red, and you have a subplan to paint the house entirely blue, our subplans do not mesh. But if I have a subplan to buy red paint at Home Depot, and you have no plan about where to get red paint because you don't care, our subplans mesh. Importantly, if we are committed to

[28] This stands in contrast to a view like that of Jankovic (2014), who argues that communicative intentions are shared intentions.

meshing subplans, then this places contraints on how our subplans get filled in as we engage in the joint activity. For example, if you know my subplan is to paint the house entirely red, you cannot just form a subplan to paint it entirely blue (though you can of course try to change my plan). The same holds true for conversation. If we want to have a conversation, but we both have subplans to talk about different things, it is not going to work out very well. Similarly, if we have subplans to address unrelated questions at some point in the conversation, our subplans do not mesh, and we will have a less than successful conversation.

Following Bratman, let's define intentions as "elements of stable, partial plans of action concerning present and future conduct" (p. 2). I use the terms "discourse plans" and "discourse intentions" to subsume all kinds of plans speakers in a discourse might have. These include run-of-the-mill communicative intentions, but also plans involving the discourse itself, such as introducing a question under discussion or a new discourse referent. Applying Bratman's definition of intention to present purposes, understanding discourse intentions as parts of plans makes sense of the fact that discourses must cohere and so must our discourse intentions. If we are looking at things from a dynamic perspective, what is of interest is not just individual communicative intentions for particular utterances taken in isolation, but discourse intentions qua parts of a larger activity, the conversation.

Plans (or really, subplans or intentions) constrain the possible moves one can make. For example, planning on buying blue paint then makes it inappropriate to buy red paint, though it may still be appropriate to buy blue paint in a gloss or flat finish. Plans are also defeasible; we can of course change our minds, at least under many conditions. But they nevertheless enforce considerable constraints, otherwise there would be no point in having plans in the first place. The idea that subplans constrain the next possible moves is suggestive of common ways of thinking of the conversational context (stemming from Lewis (1979)), that is, the state of the context constrains what moves one can make next in a conversation. I embrace this suggestion, taking plans and plan recognition as a starting point for a dynamic pragmatics. Interlocutors have discourse plans, and these plans must mesh with each other. Combining this with the idea that conversational contexts are an important theoretical tool that represent what must be tracked in a conversation (and thus an important tool in achieving the goal of having a conversation), we get that the conversational context is updated based on speakers' intentions as part of their discourse plans.

One other thing that plans do is support expectations, both for the agent who has the plan and for the others involved in the activity. For example, my plan to go to the store later supports my expectation that I will do so, and if made plain to my husband, also supports his expectation that I will do so. Expectations play an important role in enabling us to coordinate our actions. Discourse plans also support expectations; the discourse plans of speakers support the discourse expectations of both the speaker and hearers. This allows interlocutors to coordinate their subsequent conversational moves. For example, asking a question supports an expectation that we will attempt to answer it. As Thomason emphasizes, this helps us interpret a speaker's subsequent utterance as a response to the question. Similarly, introducing an object under discussion supports

expectations that it will be picked up, and helps us interpret a speaker's subsequent use of a pronouns as picking up on that object.[29]

I take Gricean reasoning for the derivation of implicatures to be part of a more general pragmatic mechanism, i.e. plan recognition. Thomason wants to get rid of the Gricean maxims entirely in favor of direct plan recognition, but I prefer to think of using the Gricean maxims as ways in which hearers can come to recognize the speaker's plans. Likewise, from the speaker's perspective, since the cooperative principle and its associated maxims guide the the activity of conversation, when the speaker makes her plans, she can expect hearers to rely on the maxims in recognizing them.

4.2 Updating with novel discourse referents

With the general planning background in place, it remains to be explained why when a speaker says something like 'a woman walked in', it is generally part of a recognizable plan to add a new discourse referent for a woman to the context. (I take this to be a more theoretically loaded way of formulating the informal idea that the speaker has a plan to talk about a woman who is novel to the conversation).

In asserting an existential proposition like (1a), the speaker has explicitly asserted the existence of some object. This stands in opposition to the truth-conditionally equivalent 'The room was not devoid of women' and 'Not every woman did not walk in', which aside from being maxim of manner violations absent clear contextual reasons for using circumlocutory phrases, entail rather than assert the existence of a woman who walked in. Though truth-conditionally equivalent, sentences like (1a) contain a singular *denoting phrase* in the sense of Russell (1905), whereas their truth-conditional equivalents do not. Denoting phrases in some sense pick out objects, though perhaps by denoting rather than (semantically) referring to them. An existential proposition, therefore, is one that is expressed by a sentence containing an indefinite.[30]

Since interlocutors track objects that are under discussion in a discourse, explicitly asserting the existence of an object is (defeasible) evidence of a plan to update the set of discourse referents tracked in the context. The question is just what sort of plan is it—a plan to add a new discourse referent or add information about an existing one? As I said above, I take the Gricean maxims to be principles that speakers and hearers can take for granted in the forming and recognizing of plans, respectively. I take novelty to be a type *Quantity* implicature, where we are considering informativeness relative to the set of discourse referents in the context. The idea is that in most cases, if the speaker intended to update an existing discourse referent (informally, talk about something already under discussion), she could have said something more informative. The contextual alternatives to using an indefinite in a sentence, insofar as it is a singular denoting phrase, are: names, pronouns, definite descriptions, and demonstratives, since these

[29] One thing I have not mentioned here is that plans typically have a hierarchical structure, which general intentions embedding more specific intentions. It could be that the best way to make sense of discourse intentions to introduce a new discourse referent is to think of them as embedded in a more general communicative intention regarding the proposition expressed.

[30] The difference is captured at the level of the propositional content only if we have Russellian or structured propositions. Nevertheless, even a system with unstructured propositions can distinguish between sentences containing singular denoting phrases and those which do not.

are other ways in which to pick out single objects. In the case of (1), the alternatives for 'a woman' are: 'the woman', 'that woman' (and related demonstratives like 'that woman I was telling you about'), 'she' and possibly 'N', where 'N' is the name of the woman the speaker has in mind, if any. These don't form a Horn scale (this is not a *scalar* implicature), but they do form something like a contextually determined alternative set of the type posited by Hirschberg (1985). All the devices listed in the alternative set are common ways of picking up on an object already under discussion. This is not to presume that each of them presupposes familiarity, but merely that by some feature of their semantics, they can be used to denote an object already under discussion. Since the speaker did not use any of these (absent contextual or explicit cancellation, or other contexts in which for whatever reason a speaker must do more to perspicuously indicate her plan to introduce a new object), she must have intended to talk about something new. This supports a discourse expectation—both on the part of the speaker and the other interlocutors—that someone is potentially going to say more about this object, since in discourses new objects aren't continuously introduced without being picked up. So this is something discourse participants potentially need to track, and thus a new discourse referent is added to the context. When we get to the pronoun in (1b), then, it is licensed, because there is a discourse referent in the context for it to pick up.[31] Since subplans constrain the possible next moves, and because interlocutors' subplans have to mesh, introducing a new discourse referent constrains subsequent felicitous discourse moves. For example, since pronouns pick up the most salient discourse referents, after the introduction of the discourse referent for 'a woman' with (1a), interlocutors cannot simply use 'she' to pick up on a different female discourse referent introduced earlier in the discourse without doing extra work.

Plans are not the only things that support expectations. I think it is a general truth about conversation that when it is reasonable to have a high expectation that someone in the conversation will (shortly) go on to say something more about a particular object (conceived as a discourse referent rather than actual referent), an anaphoric pronoun is licensed. Imagine we are in a seminar room together. There are lots of things in the room: students, professors, a table, chairs, blackboards, windows, walls, etc. These are all the ordinary things in the room, and supposing we have just sat down in this room on an ordinary day, with no one having made any gestures, and nothing remarkable having occurred, there is no reason to have a high expectation that

[31] An anonymous reviewer raised the point that I am excluding other potential alternatives such as 'this woman' (used in its specific indefinite sense) and 'a woman I haven't mentioned yet'. To this I can add others, like 'one or more women', 'at least one woman', 'some woman'. If, for example, the alternatives include 'a woman I haven't mentioned yet', then this should equally implicate non-novelty. (This is a version of the symmetry problem.) I take it the alternatives should be constrained to denoting phrases that people actually tend to use in discourse, thus excluding 'a woman I haven't mentioned yet'. Incidentally, I do think that it is a prediction of my view that if we did start to use such phrases in every day conversation, simple indefinites like 'a woman' would no longer indicate a plan to introduce a new discourse referent. However, if my view is right, it is also a prediction that speakers have little reason to start using such clumsy phrases. The rest of the expressions just mentioned are things people tend to say, but there is nothing relevant here to infer from the speaker not using them. That is to say, these are all equally good ways of asserting the existence of a woman who walked in. Not using one of these doesn't indicate anything for present purposes. (I say for present purposes because I am interested in deriving the novelty implicature. There may be other implicatures that can be derived from the subtle differences in meaning between these indefinite expressions, but that is a project for another time.).

for any particular object in the room, someone is about to say something about that object. Now suppose a goat walks in (and this is not an ordinary part of our seminars). Nothing has been said about the goat yet, but it is perfectly acceptable to for someone to say, discourse initially, "It is about to eat your shoe".[32] One way of explaining this, the way I think it should be explained, is that as soon as the goat walks into the room, it is reasonable for each conversational participant to have a high expectation that someone says something about the goat. So a discourse referent for the goat is added to the context, which explains why the pronoun is licensed. It needn't have been something as unusual as a goat. Suppose someone drops their phone with a loud crash. I might felicitously say "please put it away" or "I hope it's not broken".[33] No one *had* to say anything about the goat or the phone — it wasn't guaranteed to be talked about, but it was more likely they would. Similarly, no one has to say anything else about an object introduced by an indefinite, but it is more likely they will. On the other hand, even though the existence of the chairs, blackboards, and windows is entailed by the context, nothing has distinguished them enough to warrant the introduction of the relevant discourse referents. The same goes for objects entailed by what someone asserts, but not explicitly mentioned.

Turning to negation, when someone asserts a negated sentence, the speaker has not asserted the existence of any object at all (quite the opposite); there is no reason to think this is a plan that relates to the objects under discussion. It would be contradictory to think it involves a plan to update a discourse referent, since the discourse referents exist according to the discourse, and for the same reasons it can't be indicative of a plan to add a new one. People don't tend accept that an object doesn't exist and then go on and talk about it as though it does exist, so there's no reason to think it supports discourse expectations to say something more about the non-existent object even in the absence of the speaker's plan. There is no reason to add a discourse referent to the context. Thus, the dynamic pragmatic account predicts widespread external infelicity, and the basic contrast between (2) and (3).

4.3 The licensing question and the semantics of pronouns

The dynamic pragmatics just outlined is intended to be an answer to the licensing question (the dynamic pragmatics will be expanded on shortly to explain how it accounts for the external felicity cases and other recalcitrant cases). It is compatible with different semantics for pronouns. Very often the licensing question is confused with the question of presupposition satisfaction, and so the presuppositions of pronouns are thought to be crucial to the answer to the licensing question. I think these questions are in principle separable, though of course on a view in which pronouns do have presuppositions, presupposition satisfaction will play a role in answering the licensing question. On my view, if I wanted to, I could for example adopt Neale's Russellian d-type theory, in which pronouns have no presuppositions at all, but implement it in a dynamic pragmatics that explains when pronouns are licensed by the context. The central idea is that in this type of account, the explanation for felicitous pronouns is not

[32] This example is based on Stalnaker (1978).

[33] I use examples involving 'it' because 'it', unlike 'he' and 'she', cannot be used as a deictic pronoun.

about presupposition satisfaction or the requirements that the presuppositions of pronouns put on the context, but rather about the fact that pronouns are anaphoric—they look back to information previously introduced in the conversation for their meaning—and that the dynamic pragmatics of context is a way to explain how information flows throughout a discourse, thereby giving an explanation of when anaphora is felicitous and when it is not. This is not the view I ultimately endorse, but I find it worth pointing out that what I say is compatible with such a view, and we can separate the licensing question from the presupposition one, and along with it, its answer. Thus one could adopt the dynamic pragmatics I endorse without adopting the semantics for pronouns that I endorse.

For concreteness, I adopt the following semantics of pronouns. It is beyond the scope of the present paper to argue for it.[34] Pronouns are d-type, and anaphoric descriptions are treated as existential quantifiers that presuppose discourse uniqueness, that is, that there is a unique discourse referent in the context that satisfies the (completed) descriptive material. For example, in (2), the pronoun 'it' goes proxy for the description *Bill's car*; this presupposes that there is a unique car of Bill's under discussion, and is true iff Bill has at least one car that is black. Discourse referents serve two roles. First, they provide the descriptive material for the pronoun. That is, the description that goes proxy for the pronoun is the one that contains all the properties associated with the discourse referent to which the pronoun is resolved. Second, they satisfy the discourse uniqueness presupposition. Formally, this is accomplished using the quantifier domain restriction theory of Stanley and Szabo (2000), with the discourse referent providing the contextually salient individual for the quantifier domain restriction variable (see "Appendix A" for details).

4.4 External felicity explained

Lewis (1979) famously observed that conversations were unlike baseball games because the former but not the latter allow for accommodation. Thomason (1990) points out that conversations are not the only things that allow for accommodation. For example, sometimes acting like the leader in a situation, such as figuring out where a group of conference participants should go to dinner, makes one the leader of the group. I think Thomason is right. Conversations are unlike baseball games in this respect, but much like many other joint activities. Sometimes an individual will merely hint that they want to paint the house red, and then go out to the store and buy red paint as though this had been agreed upon. Whether that is a *felicitous* move to make in the course of a joint activity of painting the house together depends on the details of the activity and its participants. The same goes for conversations, and specifically for acting as though a discourse referent has already been introduced, i.e. for using an antecedent-less pronoun.

I've argued in previous work (Lewis (2012, 2014)) that it should be expected in the kind of pragmatic framework that I am proposing that there are cases of felicitous pronouns without proper antecedents. Conversational participants are not perfect

[34] For a detailed discussion and defense of this view of pronouns (and definite descriptions), see Lewis (ms).

planners, and under certain circumstances, it is natural for plans to change between sentences. One might be planning to speak about something and then change one's mind, or interlocutors might not perfectly agree on what to talk about. Furthermore, although speakers are supposed to make their plans recognizable, sometimes a speaker's plans are not going to be perfectly perspicuous. Sometimes it might take the speaker continuing to speak for what they intend to become clear (this is true whether we are talking about discourse intentions regarding objects under discussion or meaning). In these circumstances, a discourse referent is accommodated, that is, the context is changed so that a discourse referent is added to satisfy the presupposition of the pronoun. Rather than being *ad hoc*, this kind of accommodation is expected in a dynamic pragmatics. Sometimes speakers will say something that raises everyone's expectations that they or someone else will go on to say something more about a particular object under discussion; in this case, a discourse referent is added before we get to any pronoun. But sometimes speakers will make as if this plan was already enacted—when it is clear enough what their intentions are, or when there was a reason to change plans (such as a change in speaker or a pause to think)—and discourse referents are accommodated.

Putting this all together, let's look at Barbara Partee's classic marble example, which is usually taken as support for dynamic semantics. Despite the truth-conditional equivalence between the first sentences of the following discourses, only the first licenses the pronoun in the second sentence.

(36) a. I dropped ten marbles and found all but one.

b. It's probably under the couch.

(37) a. I dropped ten marbles and found nine of them.

b. #It's probably under the couch.

The difference is predicted by the account outlined in the previous subsection. (36a) expresses a proposition that asserts the existence of one marble that is not found (including the object denoting phrase 'one marble'), whereas (37a) merely entails the existence of the unfound marble. Following the reasoning laid out above, it can be inferred from the assertion of (36a) that the speaker wants to talk about the missing marble, thus introducing a discourse referent for it. No such inference can be made from the assertion of (37a). However, it is also often noted that (37) improves dramatically—in fact many judge it perfectly felicitous—if there is a pause in between (37a) and (37b). If dynamic binding is what explains the licensing of a pronoun and how it gets its value, this is mysterious. But if we explain this phenomenon as I have, then the improvement of the second with a pause (or a change of speakers) falls out of the same pragmatic account that explains the difference in felicity. It is odd for a single speaker to change plans mid-sentence, or in between two sentences said in quick succession, expressing a single thought. But pausing to think, or a change in speakers after a pause, can be indicative of a (slight) change in plans. Hence a discourse referent is added by accommodation. Accommodation does not occur every time a presupposition is not satisfied (otherwise nothing would be infelicitous). Accommodation also does not occur whenever a speaker's intentions are recognizable; this is a necessary rather than sufficient condition. Accommodation occurs under constrained circumstances, when attributing a change in plans is reasonable, the intended discourse referent

can be clearly inferred from previous discourse or other aspects of the conversational environment, and often, when it's clear why the circumlocutory route was taken.

Here is another example, from the movie *When Harry Met Sally*:

(38) a. Harry: I'm getting married.
b. Sally: You are? (Long pause in which Harry says "mmhmm"…) *You* are. (Another pause.) Who *is* she?

In this example, it is clear that Harry has a plan to tell Sally about his impending marriage. Marriage is already a topic of conversation at this point, since Sally has just told Harry that she has no interest in marriage at all. But Sally finds Harry extremely annoying and is shocked that anyone would want to marry him, and therefore clearly has a plan to switch the conversation to talk of his fiancée. The long pause indicating Sally thinking, as well as the change in speakers, make the conversational participants (and anyone watching the movie) primed to accept a change in conversational plan. And since it is utterly unambiguous about whom Sally is asking, the pronoun sounds perfectly fine.

Exactly the same phenomenon is going on in the external felicity cases. These cases all have something in common: it is clear what the intended discourse referent is for the pronoun, but moreover it is clear *why* the speaker began with a negated sentence even though this is not usually a good way to introduce a discourse referent, and it's clear *why* the speaker wants to (at least temporarily, in some cases) continue the conversation *as if* there was a specific object under discussion. These are not things that can be attributed solely to the presence of certain linguistic expressions such as modals, negation, or presupposition triggers like 'anymore', though these of course play a role when they are present. Rather it is general principles about what it is rational to take speakers to be doing in these sorts of discourses that explain their felicity, and they do so in a way that is entirely consistent with both the unembedded cases and the other (infelictous) negation cases.

The negation cases of external felicity divide into three (non-mutually exclusive) types that support this picture. First, there are cases in which there is reason to think that the entity in question exists after all, such as (39):

(39) a. Bryan doesn't have an apartment in Paris.
b. He gave it up years ago.

The assertion of (39a) doesn't add a discourse referent for Bryan's former apartment to the context, but when the speaker utters (39b), it is clear that she is talking about the apartment in Paris that Bryan used to have. Since an apartment in Paris was already mentioned, and not having one is compatible with once having had one, the speaker's intentions in (39) are clear and a discourse referent can be accommodated. The discourse is even better when 'anymore' is included in the first sentence, since this straightforwardly presupposes that Bryan used to have an apartment in Paris, priming the conversational participants for accommodation when we get to (39b). This explanation also predicts a contrast between (39) and examples like (40):

(40) a. Bryan doesn't have an apartment in Paris.
b. #It was large.

Nothing about (40) changes the fact that not having an apartment in Paris is compatible with having had one, and largeness is a reasonable property to attribute to Bryan's former apartment. Furthermore, the pronoun still presupposes the existence of a unique discourse referent, so if it was just a matter of accommodation being triggered by presupposition failure, we would expect accommodation here. But what makes accommodation impossible is that nothing about (40) reveals a perspicuous discourse plan, i.e. it doesn't make it clear that the speaker is talking about an apartment that Bryan used to have, either by using 'anymore' or by talking about giving one up, and so the preconditions for accommodation aren't met. (Of course, if other aspects of a conversation in which (40) is embedded makes it clear that the speaker is talking about an apartment that Bryan used to have, the preconditions for accommodation might indeed be met, and this discourse is predicted to be just fine.)

Similar considerations apply to (9), repeated here:

(41) a. Moses doesn't have a daughter who owns a sushi restaurant.
b. It's a pizza restaurant.

Again, once we get to (41b) it's clear what the speaker's plan is. The speaker was using (41a) to deny that Moses has a daughter who owns a *sushi* restaurant, not to deny that he has a daughter who owns a restaurant. Thus the preconditions for accommodation are met—the existence of a restaurant that Moses's daughter owns is easily inferred, and the speaker's reasons for speaking this way are completely clear. It is important to note that we can easily get this reading without focus on 'sushi', so it's not just a matter of the interaction between negation and focus. This is also not a case of predicate negation—the negation is nowhere near the predicate *sushi* in either surface or logical form. Like the other kinds of examples we are going to consider below, examples in this category are generally best when the second sentence is in some way an explanation of why the speaker asserted the negation in the first place.

The second and third cases of external felicity are examples of the same general kind: cases in which there is good reason to talk about a hypothetical entity even though no such entity exists. I posit that there are two conditions under which this is acceptable: when the speaker is explaining why there is no such entity and when the speaker is answering a question under discussion that presupposes the existence of the relevant non-existent entity (these are not mutually exclusive). Consider (4), repeated here:

(42) a. There wasn't a thief here.
b. He would have had to have been magical (to break in without leaving a trace).

Asserting (42a) is not a perspicuous indication of a plan to go on and talk about a thief. As I argued above, no discourse referent is added because the speaker has not even asserted the existence of anything (quite the opposite). But when we get to (42b), there are good reasons to take the speaker to be talking hypothetically about a thief and

therefore add a discourse referent for a thief to the set of discourse referents. One, the speaker is speaking counterfactually by using 'would'; second, the speaker is offering an explanation as to *why she asserted there was no thief*. It is like the speaker is making a little *reductio ad absurdum* argument: Suppose there was a thief; the thief would have had to have been magical; there are no magical thieves; therefore, there was no thief. Third, the natural context for saying (42) in the first place is one in which someone has suggested or asserted that there was a thief, and so the question under discussion is either whether there was a thief or one that presupposes there was a thief (e.g. *were your pens stolen?*). Hence, it is reasonable to be talking hypothetically, as though there was a thief, and attributing to him a property it is conversationally understood that he couldn't have by way of explanation of (42a).

More specifically, what (42b) does is introduce a *derived* or *subordinate* context in the sense of Stalnaker (2014). Derived contexts, for Stalnaker, are separate context sets that are introduced when talking about a set of possibilities that is distinct from the global context set.[35] For example, if one is talking about a particular agent's beliefs, this introduces a derived context set for their belief worlds. Likewise, if one has made a supposition, this introduces a derived context set that is compatible with the supposition. In this case, (42b) makes a counterfactual supposition, and so we get a derived context set including possibilities that are compatible with a thief having been at the place in question. Discourse referents have to exist according to the context set in play, and so discourse referents that are introduced while a derived context set is in play are no longer available after it is no longer in play. The introduction of a discourse referent for a thief here, therefore, does not entail that a thief exists (this would be contradictory), but does entail that a thief has to exist in all the possible worlds in the context set that is currently in play.

This explanation correctly predicts a contrast with examples like:

(43) a. There wasn't a thief here.
b. #He was sneaky.

Here the speaker has no reasonable or recognizable plan; there is no inferrable reason why the speaker is speaking as though there was a thief in uttering (43b). (43b) is not hypothetical, and so does not involve the introduction of a derived context set, nor does it offer any sort of explanation of why the speaker asserted (43a).

It is not merely the presence or absence of the subjunctive 'would' that makes these felicitous, as noted in § 2 above. The account also explains why (44) is bad:

(44) ??There wasn't a thief here and he would have to have been magical to get in without leaving a trace.

It is bad for a speaker to change plans mid-stream, when expressing a single thought without a pause. The first conjunct does not indicate that the speaker has a plan to introduce a discourse referent for a thief, and the "and" indicates that what follows *is a continuation of the same thought*, without any indication that it is an explanation

[35] They are local contexts in some sense of the term, but not in the sense that they are constructed by way of the semantic composition of a sentence.

of why the first conjunct was asserted; hence the anaphora is unlicensed.[36] It also explains why (19)–(21) in § 2 are infelicitous, as there is no recognizable reason to speak hypothetically about the non-existent entity; they neither offer clear explanations as to why there is no such entity nor do they answer a QUD that presupposes the existence of the relevant entity. Of course, if there was more contextual information that made the second sentences of these discourses satisfy these criteria, the present account predicts that they would be felicitous. I take this to be the right result. For example, consider (20), repeated here.

(45) ?There wasn't a visitor today. He would have been happy to be here.

Out of the blue, it seems bad because being happy to be somewhere is usually a reason to visit it, rather than not visit it. However, if the conversation takes place against the background presupposition that everyone in the domain of discourse prefers to be miserable, (45) is much improved.

Examples like (6), repeated here, are cases that are good when the question under discussion presupposes or otherwise strongly suggests the existence of the relevant entity:

(46) a. Mary doesn't have a car.
b. So she doesn't have to park it.

(46) is odd if uttered out of the blue. As I mentioned above, the sort of context that makes (46) sound good, and the kind of context I take it that we are tacitly imagining, is one in which the question under discussion has something to do with Mary parking. For example, this would be an appropriate thing to say if someone had suggested that Mary is late because she is having trouble finding parking.

It is often observed that negated sentences generally require there to be an explicit question under discussion that they address. Psycholinguistic experiments support that negations without a proper context take longer to process than their positive counterparts without a proper context.[37] Often the QUD is positive, with the negated sentence providing a negative answer to it—e.g. *Does Mary have a car?* Other times the QUD is negative — e.g. *Who doesn't have a car?* And other times the QUD presupposes or is otherwise closely related to the positive counterpart to the negative sentence—e.g. *Is Mary having trouble finding parking?* or *How is Mary getting to the city today?* In all the cases except for the negative QUD, the context makes salient the existence of a car that belongs to Mary. This alone isn't sufficient for there to be a discourse referent or to enable accommodation, otherwise we'd see prolific licensing of anaphora on indefinites under the scope of negation. But it does help set up a context that is ripe for accommodation if other factors come in. The other relevant factor here is that the sentence containing the anaphoric pronoun answers

[36] This also explains why (44) with "because" replacing "and" is acceptable.

[37] For example, Horn (1989) cites psycholinguistic research that shows that processing time is longer for negative sentences than positive ones, but not if a proper context of denial is set up. Furthermore, many of the scholars he cites agree that negative sentences can be odd out of the blue when their positive counterparts aren't. Tian and Breheny (2016) argue that "negation is a cue for retrieving the prominent QUD" (27).

the QUD or a sub-QUD (a question that helps answer the QUD). As long as interlocutors continue to address the QUD, a derived context set is created. For example, suppose the QUD is: *Is Mary having trouble finding parking?* (46a) implicates a negative answer to the question, and so it should be removed from the stack of questions. But (46b) continues to address it, and so now we are in a derived context set, one in which the possibility of Mary having a car is still a live option, because we are still addressing whether or not she is having trouble finding parking. The discourse referent for Mary's car is available so long as interlocutors are in this derived context.

Truth-conditionally, however, there is nothing about (46b) that entails the existence of a specific car that Mary doesn't have to park. Recall that I am treating pronouns as restricted existential quantifiers. The complete description is provided by the information associated with the discourse referent. Here the properties are *is a car* and *is owned by Mary*. (46b) is therefore equivalent to 'there does not exist a car that Mary owns that Mary has to park', which intuitively yields the right truth conditions.

As explained above, in both the modal and negation cases, the discourse referents only have a temporary life span—they are only live so long as the modal or negated talk (on the same topic) continues, i.e. as long as the derived context set is in play. This is an improvement over treating it as accommodation in a local context that occurs as part of semantic composition, because it means that long discourses don't have repeated local accommodation. That is, we can explain the felicity of a discourse like (47) without positing that each sentence repeatedly locally accommodates a discourse referent for Mary's car:

(47) a. Mary doesn't have a car.
b. So she doesn't have to park it.
c. She doesn't have to clean it.
d. She doesn't have to buy insurance for it.

The proposed theory also offers a partial explanation of why some, but not all, cases of double negation license pronouns outside its scope. For example, as was observed in § 2, double negation cases like (14) are infelicitous, but some, like (15), are perfectly fine. First of all, it's not clear that this data point, though generally accepted in the literature, is exactly right. Both are actually pretty bad if we're really imagining someone saying them out of the blue. But it is easier to tacitly imagine a natural default context for (15), one in which someone has asserted something to the effect that John doesn't own a car, and the speaker of (15) is denying that. The present theory can explain why anaphora in double negation cases is often infelicitous; out of the blue, a double negation is a non-perspicuous way of revealing a plan to go on and talk about some object or other (out of the blue, it is not clear what doubly negated locutions are doing at all!). But in a context in which it is clear why the speaker uses the double negation, the locution is a good indication of such a plan, and so conversational participants are right to add a discourse referent. If Alfred has just denied that Bob saw anyone walking in the park, Bob can felicitously say (48).

(48) It is not the case that there is not a man walking in the park. He is whistling.

A potential objection to the view I am defending in this section is one that is commonly raised against neo-Gricean accounts in general, and that is about its ability to make predictions about new cases. What are the independent criteria by which a speaker's plan is judged to be perspicuous enough to introduce a discourse referent, either by way of saying something containing a potential antecedent or by way of saying something containing an antecedentless pronoun? For example, what does the present theory predict about indefinites in other non-upward monotonic positions, such as under the scope of 'fewer than two'?

(49) a. There were fewer than two students who chose to read a book on dynamic semantics.

b. ??It was Chierchia's book.[38]

Recognizing a speaker's discourse plan to introduce a new discourse referent by way of reasoning based on the maxim of quantity is triggered in the first place by an existential assertion, which (49a) is not. So my theory clearly predicts that (49a) does not introduce a novel discourse referent.[39] What about the accommodation of a discourse referent for the pronoun? Why can we not infer that exactly one student chose to read a book on dynamic semantics, and 'it' is supposed to refer to that book? I have posited several conditions under which accommodation naturally occurs: that the sentence containing the pronoun explains why the negated sentence was said in the first place or it continues answering a question under discussion. It is very hard to imagine a context in which this is what the speaker of (49b) is doing. Furthermore, taking the model of discourses like (39) and (41), accommodation of a discourse referent for a pronoun can occur when it's clear the speaker was denying some property of an object that does exist, rather than the existence of the object (even if they used sentential negation). This also does not apply here. So my account predicts the (general) infelicity of (49b). Of course, my view *also* predicts that we should expect some variation. In a specific context in which there is a clear reason to use the 'fewer than two' locution even though the second sentence reveals that the speaker knows there was exactly one, then my theory predicts that (49) should be much improved, if not perfectly felicitous. This is supported by the fact that (49) is much improved if we connect the two sentences with 'because':

(50) There were fewer than two students who chose to read a book on dynamic semantics because it was Chierchia's book (and Chierchia's book is challenging).

[38] Thank you to an anonymous reviewer for raising this point and for this example.

[39] We shouldn't forget that there are some contexts in which contextual factors make it such that existence does not need to be asserted and a discourse referent will be established anyway because the reason for a speaker taking a non-standard route to introducing a new object under discussion will be clear. These will be highly contextual, particularized pragmatic derivations. So in a specific context that warrants saying something like 'fewer than two' when the speaker clearly means 'exactly one', a discourse referent will be introduced.

Because the proposed theory requires looking at the details of the context, sometimes intuitions on discourses taken out of context will be muddy. It also predicts that sometimes the very same sequence of sentences will be felicitous and other times infelicitous. Finally, since perspicuity is a spectrum rather than an on/off phenomenon, this grounds the fact that felicity is a spectrum rather than an on/off phenomenon, which I take to be the right prediction. (One piece of evidence for this is the range of judgments of infelicity that are represented in the literature by single question marks, multiple question marks, and hashes.)

5 Internal dynamics

Dever (2013) raises doubts about the plausibility of dynamic pragmatics, particularly in its ability to account for local contexts. Regarding the dynamic semantic treatment of negation, he writes:

> To update with $\neg A$, we locally/internally update with A, for the purposes of determining what should be eliminated from the prior context. The local context $\sigma \uparrow A$ doesn't survive outside the negation, but can have observable effects inside the negation (such as licensing anaphora within the negation on an indefinite also within the negation—the mere locality of the local context then explains the *unavailability* of anaphora outside the negation). (114)

The ability to account for local contexts is a potential challenge to my account, or any dynamic pragmatic account. I take this challenge seriously; dynamic semantics has no problem with sentence-internal dynamics since sub-sentential expressions also have semantic values that are context change potentials. Sub-sentential dynamics is therefore neatly accounted for by exactly the same mechanisms as cross-sentential dynamics. On the other hand, the account I have given in terms of pragmatic reasoning and discourse plans has to say something about how interlocutors reason about a speaker's discourse plan when it comes to unasserted, sub-sentential expressions.

Despite the fact that this is an important challenge that needs to be met by a dynamic pragmatic account generally, I don't think it's a challenge that needs to be met here, when dealing with negation. It is my contention that there are no internal dynamics that are specific to negation. That is, there is no need to appeal to local contexts to explain what goes on within the scope of a negation operator, and *contra* Dever, there is no licensing of anaphora "within the negation on an indefinite also within the negation". Therefore, the ability of a dynamic pragmatics to account for local contexts or sub-sentential dynamics is beyond the scope of the present work.

To see why I propose that there is no sentence internal dynamics when it comes to the negation operator, begin by noting the following contrast:

(51) a. It is not the case that a student came and asked for her grade.
b. # It is not the case that a student came and she asked for her grade.

This suggests that on the most salient reading of sentential negation, it does not take wide scope over sentence (CP) conjunction, as it would have to get a felicitous reading of (51b). This means that (51a) is a case of verb phrase (VP) conjunction, which can

be treated as an ordinary case of a bound pronoun. That is to say, (51a) is just like 'A student asked for her grade', 'no student asked for her grade', or 'every student asked for her grade', which are textbook examples of bound pronouns.[40]

There are contexts in which (51b) is felicitous, but they also do not require internal dynamics. (51b) is felicitous on an echoic denial use, as in when someone has just said "A student came in and she asked for her grade" and the speaker wants to deny it. Echoic denial usually allows for the grammatical features of the sentence it is denying (e.g. it still takes positive polarity items). I also assume that it inherits the semantic interpretation of the sentence it denies, including the resolution of context-sensitive elements; that is, the pronoun in the scope of the denial is interpreted in the exact same way as in the asserted positive sentence and its meaning is parasitic on the asserted sentence. Again, the the felicity of the pronoun is not explained by the introduction

[40] Comparing a syntactic tree for a standard bound pronoun sentence 'a student asked for her grade' with 'it's not the case that a student came and asked for her grade' shows that the pronoun is c-commanded by the quantifier 'a student' in both cases, which means it can syntactically bind it. A node A c-commands a node B in a syntactic tree iff the lowest branching node dominating A also dominates B and A does not dominate B nor does B dominate A. In these examples, the pronoun 'her' is c-commanded by the DP 'a student' in both trees. The tree with VP conjunction could also be written as a non-binary branching tree, depending on the syntactic theory to which one ascribes; this makes no difference to the c-command relationship of DP and pronoun.

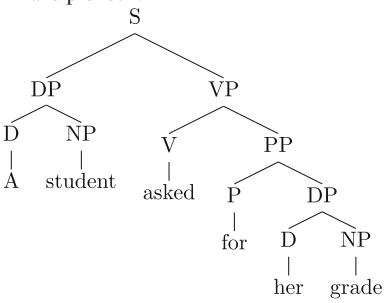

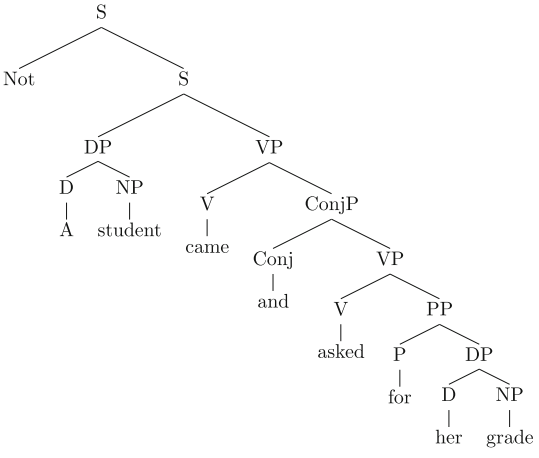

of a local context.[41] Finally, I think there is another condition under which (51b) is felicitous, and that is in exactly those situations in which the anaphora is globally licensed. For example, in contexts in which it is clear that the denial is being targeted at the predicates rather than the entire sentence, it is perfectly fine:

(52) It is not the case that a student came in to my office and she asked for her grade. Rather, she called me on my personal cell phone and she *demanded* to know her grade.

This are analogous on my view to cases like (39) and (41).

To see that this is not just a matter of the locution 'it is not the case that', consider the following parallel contrasts:

(53) a. No student came and asked for her grade.
b. #No student came and she asked for her grade.

(54) a. A student didn't come in and ask for her grade.
b. #A student didn't come in and she asked for her grade.

(55) a. I doubt that a student was here and left her essay.
b. #I doubt that a student was here and she left her essay.

(56) a. There's no way a student was here and left her essay.
b. #There's no way a student was here and she left her essay.[42]

Given the way in which dynamic semantics defines negation and given the fact that (51b) is grammatically possible, it is actually a surprising result from the point of view of dynamic semantics that there is not a salient reading of (51b) that is equivalent to (51a), since the pronoun should be able to be easily resolved in the local context created by negation. On the other hand, static semantics correctly predicts the possible readings.

6 Conclusion

Dynamic semantics by design gives a neat, simple answer to the licensing question. I've argued that the data involving negation is much less neat than first meets the eye, and is better explained by a dynamic pragmatic account. Of course, this is not meant to be a decisive argument against dynamic semantics in general. It is an argument against dynamic semantics as an answer to the licensing question. If it turns out that dynamic semantics is needed for other reasons, then the natural view would be that we need a dynamic semantics and something like the dynamic pragmatics proposed in this paper. The upshot, however, is that not only *can* a dynamic pragmatics provide an answer to the licensing question and integrate with a static semantics, but if I am right, this is the way we *should* be answering the question.

[41] See Geurts (1998) and van der Sandt and Maier (2003) for more on denial.

[42] For the 'I doubt' and 'There's no way' locutions the echoic denial reading is more salient. To get the infelicitous reading, the reader must imagine a case in which this is not being uttered as an echoic denial.

Acknowledgements Thanks are due to Josh Dever, Michael Glanzberg, Jim Pryor, Jessica Rett, and Anders Schoubye for helpful comments and discussion. Thanks as well to audiences at the *Philosophy of Language and Linguistics Conference* in Dubrovnik, Croatia in September 2014, the *Workshop on Semantic Content* in Barcelona, Spain in November 2014, and the *Dartmouth Philosophy of Language Workshop* at Dartmouth College in September 2015, and members of the NYU *Mind and Language* seminar in February 2018 for questions and discussion, as well as two anonymous reviewers for this journal for their comments, all of which helped improve the paper vastly.

A formal implementation

In this appendix, I briefly outline the formal implementation of both context and the semantics of pronouns, showing how these work together to account for basic cases of discourse anaphora.

1. **Context:** A context C contains the following elements:

- DR, the set of discourse referents, modeled as indices. These indices are the domain of the assignment functions.
- WG, a set of world/assignment function pairs ⟨w, g⟩ such that w ∈ W, where W is the set of all possible worlds, and g is one of the possible assignments of indices in DR to entities in w, for each possible such assignment g.[43]
- CG, the common ground, modeled a context set, i.e.{w ∈ W | w is possible given conversational presuppositions}
- QUD, the set of questions under discussion.

2. **Syntax:** The syntax of the language includes variables x, y, z, with or without subscripts, numerical indices, logical constants (proper names), predicates (from English), the logical connectives ∧, ∨, ¬, the indefinite article *a*, the definite article *the*, the generalized quantifiers *every*, *some*, *most*, *few*, *no*.

3. **Semantics:**

- A model M = ⟨W, D_e, D_t, I⟩ where W = set of worlds w, D_e = domain of entities e, D_t = {0,1}, and I = interpretation function
- For each constant c, I(c) = some e ∈ D_e
- For n-place predicates p, I(p) = {⟨w,{⟨e_1,…e_n⟩, ⟨e_1,…e_n⟩…}⟩, ⟨w,{⟨e_1,…e_n⟩, ⟨e_1,…e_n⟩…}⟩…}, such that w ∈ W and e_1,…e_n ∈ D_e, and the second member of the n-tuple is a set of entities if p is a 1-place predicate and a set of n-tuples of entities otherwise.
- The logical connectives have their ordinary static denotations.[44]
- The quantifiers have standard denotations from generalized quantifier theory, e.g.:

a. $[\![\text{[every x: } \phi](\psi)]\!]^{M,w,c,h} = 1$ iff ∀e in D_e s.t. $[\![\phi]\!]^{h[x\to e]} = 1$ at w, $[\![\psi]\!]^{h[x\to e]} = 1$ at w

[43] I use the set of all worlds rather than the worlds in the common ground because the information the assignment functions encode is the properties associated with the discourse referents; this has nothing to do with which worlds are still considered open according to the conversation. Keeping these separate is helpful in maintaining the ordinary notion of truth at a world in the system.

[44] I don't discuss conditionals in this fragment for the sake of simplicity.

b. $[\![$[some x: $\phi](\psi)]\!]^{M,w,c,h} = 1$ iff $\exists$e in D_e s.t. $[\![\phi]\!]^{h[x \to e]} = 1$ at w & $[\![\psi]\!]^{h[x \to e]} = 1$ at w

- The assignment functions relative to which denotations are calculated are distinct from the ones in WG. The assignment functions in WG are functions from indices to entities, the other assignment functions are functions from variables to entities or indices. I reserve 'g' for the former and 'h' for the latter.
- *An* has the same denotation as the existential quantifier *some*.
- Pronouns are equivalent to definite descriptions with null overt material, and the definite article is treated as a generalized quantifier that presupposes discourse uniqueness:

$$[\![\text{[The x:}\phi](\psi)]\!]^{M,w,c,h} = \begin{cases} \text{defined if } \exists!\text{n} \in \text{DR s.t. for all } \langle \text{w,g} \rangle \in \text{WG}, [\![\phi]\!]^{h[x \to g(n)]} = 1 \text{ at w} \\ 1 \text{ iff } \exists \text{e in } D_e \text{ s.t. } [\![\phi]\!]^{h[x \to e]} = 1 \text{ at w \& } [\![\psi]\!]^{h[x \to e]} = 1 \text{ at w} \\ 0 \text{ otherwise} \end{cases}$$

The clause defines the discourse uniqueness presupposition as the requirement that there is exactly one index *n* in DR such that the (possibly complex, massively conjoined) property associated with that index is exactly the same as the property denoted by the descriptive material in the definite description. The presupposition is a condition on definedness. The relevant descriptive material on the definite description that must satisfy discourse uniqueness is the completed description. The completed description occurs via the mechanism of quantifier domain restriction. If defined, *The F is G* is true iff there is at least one F that is G.

I treat quantifier domain restriction following Stanley and Szabo (2000). All quantifiers are restricted by a variable that shares a node with the noun at the level of syntax. The variable is of the form *f(i)*, where *i* is an individual variable that can either be bound or get a value from the context and *f* is a contextually determined function from individuals to quantifier domains, which in the intensional version is a property, a function from worlds to sets of individuals. The domain of the quantifier is determined by combining the denotation of the overt predicate (if any) and the result of *f(i)*. (In Stanley and Szabo's extensional version, they combine by set intersection. In the intensional version, we can think of them combining by something like looks exactly like predicate modification, except for the fact that the noun and variables share the same node.) For example, take the sentence 'every student was happy' in a particular context of use, formally written as [Every x: $\langle$student, f(i)$\rangle$ x] (happy x). In a particular context, f might be the function from events to a function from worlds to their participants at that world, and the value for i might be the 2019 Met Gala. Hence, *f(i)* yields a function from worlds to the set of all the participants of the 2019 Met Gala at that world, and combining that with the denotation of *students* (i.e. a function from worlds to students at that world), we get *students at the 2019 Met Gala* as the restrictor property. Thus we get the restricted universal claim that every student at the 2019 Met Gala was happy.

As I said, pronouns are treated as definite descriptions with null overt material, and so the descriptive material comes entirely from the quantifier domain restriction.[45] Let's see how this applies to our central example:

(57) a. A woman walked in.
a'. [An x: woman x] (walked in x)[46]
b. She sat down.
b'. [The x : $\langle \emptyset$,f(i)$\rangle x$] (sat down x)

Introduction and update of discourse referents are modeled as change to DR and WG. (57a) adds a novel index, say *1*, to DR, and changes WG such that the new set contains all the *1*-variants of the input assignment functions, such that *1* is now assigned to an individual in I(*woman*) and I(*walked.in*), relative to its paired world. An *1*-variant of an assignment function g is here defined as all the possible extensions of g that assign *1* to something in the relevant interpretation.

Consider the following toy model:

1. W = $\{w_1, w_2, w_3\}$
2. D_e = {Alice, Bob, Carol, David, Emily, Francine}
3. I(*woman*) = $\{\langle w_1$, {Alice, Carol, Emily, Francine}$\rangle$, $\langle w_2$, {Alice, Carol, Emily, Francine}$\rangle$, $\langle w_3$, {Alice, Carol, Emily, Francine}$\rangle\}$
4. I(*walked in*) = $\{\langle w_1$, {Alice, Carol, Bob}$\rangle$, $\langle w_2$, {Bob, Carol, David, Emily}$\rangle$, $\langle w_3$, {Alice, Emily, Francine}$\rangle\}$
5. I(*sat down*) = $\{\langle w_1$, {Alice, Bob, Francine}$\rangle$, $\langle w_2$, {Carol, Emily}$\rangle$, $\langle w_3$, {Francine}$\rangle\}$

For simplicity assume a initial context in which CG = W and DR = { }. I suppress the QUD in what follows. The pragmatic effect on the context of asserting (57a) is to add a novel discourse referent for a woman who walked in, and to eliminate all worlds in CG that are incompatible with a woman walking in:

Context after (57**a**) **is asserted:**

1. CG = W ∩ $[\![$a woman walked in$]\!]^{M,c,h}$ = W_1
2. DR = {1}
3. WG = $\{\langle w_1, g_1$: 1→ Alice$\rangle$, $\langle w_1, g_2$: 1→ Carol$\rangle$, $\langle w_2, g_3$: 1→ Carol$\rangle$, $\langle w_2, g_4$: 1→ Emily$\rangle$, $\langle w_3, g_5$: 1→ Alice$\rangle$, $\langle w_3, g_6$: 1→ Emily$\rangle$, $\langle w_3, g_7$: 1→ Francine$\rangle\}$

The discourse referent serves as the value for the salient individual i in (57b), and the QDR function that goes in for the variable f gathers all the information associated with that discourse referent, yielding the set of all possible individuals who could be a witness for the discourse referent relative to each possible world:

[45] I am suppressing treatment of gender and number in pronouns, but that could be easily incorporated.

[46] I haven't included a QDR variable on 'a woman' for the sake of perspicuity, though I think all quantifiers come with the QDR variable.

Contextual QDR function for anaphoric pronouns:
For some index n, $f(n) = \{\langle w, \{e_1 \ldots e_n\}\rangle \mid$ w is the first member of a tuple in WG and $\{e_1 \ldots e_n\}$ is the set of all entities assigned to n by some g which is the second member of a tuple that has w as its first member$\}$

This will yield a function from worlds w to the set of all women who walked in at w. (57b) is then true at a world iff there is at least one object in that domain of possible witnesses that sat down, i.e. it is true iff there is at least one woman who walked in and sat down.

Context after (57b) is asserted:

1. CG = $W_1 \cap [\![\text{some woman walked in and sat down}]\!]^{M,c,h}$
2. DR = $\{1\}$
3. WG = $\{\langle w_1, g_1\colon 1 \rightarrow \text{Alice}\rangle, \langle w_2, g_3\colon 1 \rightarrow \text{Carol}\rangle, \langle w_2, g_4\colon 1 \rightarrow \text{Emily}\rangle, \langle w_3, g_7\colon 1 \rightarrow \text{Francine}\rangle\}$

The effect here is again purely pragmatically motivated. Eliminating incompatible assignment functions captures the fact that conversational participants use discourse referents to track what properties hang together as satisfied by a single witness according to the discourse.

For the cases that involve accommodation, such as the external felicity cases central to this paper, the relevant discourse referent is accommodated in the context before the semantic machinery does its work. For example, take (39): when we get to the pronoun 'it', there is no unique discourse referent to which it can be resolved. So a discourse referent for Bryan's Paris apartment is added to the context, and then semantics works in the same way as described.

References

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

Brasoveanu, A. (2010). Structured anaphora to quantifier domains. *Information and Computation*, *208*(5), 450–473.

Bratman, M. E. (1999). *Faces of intention*. Cambridge: Cambridge University Press.

Chierchia, G. (1995). *Dynamics of meaning*. Chicago: University of Chicago Press.

Cooper, R. (1979). The interpretation of pronouns. In F. Henry & H. Schnelle (Eds.), *Selections from the third groningen round table syntax and semantics* (pp. 61–92). Cambridge: Academic Press.

Davies, M. (1981). *Meaning, quantification, necessity*. Abingdon: Routledge and Keegan Paul.

Dever, J. (2013). The revenge of the semantics-pragmatics distinction. *Philosophical Perspectives*, *27*, 104–144.

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

Elbourne, P. (2013). *Definite descriptions*. Oxford: Oxford University Press.

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

Frank, A., & Kamp, H. (1997). On context dependence in modal constructions. In *Proceedings of salt VII* (pp. 151–168).

Gauker, C. (2008). Against accommodation: Heim, van der Sandt, and the presupposition projection problem. *Philosophical Perspectives*, *22*, 171–205.

Geurts, B. (1998). The mechanisms of denial. *Language*, *74*(2), 274–307.

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

Grice, P. (1989). *Studies in the way of words*. Cambridge: Harvard University Press.

Groenendijk, J., & Stokhof, M. (1990). Dynamic montague grammar. In L. Kalman & L. Polos (Eds.), *Proceedings of the 2nd symposium on logic and language* (pp. 3–48). Budapest: Eotvos Lorand University Press.
Groenendijk, J., & Stokhof, M. (1991). Dynamic predicate logic. *Linguistics and Philosophy*, *14*, 39–100.
Groenendijk, J., Stokhof, M., & Veltman, F. (1996). Coreference and modality. In *The handbook of contemporary semantic theory*, (pp. 179–216). Black-well.
Heim, I. (1982). *The semantics of definite and indefinite noun phrases*. University of Massachusetts dissertation.
Heim, I. (1983). On the projection problem for presuppositions. In *Second annual west coast conference on formal linguistics*.
Heim, I. (1990). E-type pronouns and donkey anaphora. *Linguistics and Philosophy*, *13*, 137–177.
Hirschberg, J. (1985). *A theory of scalar implicature (natural langauges, pragmatics, inference)*. UPenn dissertation.
Horn, L. (1989). *A natural history of negation*. Chicago: University of Chicago Press.
Jankovic, M. (2014). Communication and shared information. *Philosophical Studies*, *169*(3), 489–508.
Kamp, H. (1981). A theory of truth and semantic representation. In J. Groenen-dojk, T. H. Janssen & M. Stokhof (Eds.), *Formal methods in the study of language*, Vol 1, (pp. 277–322). Amsterdam: Mathematisch Centrum.
Kamp, H., & Reyle, U. (1993). *From discourse to logic*. Berlin: Kluwer Academic Publishers.
Karttunen, L. (1976). Discourse referents. In J.D. McCawley (Ed.), Syntax and semantics 7: notes from the linguistic underground, (pp. 363–386). Academic Press.
Keshet, E. (2018). Dynamic update anaphora logic: A simple analysis of complex anaphora. *Journal of Semantics*, *35*, 263–303.
King, J. (1987). Pronouns, descriptions and the semantics of discourse. *Philosophical Studies*, *51*, 341–363.
King, J. (1991). Instantial terms, anaphora and arbitrary objects. *Philosophical Studies*, *61*, 239–265.
King, J. (1994). Anaphora and operators. *Philosophical Perspectives*, *8*, 221–250.
Kurafuji, T. (1998). Dynamic binding and the E-type strategy: Evidence from Japanese. In D. Strolovitch & A. Lawson (eds.), *Proceedings of SALT VIII* (pp. 129–144).
Kurafuji, T. (1999). *Japanese pronouns in dynamic semantics*. Rutgers University dissertation.
Lewis, D. (1979). Scorekeeping in a language game. *Journal of Philosophical Logic*, *8*(3), 339–359.
Lewis, K. S. (2012). Discourse dynamics, pragmatics, and indefinites. *Philosophical Studies*, *158*, 313–342.
Lewis, K. S. (2014). Do we need dynamic semantics? In A. Burgess & B. Sherman (Eds.), *Metasemantics: new essays on the foundations of meaning, chap. 9* (pp. 231–258). Oxford: Oxford University Press.
Lewis, K. S. (2017). Dynamic semantics. *Oxford handbooks online*. Retrieved July 18, 2020, from https://www.oxfordhandbooks.com/view/10.1093/oxfordhb/9780199935314.001.0001/oxfordhb-9780199935314-e-14.
Lewis, K. S. (ms). Descriptions, pronouns, and uniqueness. Unpublished manuscript.
Muskens, R. (1991). Anaphora and the logic of change. In J. van Eijck (Ed.), *Logics in AI* (pp. 412–428). Berlin: Springer.
Neale, S. (1990). *Descriptions*. Cambridge: MIT Press.
Roberts, C. (1996). Anaphora in intensional contexts. In S. Lappin (Ed.), *The handbook of contemporary semantic theory*. Hoboken: Blackwell.
Roberts, C. (2003). Uniqueness in definite noun phrases. *Linguistics and Philosophy*, *26*, 287–350.
Roberts, C. (2004). Context in dynamic interpretation. In L. Horn & G. Ward (Eds.), *Handbook of pragmatics*, Chap 9, (pp. 197–220) Blackwell.
Roberts, C. (2012). Information structure: Towards an intergrated formal theory of pragmatics. *Semantics and Pragmatics*, *5*(6), 1–69.
Russell, B. (1905). On denoting. *Mind*, *14*(56), 479–493.
Stalnaker, R. (1978). Assertion. *Syntax and Semantics*, *9*, 315–332.
Stalnaker, R. (2014). *Context*. Oxford: Oxford University Press.
Stanley, J., & Szabo, Z. (2000). On quantifier domain restriction. *Mind & Language*, *15*(2–3), 219–261.
Thomason, R. (1990). Accommodation, meaning, and implicature: Interdisciplinary foundations for pragmatics. In P. R. Cohen, J. Morgan, & M. E. Pollack (Eds.), *Intentions in communication*. Cambridge: MIT Press.
Tian, Y., & Breheny, R. (2016). Dynamic pragmatic view of negation processing. In P. Larrivée & C. Lee (Eds.), *Negation and polarity: Experimental perspectives language, cognition, and mind*. Berlin: Springer.

van den Berg, M. H. (1996). *Some aspects of the internal stucture of discourse. The dynamics of nominal anaphora*. University of Amsterdam ILLC dissertation.
van der Sandt, R. (1992). Presupposition projection as anaphora resolution. *Journal of Semantics*, *9*, 333–377.
van der Sandt, R., & Maier, E. (2003). Denials in discourse. Unpublished manuscript.
Veltman, F. (1996). Defaults in update semantics. *Journal of Philosophical Logic*, *25*, 221–261.
Yalcin, S. (2012). Dynamic semantics. In D. Fara & G. Russell (Eds.), *The Routledge companion to philosophy of language*, Chap 2.7, (pp. 253–279). Routledge.