Philos Stud (2012) 158:313–342
DOI 10.1007/s11098-012-9882-y

Discourse dynamics, pragmatics, and indefinites

Karen S. Lewis

Published online: 24 March 2012

1 Introduction

Everyone agrees that conversations take place in a context. This is not to merely point out that conversations occur at a time and a place, or that there are particular speakers and hearers, though this is all of course true. Conversations take place against a background of mutually recognized facts: facts about the beliefs and presumptions of the participants, facts about the information conveyed thus far, facts about what's under discussion, and so on. Although there is some disagreement as to the exact nature of conversational contexts, it is generally agreed that they record these sorts of facts.

As Stalnaker (1978) and Lewis (1979) first pointed out, the context both affects and is affected by the utterances in a conversation. Consider a conversation in which I ask (1) and another one in which I ask (2):

(1) Do you know of a nice, flat road for bike riding?
(2) I need a large, flat surface for a physics experiment. Do you know of any?

In the first case it is perfectly acceptable to assert (3) in reply, but in the second case it is unacceptable:

(3) 4th Avenue is flat.

Uttering (1) or (2) affects (among other things) the standards of precision in the conversational context. What counts as flat for bike-riding is a lot more forgiving

K. S. Lewis (✉)
Department of Philosophy, University of Southern California, Los Angeles, CA, USA
e-mail: karenlew@usc.edu

than what counts as flat for a physics experiment. In turn, the standards of precision in the context affect the acceptability of (3).[1]

In this way, discourses are *dynamic* things.[2] An utterance affects the context in various ways. For example, the context might be updated with the information conveyed, a new topic of discussion, an increase in the standards of precision, or the fact that the particular utterance occurred. In turn, the context constrains whether an utterance is acceptable or even true.

Over the last 30 years, the dynamic nature of discourse has thrown into question whether traditional truth-conditional semantics is the right approach to a theory of meaning. A number of philosophers and linguists have argued that we should abandon the traditional notion of content and adopt one that acts directly on contexts. This approach, called *dynamic semantics*, takes the semantic value of sentences to be their *context change potentials* (CCPs) and sub-sentential expressions their contribution to the CCP of the whole. Technically, a CCP is a function from contexts to contexts. Intuitively, it is like an instruction or a recipe for updating the context. By contrast, traditional *static* semantics generally takes the semantic content of a sentence to be its truth conditions, and the sub-sentential expressions their contribution to the truth conditions of the whole.[3]

The central difference between static and dynamic semantics comes down to whether certain changes to the conversational context are (at least in part) dictated by the semantics, or whether they are explained by pragmatics.[4] Dynamic contents encode (some of) the effects of an utterance on an arbitrary input context. By contrast, static contents do not encode any updates to the context. On a static view, the effect(s) of content on the context has to be explained pragmatically. *These are fundamentally different sorts of explanations*. Semantics describes facts about natural language. Pragmatics, on the other hand, describes facts about rational agents who engage in co-operative activities.

The sort of updates to the context that are central to the dynamic view are systematic or robust ones: updates that are observed across contexts and languages. Dynamic semanticists claim that we need a semantic theory to adequately account for these sorts of updates, particularly in cases in which the updates are not merely a matter of recording the truth-conditional information conveyed by an utterance. As just stated, the context records not only the information conveyed by utterances in the conversation, but information about the conversation itself: what topic is under discussion, etc. Call this *discourse information*. Dynamic semantics generally holds that pragmatic accounts cannot capture updates regarding discourse information. I take the reasoning behind this to be something like the following. If static contents

[1] Whether the change in context affects the truth or merely the acceptability of (3) depends one one's view of predicates like *flat*.

[2] It's a little weird to think of untokened discourses as dynamic; it is more precise to think of tokened discourses as dynamic and untokened discourses as things that would be dynamic if they were tokened.

[3] This is not an entirely accurate characterization of static semantics, since static semantics includes views in which contents encode more than merely truth conditions, such as structured views of content. The point is that static contents are not things that act directly on contexts.

[4] See chapter 1 of Lewis (2011) for arguments for this point.

of utterances are truth conditions, and pragmatics standardly acts on contents, then it seems all we can ever yield is updates about truth conditions, and not about other facts such as discourse information.

The view defended in this paper shows these claims about the limits of pragmatic explanations to be false. I will examine the case of indefinite expressions—expressions of the form *a(n) x* or *some x*—which played a central role in motivating the move to dynamic semantics. In looking at this case, I will show that pragmatics not only does have the power to explain systematic, robust updates involving discourse information, but that a pragmatic explanation is the right sort of explanation of the phenomena in question.[5] The important result is that we are able to maintain a traditional static notion of content and account for the relevant updates to the context by appealing to pragmatics. This also opens up new strategies for accounting for a variety of linguistic data.

While the present paper is not intended to be the final word on the debate between static and dynamic semantics (since it is my view that such an adjudication can only happen by examining all the relevant linguistic data), it does undermine some important arguments for dynamic semantics. Specifically, it has become a trend in the recent literature to move directly from the observation that a certain linguistic construction has a systematic interaction with the conversational context (which we will see shortly in the case of indefinites) to the claim that we need to adopt a dynamic semantics to capture the data. I will argue that in the case of indefinites, a pragmatic view not only can account for the data, but in several important instances accounts for it better than dynamic semantic accounts. This suggests that the quick move from the linguistic data to a dynamic semantics is not warranted, and that other arguments for dynamic semantics (evidence from conditionals, modals, presuppositions, and so on) should be similarly re-examined.

In § 2, I identify two features associated with indefinites in discourse, *novelty* and *licensing*. I briefly review one family of dynamic views that semantically encodes these discourse properties. These views are particularly attractive because the CCP (i.e. the dynamic semantic value) of the indefinite simultaneously accounts for novelty and licensing. I then argue that these accounts cannot be right, because novelty is not a semantic feature of indefinites.[6] In § 3 I motivate a neo-Gricean pragmatic picture in which updating the conversational context is a matter of recognizing speakers' plans for the discourse. I maintain a static, truth-conditional notion of content and the standard semantics for indefinite expressions, and argue that novelty and licensing are uniformly explained by this pragmatic view involving planning. § 4 discusses some further merits of my pragmatic account over an account that semantically encodes the discourse properties of indefinites. Those readers who are interested in some formal comparisons between dynamic semantics and the sort of account I propose should consult the Appendix.

[5] While the present paper concentrates on basic, unembedded cases of indefinites, in other work (Lewis (2011)) I show how my account can be extended to cases of indefinites embedded under operators and quantifiers.

[6] This is not to argue that no dynamic semantic account can be right, but that this family of accounts cannot be right.

2 Novelty and licensing

The received tradition is to treat the indefinite as a classical existential quantifier. However, uttering a sentence containing an indefinite in a discourse generally has a greater effect than that of making a mere existential claim. For example, consider the following short discourse:

(4) a. A woman walked in.
b. She ordered lunch.

Many people have noted that when a speaker utters something like (4a), it is understood that she is using *a woman* to talk about something new according to the conversation. Call this feature of indefinites *novelty*. This contrasts with definite expressions like *the woman* and *she*, which are generally thought to continue discussing something familiar to the conversation. It is important to note that novelty is not a matter of reference or denotation; no one claims that the object in the world that actually satisfies the indefinite description has to be new to the conversation. Novelty is the claim that, roughly, a speaker is talking about something that is novel for the purposes of the conversation. Novelty is unexplained by traditional semantics, since there is nothing in the semantic value of an existential quantifier that encodes whether it describes something novel or familiar to the conversation.

Secondly, (4a) introduces a woman under discussion that can be picked up by future anaphoric expressions, like *she*. Call this feature *licensing*. It is not at all obvious how to explain in what way *a woman* licenses the anaphora in (4b). Since *a woman* is not a referring expression, we cannot simply say that *a woman* and *she* refer to the same object in the world. (As, for example, we might be tempted to say if we replaced *a woman* in (4a) with a proper name like *Michelle Obama*.) Neither can we say that *a woman* binds the pronoun, since the pronoun lies outside of its syntactic scope.

2.1 Dynamic semantic accounts

As previously noted, one reaction to the discourse properties possessed by indefinites (i.e. novelty and licensing) has been to develop a semantics that encodes these properties. This is the strategy of many dynamic semanticists, such as Kamp (1981), Heim (1982), Groenendijk and Stokhof (1991), Kamp and Reyle (1993), Muskens (1996), and Asher and Lascarides (2003). On these views, the CCP of an indefinite description dictates that a novel representation of an object under discussion should be added to the context. This representation then provides a value for subsequent anaphoric pronouns.[7]

Objects under discussion are represented in the context by discourse-level entities, i.e. representations that are neither linguistic expressions nor objects in the world (or in a model). I adopt Heim's notion of a *file card* when describing the object that plays this role.[8] A conversational participant has to keep track of the information conveyed about

[7] This description glosses over considerable formal differences between the views. These differences are not relevant for the present project.

[8] I take file cards to be equivalent to other concepts that do the same work, such as discourse referents or discourse entities. I prefer the file card metaphor because the terminology involves no temptation to confuse file cards with real objects in the world, contrary to the unfortunate term *discourse referent*.

the objects under discussion, making sure to bundle together bits of information thought to be about the same object. But this bundling of information doesn't commit the conversational participants to talking about a *particular* object in the world. For example, the fact that (4) is about a woman who ordered lunch doesn't commit the conversation to being about a particular woman who ordered lunch. In some cases, there may even be no fact of the matter which woman discourse (4) is about. But it is clear that according to the conversation, the properties of being a woman, walking in, and ordering lunch should be bundled together. The discourse does not accurately represent the world if there turns out not to be a single object that shares these properties. And the conversational participants are not accurately tracking the conversation if they think that there are two distinct things under discussion, say, a woman who walked in and another individual that ordered lunch. Thus the metaphor of a file card is useful. File cards record information that is presumed to be about a single object for the purposes of the conversation. They can be satisfied by objects in the world (or individuals in the model), but aren't required to be associated with any object in particular. Though talk of file cards is clearly metaphorical, the file card picture can be formally implemented in various ways, using familiar tools such as assignment functions.[9]

In the family of dynamic semantic views alluded to above, (4a) encodes an instruction to add a file card to the context, give it a label that has not yet been used, such as x, and then record on it *x is a woman* and *x walked in*. (4b), on the other hand, encodes an instruction to update the file card labeled x, adding to it the information *x ordered lunch*.[10] By using the notion of a context change potential, the semantics accounts for both novelty and licensing with a single explanation. An indefinite semantically encodes an instruction to introduce a new file card. This is what accounts for novelty. The introduction of this new file card also provides something for the sentence containing the pronoun to update. This accounts for licensing. Though the pronoun is not syntactically bound by the indefinite, the semantics of the indefinite bridges the gap between the indefinite and pronoun; the indefinite introduces something that the pronoun can pick up on.[11] Without the new file card introduced by (4a), the instruction encoded by (4b)—to update the file card labeled x—could not be carried out.

The dynamic picture described also accounts for some other cases in which indefinites license anaphoric pronouns beyond their syntactic scope, as in donkey sentences:

Footnote 8 continued
Anyone who prefers working in a discourse referent framework should feel free to substitute *discourse referent* for *file card* throughout. By adopting Heim's metaphor of a file card, I do not intend to take on any of the details of file change semantics.

[9] One standard example of such a view is *Dynamic Predicate Logic* (Groenendijk and Stokhof (1991)). The interested reader is encouraged to look at the Appendix for a somewhat formal comparison between DPL and an implementation of the view I present informally in § 3.

[10] In general, sentences with pronouns encode an instruction to update a particular file card; just which file card depends on the association of a pronoun with a particular antecedent, which is determined before the semantic machinery does its thing. How this association takes place is itself an interesting topic, but is beyond the scope of the present work.

[11] On certain formal implementations, this amounts to the indefinite *semantically binding* the pronoun. See the Appendix for more details.

(5) If a farmer owns a donkey, he beats it.

I am putting donkey sentences and other cases of indefinites embedded under quantifiers and operators aside for the present paper. For a discussion of these cases, see chapter 3 of Lewis (2011).

2.2 Against semantic accounts

The dynamic view just described offers an elegant account of indefinites in discourse. But if novelty is encoded in the semantic value of an indefinite, as such accounts claim, we should expect it to always be present. However, examining a wider variety of examples reveals that indefinite descriptions are not always associated with a new object under discussion. On the contrary, there are many felicitous uses of indefinites that do not describe a novel object under discussion at all. Consider the following discourses:

(6) a. A student walked into Sue's office and asked her about his exam.
b. Finally, a student needed her help!
(7) a. I went to see *Star Trek* on Sunday.
b. That's pretty much all I did all weekend: I saw a movie.
(8) a. We have this nail here.
b. Unfortunately, now we have a nail and no hammer.
(9) a. I went out to dinner with the woman from the bar last night.
b. Can you believe it—a woman went out to dinner with me!

These are not just cases in which an indefinite fails to introduce a novel file card. Rather, the second sentence in each discourse contains an indefinite that intuitively continues talking about an object *already* under discussion. Let's call these *summary uses* of indefinites. This contrasts with the other main use of indefinites, *introductory uses*, as illustrated by (4). Unlike the typical cases in which we don't get novelty (and licensing) because of the interaction between the indefinite and an operator or quantifier (such as negation), summary uses do not require the presence of a particular operator that takes scope over the indefinite. Summary uses are also not special cases constrained by how the object under discussion is previously introduced or described. In the examples, the familiar object is previously talked about using an indefinite description (in (6)), proper name (in (7)), demonstrative (in (8)), or definite description (in (9)). Furthermore, summary uses of indefinites can occur in subject or object position.

In each of (6), (7), (8), and (9), it would be a mistake to interpret the indefinites in the latter sentences as talking about a novel object under discussion. The most salient interpretation (if not the only interpretation) of (6) is that one and the same student walked into Sue's office and needed her help. The use of the indefinite expression in (6b) does nothing to hinder this interpretation, despite the fact that we all know many students might visit a professor's office. In (7), it is clearly understood that all the speaker did all weekend is see *Star Trek*, as evidenced by the fact that it would be infelicitous for an interlocutor to ask "What movie did you see?" after (7b). The same goes for the other summary cases. Pre-theoretically, if

one was to ask someone how many nails are under discussion in (8) or how many women went out to dinner with the speaker according to (9), the answers would be one nail and one woman, respectively.[12]

If a semantic account of novelty is right, the natural interpretation of these discourses should *conflict* with the meaning of the indefinites, and these examples should sound really bad. But they don't; in fact, they sound perfectly ordinary. The introductory use of indefinites obeys novelty. Summary uses do not. Novelty therefore only applies to some uses of indefinites. I take this to be a decisive reason against a semantic account of novelty.[13]

It is also worth noting that without a clear indication of a summary use (such as the word *finally*), the correct interpretation of a discourse like (6) is unclear. Consider:

(11) a. A student walked into Sue's office and asked her about his exam.
b. A student needed her help.

It is simply unclear whether there are one or two students under discussion in (11). This is unexpected on a semantic view of novelty. On a semantic account, the natural reading of (11b) should be that *another* student needed Sue's help.

Some other cases of non-novel uses of indefinites have been raised against the semantic views of novelty in the past. These cases involve embedded indefinites. For example, von Fintel (1994) discusses the following example in a footnote:

[12] After formulating these examples, I was pleased to discover that Szabo (2000) also notes an example of this kind:

(10) The detective ordered a Martini. As soon as the waiter left he knew that something was wrong. Then he realized what it was. He had just ordered a Martini from a waiter who looked exactly like the murderer he was after. (p. 37)

Szabo also uses this as evidence that novelty is pragmatic, but offers neither further argument for this fact, nor a pragmatic derivation of novelty.

[13] One thing the proponent of a semantic account of novelty might say in these cases is that they involve the *merging* of file cards. Multiple file cards can be merged into one during the course of a conversation if it is revealed that what were being treated as distinct objects under discussion are actually satisfied by one and the same object in the world. For example, suppose we are discussing Mark Twain and Samuel Clemens, unaware that they are in fact the same individual. If at some point someone informs the conversational participants that they are one and the same person, the file card for Mark Twain and the one for Samuel Clemens are merged into one. So someone might suggest that what is going on in the summary uses is that a new file card for the indefinite description is in fact being introduced, but then immediately merged with the relevant old one. I think this explanation of the summary cases fails for several reasons. There is nothing jarring about the cases. Unlike the case in which it is explicitly conveyed in the conversation that we have been mistakenly treating one object as two, there is nothing here to provoke a pragmatic repair of the context. Furthermore, a merging solution seems to be an *ad hoc* treatment of the intuitive data. Phenomenologically speaking, there is nothing in the summary cases that feels like the paradigmatic cases of merging file cards (as in the Mark Twain/Samuel Clemens case). Merging saves a technical notion of novelty by sacrificing the file card picture's ability to capture what is intuitively going with the objects under discussion (i.e. that there is one object under discussion in the summary cases, not two). An account of discourse dynamics that captures what ordinary people take to be going on in conversations is preferable to an account that makes technical distinctions that don't intuitively connect to what is going on in real life conversations.

(12) Show me a man who plays hard and I show you a man who deserves a beer.[14]

The most salient reading of (12) is the one in which for each hard-playing man shown to the speaker, it is that same man who deserves a beer. Intuitively the speaker is not proposing a game in which the interlocutor shows him a hard-playing man and the speaker shows him a *different* man who deserves a beer.[15]

One reaction to the cases of non-novel indefinites is to drop novelty as a semantic feature, but retain treating semantic values as CCPs to account for licensing.[16] On this approach, the fact that novelty is associated with indefinites in many contexts must still be accounted for. In the following section, I argue that novelty is best explained as a matter of pragmatics—an implicature of uttering sentences with indefinites in these contexts. To maintain a dynamic semantics, then, there must be two distinct explanations of the discourse properties of indefinites: a pragmatic one for novelty, and a semantic one for licensing. Given this, we lose the elegant simplicity with which the original versions of dynamic semantics accounted for all the discourse properties of indefinites. By contrast, I argue that we can maintain a traditional, truth-conditional semantics and give a uniform pragmatic explanation for novelty and licensing. Once the pragmatic account of novelty is developed, an account of licensing emerges naturally, making a semantic account of licensing superfluous.

2.3 Towards a pragmatic account

In the previous subsection I argued that novelty is not semantic, which is one reason to think it is a pragmatic phenomenon. But there are also positive reasons for thinking it is pragmatic. One of the marks of a pragmatic phenomenon is *cancellability*—something conveyed via pragmatics can be cancelled whereas something conveyed via semantics cannot (or it is at least much more difficult to find a situation in which conventional aspects of language are cancelled). If I am right, and novelty is an implicature of utterances containing indefinites, then it is what Grice called *a generalized conversational implicature*, since this is the sort of implicature that is normally associated with a sentence containing a particular lexical item or phrase across different contexts. A generalized conversational implicature can be cancelled in two ways.[17] It might be *contextually* cancelled, in that it is used in a context which makes it clear that the implicature is not present, or

[14] p. 65 fn. 59. Though clearly a case of non-novelty, this doesn't seem like the same sort of example as the summary cases. Introductory and summary uses of indefinites are not intended to exhaustively cover all possible uses of indefinites.

[15] Another case of non-novelty for indefinites, called *the requantification problem*, occurs when giving a focus sensitive or presupposition sensitive semantics for adverbs of quantification. These semantics require that some indefinites pick up old file cards. For more discussion see Rooth (1995), Krifka (2001), and von Fintel (1994).

[16] E.g., this is the approach of Krifka (2001) and Farkas (2002).

[17] See Grice (1989) pp. 39, 44. It seems that Grice himself actually misspeaks when he writes about generalized conversational implicatures. While generalized implicatures are supposed to work just like particularized ones in that they are derived from the content of an utterance, Grice writes that they are associated with "a certain form of words in an utterance". (p. 37) He must mean, rather, that the implicature is normally associated with the *content* of utterances containing a certain form of words.

it might be explicitly cancelled. Novelty bears both these features. We already saw that novelty is contextually cancelled in summary uses. But consider also the following example of explicit cancellation.

Suppose Ann has been telling Beth about a student of hers, Jane. After some discussion of Jane, the following dialogue occurs:

(13) Beth: I have to go, are we still meeting later for coffee?
Ann: I can't—I have a meeting with a student. In fact, it's Jane I'm meeting with.

Though the utterance containing *a student* at first seems to indicate that a novel student is being spoken about, this is immediately cancelled by the explicit statement that the student with whom Ann has the meeting is in fact the very same student they had been discussing. On the other hand, conventional or semantic phenomena, such as the contrastive feature of *but*, are not so easily cancelled. While (13) does not seem odd at all, (14) certainly does:

(14) Jane is short but nice—not that I generally think short people are not nice.

Though cancellability is not a decisive test for whether a phenomenon is pragmatic, it is strong evidence on the side of the pragmatic story, especially in this case. There are some cases in which cancellability can be an indication of something else. For example, Crimmins (1992) argues that the default interpretation of underarticulated expressions can be cancelled.[18] In general, default semantic interpretations can be cancelled when there's another semantic interpretation available. For example, the default interpretation of *I swam to the bank* is that I swam to the river bank, but I can correct that by explaining that I swam to a financial institution. But there's no evidence that novelty falls into any of these categories. There's no reason to think that novelty is the result of some sort of underarticulation. And if indefinites were ambiguous between a novel and non-novel reading, we would expect some language to have two different lexical items to encode those meanings. But to my knowledge, there is no such language reported in the literature. So an ambiguity thesis is unmotivated. In general, ordinary, non-ambiguous conventional meanings are difficult to cancel—especially when we're talking about meaning at the level of character rather than content. So while cancellability may not be a knock-down test, it is excellent prima facie evidence for exploring a pragmatic explanation. Furthermore, regardless of whether one considers the summary uses cases of contextual cancellation, I have already established them as a class in which indefinites are not associated with so much as a hint of novelty. Taking this into consideration along with the fact that the phenomenon is explicitly cancellable, there is overwhelming evidence that the novelty condition should be accounted for pragmatically.

[18] For example, suppose I am on the phone with you and I am in New York City and you are in Florida. If I say "it's raining", the default interpretation is *that it's raining in New York City*. But I can cancel that interpretation by explaining that I was watching TV and saw that it is raining where you are in Florida, which is what I meant to express.

3 The pragmatics of indefinites in discourse

3.1 Framework

Philosophers and linguists alike have criticized Gricean pragmatics for not being rigorous enough.[19] Many dynamic semanticists argue that a pragmatic theory of discourse phenomena is inadequate precisely because of the lax nature of pragmatic explanations. However, just because Gricean derivations are generally unsatisfactory, it does not follow that pragmatics lacks the explanatory power of semantics. The difference between semantics and pragmatics is not in the rigor of the explanations they offer, but in the type of explanation. We can accept Grice's claims about the nature of pragmatics—that it is based in principles of rational, cooperative activity—without endorsing any of his particular views.

Thomason (1990) suggests that if Grice is right about the nature of pragmatic phenomena 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".[20] I agree with Thomason in thinking that the applicable reasoning mechanism is plan recognition, and the data structure is the conversational context. In this subsection, I will lay out the necessary details of each of these elements in turn. In § 3.2, they will come together in the explanation of the discourse properties of indefinites.

I have already talked considerably about the role of the context in a conversation, so it should be evident that it is an important, independently motivated data structure. Though there is some disagreement about the nature, structure, and precise content of contexts, it is a plausible principle that the elements included in the context are just those that the conversational participants must track for the conversation to run smoothly. These must include the objects under discussion, since in any discourse, there are things under discussion. In most normal discourses, the participants continue to talk about the objects under discussion over the course of multiple utterances. To be able to follow a conversation, then, participants must keep track of the objects under discussion.[21]

We can adopt the same representation of objects under discussion as discussed in the previous section, namely, file cards. Recall that file cards record information about objects under discussion, and can be satisfied by objects in the world (or

[19] See for example Thomason (1990), Levinson (1983), Leech (1983), Harnish (1976), and Davis (1998) for some criticisms of Gricean derivations of implicature.

[20] p. 330

[21] It should be noted that some people disagree with this assessment, and think the only element that needs to be represented in the context is the common ground, i.e. the set of propositions mutually presumed by the conversational participants. Since the common ground represents everything that is presumed by the conversational participants, it will include information about the conversation itself, such as which objects are under discussion. Therefore, in principle, there should be some way to translate what I say here into an account that has only the common ground in the context. However, I do not think it is a fruitful way of thinking of things and so will not pursue it further here.

model), but are not themselves concrete objects in the world (or model). Since file cards are the only element of the context with which I am concerned in this paper, I adopt the simplified view that the context is a set of file cards. For now, we can think of this as an unstructured set, to which file cards get added but not removed. Nothing hangs on this simplifying assumption.

Planning and plan recognition are also independently motivated, as important parts of co-operative activities among rational agents. Plans have long been posited by philosophers and computer scientists to play crucial roles in activities such taking a trip, cooking, playing sports, fixing a car and so on. A well-run conversation is just like any of these other co-operative, rational activities. Take planning a trip with friends for instance. Achieving the goal of a successful trip is not a matter of performing random, disconnected actions. Rather, it requires having a coherent series of plans. Not only does one have to make plans such as where to go, how to get there, and what hotel to book, but they have to cohere in a way such that, for example, the hotel is booked in the very same place one chose to travel to. Likewise, participants in a conversation do not make random, disconnected utterances. A successful conversation also requires a coherent series of plans: not just what to talk about or how to answer a question under discussion, but also how an object under discussion relates to a question under discussion, for example.[22]

Of course, a complete plan for a typical conversation is not decided upon beforehand, but the sort of plans we will be concerned with are speakers' short-term plans, which we can call *local* plans.[23] Even in the most casual of conversations, speakers have at least basic, limited plans that drive what they say next, such as a plan to answer a question under discussion or a plan to introduce a new topic of conversation. For example, a speaker might mention the Bahamas in order to make the Bahamas the current topic of conversation.[24] This is no different from planning a trip: for instance, we may begin by researching desirable destinations, and which destination is chosen in part determines how the rest of the planning goes.

The sort of plans one makes in co-operative rational activities like conversing and trip-planning must be recognizable to the other participants in the activity. For example, I should make it clear if I plan to book a hotel room, or else we might end up with several hotels booked for the same night. Not making one's plans recognizable will likely result in a disastrous trip. Similarly, if a speaker wants to introduce the topic of going to the Bahamas into the context, the hearers must be

[22] Discourse plans are distinct from domain plans in that the latter is a plan regarding a task aside from the conversation, while the former is a plan for the conversation itself. The two can of course be related—my plan for the conversation might be intricately involved with achieving my domain plan—but they are nevertheless distinct. For example, suppose we are having a conversation about going to see a movie together tonight. I might have a domain plan to go see an action film; I might have a discourse plan (part of which is) to enter the topic of action films onto the conversational context. Clearly the discourse plan here is aimed to play a part in achieving my domain plan, but they are not the same plan. See Grosz and Sidner (1990) and Litman and Allen (1990) for further discussion on discourse vs. domain plans.

[23] Local plans might also be thought of a subplans, or elements of an overall plan. I am remaining neutral for now on the nature of plans.

[24] A speaker might have a weaker sort of plan, one that constrains the possible plans for the conversation rather than determining one particular future course, even in the short term. For example, a speaker might mention several places and continue to talk about one of them depending on what her interlocutor replies.

able to recognize this plan in order to properly track the conversation. Not doing so will likely result in a disastrous conversation. On my view, therefore, tracking the conversation—i.e. updating the conversational context—is a pragmatic process involving plan recognition.[25]

Plan recognition aids in co-operation in large part because, as Bratman (1990) argues, recognizing plans helps support expectations. Expectations, in turn, enable agents to make further plans based on them. For example, recognizing my plan to go to the departmental colloquium tomorrow supports your expectation that I'll be there, thus enabling you to plan on that assumption. It also supports my own expectation that I'll be there, allowing me to plan on the same assumption. The same holds true for discourse plans. For example, my plan to introduce the topic of the departmental colloquium supports both of our expectations that it's now (one of) the topic(s) of conversation, allowing us to both make further discourse plans accordingly.

Plan recognition is compatible with Gricean pragmatics.[26] Grice thought that pragmatic derivations were about recognizing speakers' intentions. Plan recognition, unlike intention recognition, not only emphasizes what a speaker wants the interlocutors to believe (or understand, or presume) but how the speaker wants to fit her contribution into the overall conversation. This helps explain some implicatures that are thought to be inexplicable on the Gricean maxims alone. Consider one of Davis (1998)'s criticisms of Grice's explanation of scalar implicatures. In certain cases of sentences that typically carry scalar implicatures, we don't primarily get the scalar implicature, but rather what Davis calls a "close-but" implicature, as in the following scenario:

(15) a. Did anyone in the class get an A?
b. Some students got B+'s.

(15b) does not primarily implicate *that not all students got B+'s*, but *that no students got A's*. This is certainly not explained by quantity, since *Some student got an A* and *Some student got a B+* are equally informative. Relevance alone—at least without detailing exactly what relevance is (i.e. in terms of planning)—also doesn't explain the precise implicature. (Davis doesn't consider relevance, but this may seem like a natural thought here.) But adding in plan recognition does explain the phenomenon. One common discourse plan is to answer the question under discussion. This gives rise to a general discourse expectation that when someone asks a question, the interlocutors will attempt to (at least partially) answer it. The second implicature is an answer to the proposed question, whereas the first fulfills no such discourse plan. Of course, Davis and others might point out that *some students in the class got an A*, *most students in the class got an A*, etc. are also compatible with the content of (15b) and answer the question posed in (15a). But (15b) is not a strategic way of giving positive answer to (15a). The speaker of (15b)

[25] At least in part—I leave it open that there are other ways in which the conversational context is updated; for example, updating with the fact that a goat just walked into the room is not a matter of recognizing anybody's plans. I will not be concerned with these other ways in the present paper.

[26] Thomason argues that plan recognition should *replace* Grice's conversational maxims, but I think it is better seen as compatible with Grice's project.

gives extra information that the speaker of (15a) is not looking for. This can be explained if the answer to the question is negative and the speaker in (15b) is giving a positive answer to a closely related question (i.e. *Did anyone in the class do well?*), but completely inexplicable if the speaker has the exact information the questioner is seeking. The foregoing explanation also does not require (15a) to explicitly be a question. Statements such as *Some people in my class got A's* are commonly thought to give rise to questions under discussion like *Did anyone in your class get an A?* A complete account of pragmatics in terms of speakers' strategic discourse plans remains to be worked out. But examples like the above suggest that this is the sort of thing Gricean pragmatics needs in general, and not just for the case of discourse dynamics involving indefinites.

3.2 The account

Now that the pragmatic framework is in place, let's turn to the specific explanation of the discourse properties of indefinites. I will argue that uttering a sentence containing an indefinite in typical introductory contexts implicates that a novel file card is being talked about. I will then argue that the pragmatic derivation of the novelty implicature also explains licensing.[27]

First it will be useful to review an often overlooked aspect of Gricean pragmatics. Grice thought there was a close connection between the maxim of relation and the maxim of manner (in particular, the supermaxim "be perspicuous"). Relevance and perspicuity go hand in hand because making relevant utterances is in part a matter of being perspicuous about the connection between the utterance and the rest of the conversation. Grice offers several examples in which an implicature is generated so that conversational participants can continue to assume that the speaker is obeying both relevance and perspicuity. That is to say, the implicature jointly appeals to relevance and perspicuity.[28]

The plan recognition framework further explains the intuitive connection between relevance and perspicuity. As mentioned above, local discourse plans that drive particular utterances are the sorts of plans that should be recognizable. Part of making a plan recognizable is a matter of being perspicuous in how it connects to the conversation as a whole, and thereby the conversational context that represents the state of the conversation. But this means that being clear about how one's local discourse plans relate to the overall plan for the discourse is part of making them

[27] In "Logic and Conversation", Grice suggests that some generalized conversational implicatures associated with sentences containing indefinites are a result of the maxim of quantity. I think this explanation fails in this case, because one of the alternatives to using an indefinite is using a proper name, and the two cannot be placed in an informativity scale, thus making it impossible to apply the maxim of quantity.

[28] See Grice (1989) pp. 31–32. Grice's two examples are each in their own way relevantly analogous and disanalogous to the derivation I give below. The reader should not take any details of these specific derivations as features of my view. The point is only that Grice seemed to think relevance and perspicuity often join together in generating implicatures.

recognizable, i.e. recognizable, perspicuous plans go hand in hand with relevant utterances. It is important to note that being perspicuous about the connection between an utterance and the conversation is both a matter of the content of the utterance and the manner in which the utterance is worded.

Recall typical introductory cases look like (4), repeated here as (16):

(16) a. A woman walked in.
b. She ordered lunch.

As previously mentioned, I am maintaining the traditional semantics of indefinites as run-of-the-mill existential quantifiers, and the semantic content of sentences as truth-conditions. Furthermore, while I remain neutral on the semantics of definite descriptions, pronouns, and names, I take it as a brute fact that whatever their semantics, they have the ability to pick out a particular object already under discussion, i.e. a file card already in the context. This is not to say that these definite constructions are conventionally associated with familiarity, but merely that they clearly have the ability to pick out particular objects, including particular objects that are familiar. I remain neutral on whether this is a semantic or pragmatic aspect of these definite expressions.

By assumption, the semantic value of an indefinite does not pick out a particular file card (or object in the world). Neither does it encode anything about novelty or familiarity. A sentence containing an indefinite simply makes a general, existential claim about the world. We are assuming that speakers are being co-operative and thereby making relevant contributions to the conversation, so the existential claim must in some way relate to the conversational context and the ongoing discourse plan. Furthermore, there are many ways in which to express the same content as an existential claim (for instance, with an equivalent existentially entailing claim), but the speaker chooses to use an explicit device of existential quantification (i.e. the indefinite).

Let's work through (16) to see how the derivation works. Suppose a speaker utters (16a) at some point during the course of a conversation. There are two relevant facts about the utterance. First, the content of (16a) is that there is at least one woman who has the property of walking in. Second, the speaker explicitly used the indefinite expression *a woman*. A good conversational participant will be as co-operative as possible in achieving the conversational goals. This involves letting one's discourse plans be as transparent as possible. Explicitly using the expression *a woman* is a fairly perspicuous way of revealing a plan to talk about a woman. Other semantically equivalent expressions—such as *not every woman not*—are not nearly as perspicuous revealers of plans. These sorts of expressions, to borrow a phrase from Grice, are unduly prolix. Moreover, they bury the word *woman* under two negations and a universal quantifier for no apparent reason. Of course, the most perspicuous way to reveal one's plans would be to constantly spell them out in detail. However, there are brevity constraints in discourse; utterances should be as transparent as possible while at the same time not being excessively long or verbose. This is a practical constraint of agents with limited resources, like time and patience. Levinson (2000) writes that "human speech encoding is relatively very slow: the actual process of phonetic articulation is a bottleneck in a system that can otherwise

run about four times faster".[29] If Levinson is right, a speaker really has to balance being transparent with being brief. Using a particular term to indicate a plan to talk about some object of that kind seems to be about as close to an optimal solution to that tension as we can get.

Co-operative conversational participants, in an effort to track the conversation, may ask themselves how the speaker's utterance of (16a) relates to the conversational context. Does the speaker want to convey information about a woman already under discussion, or is this woman novel to the discussion? If the speaker had wanted to pick out a particular woman already under discussion, she had a much better way to do so, one far less prone to interpretive error: she could have used a pronoun, definite description, demonstrative, or name.[30] But the speaker didn't do so. So unless there is some other clear reason for the speaker making an existential claim rather than one containing a definite expression, (16a) is indicative of a plan to convey information about a *new* woman under discussion. Talking about a novel woman is often indicative of a plan to go on and say something more about that woman. So co-operative interlocutors who want to responsibly track the conversation, upon recognizing the speaker's plan to talk about a new woman, will add a novel file card with the information *x is a woman* and *x walked in* (where *x* is a previously unused symbol).[31] The addition of the appropriate file card connects the existential claim to the set of file cards, making it relevant to the conversation in general. But this very same act explains why the subsequent anaphoric pronoun in (16b) is licensed—there is now an appropriate file card in the context which it can pick up on.[32]

The notion of introducing a novel object as indicative of a discourse plan to say something further about it needs to be clarified at once. Of course, there are plenty of reasons why a speaker might bring up an object that she won't mention again. There is no requirement on co-operative discourse that every object mentioned play an extended role in the conversation. However, if speakers are not making a series of disconnected claims, many of the objects mentioned will play an extended role in the conversation. Just imagine a conversation in which no objects (or the majority of objects) mentioned were ever picked up on again, as in (17):

(17) I saw a man. A woman bought a beer. A bird flew into a tree. A baby cried.

The discourse begins to sound ridiculous, like a laundry list of statements about the world rather than a proper conversation (or like a literary description of a scene, but

[29] p .6

[30] Recall that I am not assuming that these expressions are conventionally associated with familiarity, but merely that they have the ability to pick out familiar objects.

[31] Robin Jeshion (pc.) suggested there might also be something like a scalar implicature going on here, since the content of the sentence is neutral as to the number of women potentially under discussion, and so we might expect to get a numerically neutral file card, i.e. one that is not specifically singular or plural, rather than a singular file card. I think this is a perfectly natural addition to the pragmatic derivation. I am also open to the idea that it is in fact a neutrally numbered file card that gets added and the fact that only a singular pronoun is licensed is a grammatical, and not a semantic or a pragmatic, constraint. At this point, I am not sure how to decide between the two explanations.

[32] On the static view, how the pronoun gets its value from the file card is a matter of the semantics of the pronoun. The present account is compatible with various views on the subject. I explore one formal account—treating the pronoun as a free variable—in the Appendix.

literary devices often work precisely because they subvert linguistic expectations). In an ordinary conversation, one's gut reaction to something like (17) is to ask, "why are you telling me this?". Such an interlocutor is not making relevant contributions to the conversation. So while mentioning a novel object does not require a speaker to go on and talk about it, it does raise the conversational participants' credence that she will go on and talk about it. So mentioning something novel is really indicative of a plan to *potentially* go on and talk about it. Such a discourse plan effectively eliminates all the potential ways the discourse could go in which the object in question is definitely *not* an object under discussion. Recognizing a plan to potentially go on and talk about something based on the reasoning outlined above is reason to add the appropriate file card to the context, since participants don't know which of the objects mentioned will be picked up on again. In fact, even the speaker need not have a firm intention on whether she plans to go on to talk about an object. Perhaps, for example, she is waiting to see what her interlocutor replies before deciding what to say.[33]

The present explanation predicts the same output contexts for (16a) and (16b) as the dynamic accounts described in § 2. Recall that they were able to account for novelty and licensing by encoding the instruction to add a new file card in the semantics of the indefinite. By contrast, I have just argued that there are pragmatic reasons for updating the context with a new file card when a sentence containing an indefinite is uttered (on an introductory use). This replicates the elegant results of the dynamic semantic accounts without making the mistake of semantically encoding novelty and without abandoning traditional treatments of semantic content.

The pragmatic account just given also correctly predicts that existentially entailing sentences like (18) don't introduce a new file card:

(18) It's not the case that every woman didn't walk in.

While on the truth-conditional account of content we've been assuming, (18) has the same content as (16a), it does not explicitly contain a device of existential quantification. As I said before, (18) would be a bad choice for a speaker who wanted to make a plan to talk about a woman recognizable.[34] In general, the account predicts that, absent a special context, a file card gets added for only the descriptive content of the indefinite used. So, in general, file cards are introduced based on an explicit linguistic act.[35]

[33] Planning to talk about an object under discussion is not the only area in which speakers may have a range of plans rather than a specific plan in mind. Other examples include specificity resolution for questions and the interpretation of epistemic modals. (See von Fintel and Gillies (2008)).

[34] On some views of structured propositions such as that of King (2007), (18) and (16a) do not have the same content. On such a view, existential sentences are only those that explicitly contain an existential term like *some* or *a*. On this view, the novelty implicature runs entirely off content, though the explanation is otherwise the same.

[35] This is not to preclude the possibility of adding new file cards based on the perceptual environment or by accommodation. Under certain circumstances, I do think file cards can be added by accommodation. I will turn to some such cases shortly. File cards are also added for things salient in the perceptual environment of a conversation. For example, to borrow a well-known example from Stalnaker, if a goat walks into the room during our conversation, we all note that a goat walked into the room and record that fact. In our current terminology that amounts to adding a file card for the goat that walked in.

In the same way, the account explains the difference in anaphoric potential between other truth-conditionally equivalent sentences like (19) and (20):

(19) David is married.
(20) David has a spouse.

According to the view just defended, (19) does not introduce a file card for David's spouse, whereas (20) does. This gives the intuitively right result, since in the general case (19) does not license pronominal anaphora on David's spouse, whereas (20) does. It is worth emphasizing that this is a quite natural result of the plan recognition model. As I have repeatedly said, a good conversational participant makes her plans recognizable, and it is intuitive to think that not only the content of what she says but also the manner in which she expresses it are indicative of her discourse plans. In this case, the fact that the term *a spouse* was or was not explicitly used lets the conversational participants know about the speaker's local plans to (potentially) go on and talk about David's spouse. Of course a speaker can utter (19) with a discourse plan of saying something about David's spouse (e.g. "David is married. His wife is smart."), but she has done nothing to indicate a local plan to make David's spouse an object of discussion. This last point is worth emphasizing. The sorts of plans that play a crucial role in the pragmatic explanation are recognition of these local plans, because they allow the participants to track the discourse, i.e. know what to expect will likely be a topic of conversation, an object under discussion, or a question being addressed. It is my contention that speakers use and participants recognize maximally strategic plans (tempered by practical concerns, like brevity). The recognition of the possibility of any old plan on the part of a speaker to potentially talk about something is not sufficient reason for adding a file card. For example, in this case, I may recognize that a speaker who utters (19) may intend to go on and say something about David's spouse over the course of the conversation, but such a plan has not been indicated in a maximally strategic way on the part of the speaker.

The account also predicts that only one file card gets added for each indefinite. For example, (16a) results in the introduction of a file card for only a woman, not any of the objects normally associated with a woman or entailed by the existence of a woman, like hands or a heart. Some people have argued that file cards are entered into the context for every object whose existence is entailed by the state of the conversation. If this was the case, upon making an utterance using an indefinite, a file card would get entered for every object whose existence it entailed. Furthermore, file cards would get added after processing existentially entailing sentences like (18) and (19).[36] I disagree. If we are to take contexts seriously in

[36] For example, see Roberts (2003). This sort of position is generally motivated by a desire to account for the familiarity requirement on definites in the face of data that demonstrates that definite descriptions can be used without explicitly introduced antecedents. The position is also not entirely fruitful, since there are just as many examples of definite descriptions felicitously used where the existence of the object in question is not entailed by the context. For example, while the position on file cards may solve the problem for cases like "I went out for dinner last night. The waiter was rude.", it doesn't solve the problem for cases like "I went out for dinner last night. The salmon was divine.". Restaurants almost always have waiters, but they do not almost always have salmon.

pragmatic explanations, they have to be manageable. If the conversational context is going to be a tractable reconstruction of the information exchanged during a conversation, and something that has a key role in things like deriving implicatures and predicting felicitous anaphora, it must be a cognizable construct of reasonable size. Adding file cards willy-nilly for every object associated with an uttered indefinite description will soon result in an unruly jungle of a conversational context, one that seems an unlikely candidate for a psychologically real representation of a conversation.[37] Thus I take it to be the correct result that the present explanation predicts only the addition of a single file card of the appropriate sort for each indefinite.

Some might object that the novelty implicature does not seem like an ordinary Gricean implicature, especially since it conveys information about the discourse itself rather than other information about the world. There are several reasons why this objection isn't particularly worrisome. First, pragmatics has not been typically occupied with accounting for discourse information, and so there is little precedence for this sort of implicature. The important result in this paper thus far is that the same pragmatic principles that explain ordinary implicatures can also explain the communication of discourse information. That this is not exactly the ordinary sort of implicature is exactly the point. Nevertheless, one might still object that speakers don't have these sort of meta-intentions to communicate propositions about the discourse itself. I do think it is a difficult (and indeed perhaps unanswerable) question whether a speaker genuinely has an intention to communicate a complete proposition like *I am speaking about a new woman under discussion now*. However, speakers certainly have intentions about whether they want to continue talking about an old object or introduce a new one. I am open to the idea that what gets implicated is not propositional and so not a traditional Gricean implicature in this sense. Finally, there is *some* precedent among traditional implicatures for implicatures about the discourse itself. For example, a speaker might implicate that something his interlocutor just said was rude or irrelevant.

The novelty implicature also bears other markings of a traditional implicature, e.g., it is non-detachable in just the way we'd expect given that it is generated by relevance and manner. That an implicature is non-detachable means that there is no semantically equivalent utterance that lacks the implicature in question. Grice believed this feature to hold of all implicatures except those derived from the maxim of manner, since those are also based on the words used. Since the novelty implicature works off the presence of an explicit device of existential quantification, it is present when *a* is replaced by another term with the same content such as *at least one* and *some*. That is, the implicature is non-detachable so long as there is an explicit device of existential quantification. But when there is no such device, no implicature is generated and so it is absent from utterances like (18) and (19) though they may have the same content as (16a) and (20).[38]

[37] In Roberts (ms), she argues that the file cards that are actually tracked by the conversational participants are just a limited subset of all the file cards available in the context, the file cards that are relevant to the current question(s) under discussion.

[38] The implicature appears to be detachable in the case of *one or more*, but I take this to be mere appearance. Utterances containing the locution *one or more* do not appear to license singular anaphora,

Someone might object that these indefinites that all supposedly have the classic semantics of the existential quantifier do not in fact license anaphora in the same way. For example, while in the majority of contexts discourses like (16) are felicitous, the felicity of discourses like (21) below varies somewhat relative to different contexts.

(21) a. At least one woman walked in.
b. She ordered lunch.

It is beyond the scope of the present paper to examine all the differences between indefinite expressions. Intuitively, however, a pragmatic story should be able to accommodate the difference between the two discourses just mentioned, since one uses a phrase that makes the possibility of more than one woman walking in explicit, whereas the other leaves it implicit. However, there is data, due to Farkas (2002), that strongly suggests a difference in the semantics between *a* and *some*, because of a difference in their embedding behaviour. For example, (22) has a (very salient) generic reading, whereas (23) does not:

(22) A man always buys the first round of beer.
(23) Some man always buys the first round of beer.

This data doesn't require adopting a dynamic semantics, but it does suggest that there is some semantic difference between *a* and *some* (as well as the other indefinites). I leave this as an open question for further examination.

3.3 Summary uses

The explanation of the novelty condition is based on the idea that asserting the existence of some thing(s) will not be relevant, that is, it will not meet the conversational participants' discourse expectations that it connects to the set of file cards, unless it is construed as indicative of a plan to introduce something new. This is because there are other linguistic expressions available that do have the ability to pick out particular existing file cards, and do so in a clearer, more effective way, such that they'd be preferred by co-operative speakers. Failing to have used such a perfectly good locution generally indicates that the speaker did not have a plan to update an old file card. If the speaker had such a plan, making an existential claim

Footnote 38 continued
despite the fact that they are truth-conditionally equivalent to claims containing indefinites. One would expect given my account that they do introduce a new file card and license subsequent anaphora. However, it seems that the lack of felicitous singular anaphora is a grammatical constraint quite independent from anything discussed here. Utterances containing *one or more* do in fact introduce a new file card, but it's one that licenses plural anaphora. (A complete account of anaphora involves positing something like plural file cards anyway.) While "One or more women walked in. #She sat down." is bad, "One or more women walked in. They sat down." is fine, despite the first sentence's non-committance as to whether there is one or multiple women. Even if one takes *one or more* to have the same semantic value as *a*, *some*, and *at least one*, it is undeniable that the former takes a plural noun and requires plural agreement on the verb, while the latter take singular nouns and require singular agreement. Thus it should come as no surprise that the former takes a plural pronoun while the latter take singular ones.

would be needlessly general. It would be powerful confirmation for this view if when there *was* a reason to assert the existence of something regardless of the speaker's plan in terms of the set of file cards, novelty was not implicated. This is just what we see in summary uses of indefinites, as in (6) and (9) from § 2, repeated here as (24) and (25):

(24) a. A student walked into Sue's office and asked her about his exam.
b. Finally, a student needed her help!
(25) a. I went out to dinner with the woman from the bar last night.
b. Can you believe it—a woman went out to dinner with me!

In the summary uses the existential, general meaning of the indefinite is emphasized. (24b) is appropriate to utter in a context in which Sue had been waiting and hoping for *some student or other* to need her help—she isn't happy or relieved because that particular student came to her office in need of help, but that some student at all needed her help. The speaker has a special reason to use an existential claim, since it expresses something a definite expression cannot. If we replaced *a* with *the* in the summary uses, they would each convey something different, if they made sense at all. Since there is this special reason to use an indefinite and only an indefinite, we have reason to believe novelty won't be implicated. And this is exactly what we see. Part of what cancels the implicature that the indefinite is being used to introduce something new is the explicit marking with some sort of term of summary—the examples of the kind I have given are much worse if they are not accompanied by terms like *finally*, *can you believe it*, or at least a special intonation. Since novelty is a generalized implicature, some sort of explicit indication to show that it is cancelled is just what one would expect.

Summary uses of indefinites are almost purely existential: they convey information about an object of a certain kind, but which particular object of that kind is irrelevant for the purposes of the conversation. Of course, the conversational participants understand, in (24) for example, that the student who needs Sue's help is in fact the student introduced in the first sentence. A competent conversational participant will update the file cards accordingly. It is in this sense that the indefinite picks up on something old. It is also in this sense that the existential claim is relevant—it meets the conversational participants discourse expectations that it connects to the file cards.[39] Unlike the introductory uses, which implicate that something new is being talked about, revealing the speaker's intention to discuss a new object, the summary uses implicate that something old is being talked about, revealing the speaker's intention to add information about an old object (as in (24b)) or comment on an old object (as in (25b)).

[39] In some of the summary uses, such as (25b), the sentence doesn't add any information, but rather registers the speaker's surprise. While it is not straightforwardly an informational update, it still intuitively clear how such a statement is relevant to the set of file cards, since it comments on an existing card in the set. Such cases are probably more acceptable if repeating the known information helps to answer a question under discussion. The connection between objects under discussion and questions under discussion is an interesting one, but pursuing it is beyond the scope of the present work.

4 Applications and advantages

We are now in a position to explain the infamous marble problem:[40]

(26) a. Jodie dropped ten marbles and found all but one.
b. It's probably under the couch.
(27) a. Jodie dropped ten marbles and found nine of them.
b. # It's probably under the couch.

This pair of examples is typically thought to pose serious problems for traditional semantics. Since (26a) and (27a) are truth-conditionally equivalent, dynamic semanticists have argued that traditional semantics lacks the resources to predict why anaphora is licensed in the first case but not in the second.[41] The present account makes precisely the correct prediction. In (26a), the fact that the speaker mentions the one missing marble is indicative of a plan to continue talking about the missing marble. Thus a file card for the one missing marble is added to the conversational context. When we get to (26b), then, there's an appropriate antecedent for the pronoun—the file card for the one missing marble. (27a) does not mention the missing marble, but merely entails its existence. Since I have argued that the context is built based on recognizing speakers' strategic discourse plans, the lack of mentioning the missing marble in (27a) reflects a lack of a local plan to go on to talk about it. Thus when the conversation reaches (27b), there is no file card for the missing marble and thus no antecedent for the pronoun. And so the anaphora is correctly predicted to be infelicitous.

Furthermore, the present explanation has a significant *advantage* over the dynamic semantic accounts from § 2. It has been noted by many that the anaphora in (27b) significantly improves after a noticeable pause between (27a) and (27b). There is no natural explanation for this phenomenon on the type of dynamic semantic accounts described in § 2.[42] However, the present account offers a natural explanation of the phenomenon. An extended pause in speech is usually an indication that the speaker is thinking. Thinking is part of planning, and when a speaker pauses to think, one significant possibility is that she is thinking about what to say or what to do next. Therefore, conversational participants are more likely to accept a change in plans, since something has occurred that often indicates a change in plans.

In this case, a file card for the missing marble is added to the context by *accommodation*. Accommodation (introduced by Lewis (1979)) is the idea that the

[40] The example is originally due to Barbara Partee, and discussed in Heim (1982), Kamp (1988), and Groenendijk and Stokhof (1999), among others.

[41] There are other replies to this example on behalf of traditional semantics. For example, Stalnaker (1998) argues that one can account for the difference based on the referential intentions of the speakers. See chapter 4 of Lewis (2011) for arguments against a referential intentions account of anaphora on indefinites.

[42] This is not to say that a dynamic semanticist could not offer an explanation, but to claim that there is no explanation that stems naturally from the family of dynamic semantic views considered in § 2.

conversational context can be adjusted to make an utterance acceptable or true if the context would otherwise cause the utterance to be unacceptable or obviously false. Unrestricted, accommodation vastly overgenerates and sometimes seems *ad hoc*: it has the power to make almost anything an acceptable utterance. But it is commonly agreed that accommodation is a real phenomenon, albeit constrained by certain conditions. Recognition of a change in plans seems to be a paradigmatic case in which conversational participants should be ready and willing to accommodate.

The result is that the felicity of (27b) is going to vary depending on how salient it is that plans are being changed. This correlation of change in felicity with change in plans is further supported by other examples in which anaphora goes from bad to good with factors that indicate a change in plans. For example, it is generally accepted, and the present theory predicts, that the following discourses are infelicitous.

(28) a. Jodie is married.
b. # He is nice.

(29) a. John is a bicycle-owner.
b. # It is nice.[43]

But there are certain circumstances in which these cases sound better, or even perfectly felicitous. For example, here again, (28) and (29) improve a lot if the speaker pauses for a significant period in between the two sentences.

Let's consider a more fleshed out instance of the phenomenon so we can see more possible factors that play into it. There's a scene in the movie *When Harry Met Sally* in which Harry and Sally are making small talk on an airplane, and the following discourse takes place.

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

I think there are at least two factors that improve this example. Aside from the pauses, which I've already discussed extensively, there is no confusion as to Sally's possible discourse plans here. Sally knows her discourse plans will be made clear even without the explicit introduction of a file card for Harry's fiancé. Second, because there are two different speakers in the conversation, the first speaker's intentions haven't been violated. Marriage is already a topic of conversation at this point in the conversation (Sally has just said that she is not interested in marrying anyone right now), and so the most obvious reason for Harry saying what he does is to convey that he, unlike Sally, is interested in marriage. But Sally, who finds Harry extremely annoying, is surprised that anyone would want to marry him, and so has her reasons for changing the focus of the conversation to Harry's fiancé. If Harry had simply said all in one breath, "I am getting married and she is a lawyer", it wouldn't sound nearly as good as (30). This result is also predicted by the planning account, as conversational participants should be more willing to accept a discrepancy in plans interpersonally than intrapersonally, since people tend to fail

[43] These examples are from Heim (1982).

to comply with other people's plans more often than they fail to comply with their own.

Dynamic semantics treats the relationship between indefinites and pronouns as extremely systematic; the relationship is a matter of semantics. By contrast, on my account, the relationship is motivated by recognizing speakers' discourse plans in using indefinites. While appealing to accommodation to explain these cases on the dynamic semantic view seems *ad hoc*, accommodation is a natural part of my view: under appropriate circumstances, interlocutors are primed to recognize that discourse plans have changed and adjust the set of file cards accordingly. The notion of plan recognition and speakers' intentions predicts the variance in acceptability of such discourses; dynamic semantics fails to predict the same variance.

5 Conclusion

I have argued that the discourse properties of indefinites are best accounted for pragmatically, as an implicature derived from recognizing the discourse plans of the speaker. A significant advantage of the present account is that the phenomena are accounted for without complicating the semantics *or* the pragmatics. The general picture of pragmatics adopted in this paper is one dynamic semantics needs anyway, for things like deriving implicatures or ambiguity resolution. Dynamic semanticists do not deny the need for a pragmatics on top of their semantics. So my pragmatic view is no more complicated than one a dynamic semanticist need adopt anyway. But my semantics is considerably simpler. The dynamic semantic theories considered in § 2 build updates with both truth conditions and discourse information into the semantics. By contrast, I take semantic contents to be truth conditions alone.

The pragmatic account of indefinites in discourse is not intended to be a full-fledged theory of cross-sentential anaphora. Such a theory requires making a commitment on the semantics of pronouns, and the present account is compatible with various views on the topic. In the Appendix, I examine one particular formal treatment, but this is meant to be illustrative of the sort of formal accounts that the present view makes possible rather than a commitment on the semantics of pronouns.

The pragmatic account opens a whole new type of strategy for explaining cross-sentential anaphora. Previous static strategies for cross-sentential anaphora largely focus on the semantics of pronouns (namely e-type, d-type, and context dependent quantifier accounts).[44] Other strategies that have aimed to replicate the dynamic-style accounts have focused on the referential intentions of speakers when using indefinites. I argue against such accounts elsewhere, but, in brief, they fail because not all indefinites that license anaphora are used with referential intentions. In contrast to standard static accounts, the dynamic semantic accounts described in § 2 have a powerful and elegant explanation of the data: an indefinite introduces a file card that provides a value for the pronoun. But these accounts mistakenly

[44] See for example Evans (1977), Neale (1990) and King (1994), among others.

semantically encode a pragmatic phenomenon. They also require abandoning the traditional notion of semantic content, replacing it with context change potentials. The present pragmatic account allows static views to explain cross-sentential anaphora via the introduction of a file card which provides a value for the pronoun. Such strategies can be fruitful because they maintain the elegance with which certain dynamic accounts explain the discourse properties of indefinites both without abandoning traditional notions of content and without semantically encoding properties that are in fact pragmatic.

Acknowledgments Thanks are due to Josh Armstrong, Andy Egan, Thony Gillies, Gabe Greenberg, Robin Jeshion, Michael Johnson, Jeff King, Ernie Lepore, Eliot Michaelson, Sarah Moss, Jason Stanley, Will Starr, Catherine Wearing, the members of the Rutgers philosophy of language work-in-progress group, and the participants at the Bellingham Summer Philosophy Conference 2011 for helpful comments and discussion. Earlier drafts of this work also benefited from the comments of audience members at the Rutgers Center for Cognitive Science semantics discussion group and the Rutgers Philosophy Department graduate student colloquium. All mistakes are of course my own.

Appendix: DPL and PL_{+D}

In order to get a sense of how my view compares to the dynamic semantic views described in § 2, I will compare a standard example of such a dynamic semantics with a simple formal implementation of my view. The dynamic semantics I will look at is Dynamic Predicate Logic (henceforth DPL, Groenendijk and Stokhof (1991)). I will compare it with static, ordinary first order predicate logic augmented with the pragmatic view argued for in this paper, a system I call Predicate Logic plus Dynamics (PL_{+D}). First order predicate logic is almost certainly not the best logic for capturing natural language semantics. I mean to endorse neither of these systems as the ultimate examples of their genre. However, they serve as convenient illustrative examples, since most readers are familiar with predicate logic, and we need not look at a more complicated logic to capture the point I wish to make.

DPL

The syntax of DPL is that of ordinary first order predicate logic, and so I will not review it here. A DPL model is a standard model for predicate logic, a pair $\langle D , F\rangle$, where D is a non-empty set of individuals, and F is an interpretation function. Contexts are sets of assignment functions, which are total functions from variables to individuals in the model. Semantic values on this formulation of DPL are relations between input and output assignment functions, but this is equivalent to stating them in terms of functions from contexts to contexts (from sets of assignment functions to sets of assignment functions). See Groenendijk and Stokhof (1990) for a detailed discussion of this equivalence. Here are the semantic clauses relevant to the examples discussed in this paper. (I suppress reference to a model since it shouldn't cause any confusion for the present examples. I use the letters *g, h*, and *k* for arbitrary assignment functions.)

DPL semantics (relevant clauses)

1. $[\![t]\!]_g = g(t)$ if t is a variable
   $= F(t)$ if t is a constant
2. $[\![Rt_1 \ldots t_n]\!] = \{\langle g, h\rangle |\ h = g\ \&\ \langle [\![t_1]\!]_h \ldots [\![t_n]\!]_h\rangle \in F(R)\}$
3. $[\![\phi \wedge \psi]\!] = \{\langle g, h\rangle |\ \exists k : \langle g, k\rangle \in [\![\phi]\!]\ \&\ \langle k, h\rangle \in [\![\psi]\!]\}$
4. $[\![\exists x \phi]\!] = \{\langle g, h\rangle |\ \exists k : k[x]g\ \&\ \langle k, h\rangle \in [\![\phi]\!]\}$

Let's work through example (4), repeated here, and its (D)PL translation:

(31) a. A woman walked in.
b. She ordered lunch.

(32) a. $\exists x(woman(x) \wedge walked.in(x))$
b. $ordered.lunch(x)$

DPL treats discourses as conjunctions of the sentences in the discourse. Given the definition of conjunction, this amounts to finding the semantic value of the first sentence of the discourse, and using the output assignment functions of that sentence as the input functions to the second sentence in the discourse. So let's just go sentence by sentence. Plugging (32a) into the semantic clause for existentials, we get $\{\langle g, h\rangle |\ \exists k : k[x]g\ \&\ \langle k, h\rangle \in [\![woman(x) \wedge walked.in(x)]\!]\}$. This result will be much more perspicuous if we first reduce the last part (the semantic value of $woman(x) \wedge walked.in(x)$). This is a conjunction, so we apply clause 3, which guarantees that the output assignments of the first conjunct serve as the input to the second. Since each conjunct is an atomic formula, the formulas act as tests on the input context, passing through only those assignment functions from the input that assign x to something in the interpretation of the appropriate predicate. Since we're dealing with all atomic formulas, for which in each input-output assignment function pair the output is identical to the input, we need not worry about going through the steps of calculating the value of the conjunction. After applying clauses 3, then 2, and then 1 inside the existential, we now have $\{\langle g, h\rangle |\ h[x]g\ \&\ h(x) \in F(woman)\ \&\ h(x) \in F(walked.in)\}$. Now it's easy to see what the existential does. (In fact, all the extra steps we just took are because the definition must apply in the general case in which there may be another quantifier inside the scope of the existential.) An existential is an assignment shifter—it takes each of the assignment functions in the input and returns all of those that differ at most from it in that they assign x to an individual in the interpretation of *woman* and *walked.in*. The output context of (32a), therefore, is the set of x-variants of the input context that assign x to a woman who walked in. The output context of (32a) serves as the input to (32b), which, as an atomic formula, tests the context, allowing through just those assignment functions that assign x to something in the interpretation of *ordered.lunch*. The result is that the final output context contains all and only assignment functions that assign x to a woman who walked in and ordered lunch.

The discourse is true if and only if the output context is not empty, that is, there is at least one woman who walked in and ordered lunch. Thus DPL succeeds in giving the intuitively right truth conditions for the discourse and accounting for both

novelty and licensing. We can think of the assignment functions as recording information about the objects under discussion. Changes to the assignment functions reflect changes in information about objects under discussion. As outlined above, the CCP of an indefinite resets the potential value of a particular variable. This is what accounts for novelty—*x* is treated like a brand new variable, with which no previous information is associated. At the same time, it explains licensing. Though the pronoun is not syntactically bound by the indefinite, on this view it is *semantically* bound: the indefinite shifts the assignment functions and it is this new set of assignment functions relative to which the pronoun is interpreted. Finally, changing the context so that it includes only assignment functions that assign *x* to something in the interpretation of *ordered.lunch* is tantamount to updating the relevant file card with that information.

PL_{+D}

PL_{+D} is just ordinary first order predicate logic augmented with two pragmatic principles. It is intended to cover the same fragment of English as DPL. Models and contexts are defined in the same way as for DPL. Strictly speaking, treating contexts as sets of assignment functions doesn't exactly match the informal description of contexts employed in this paper. I've said we are taking a context to be a set of file cards, but this implies that each assignment function represents a file card, which it clearly doesn't. Technically, a more appropriate formal implementation of a file card is as a set of (partial) assignment functions, and so our formal contexts should be sets of sets of assignment functions. I will suppress this complication for present purposes, since the simpler version offers a better comparison with DPL, and still captures information about objects under discussion, even if the context is not strictly speaking a set of file cards. (One can think of the set of partial assignment functions for each variable as the file cards.) Following the ordinary Tarskian interpretation of predicate logic, the semantic value of formulas in PL_{+D} are sets of truth-making assignment functions. A formula is true (relative to a model) iff its denotation is not empty.

PL_{+D} semantics (relevant clauses)

1. $[\![t]\!]_g = g(t)$ if t is a variable
   $= F(t)$ if t is a constant
2. $[\![Rt_1 \ldots t_n]\!] = \{g | \langle [\![t_1]\!]_g \ldots [\![t_n]\!]_g \rangle \in F(R)\}$
3. $[\![\phi \wedge \psi]\!] = \{g | g \in [\![\phi]\!] \,\&\, g \in [\![\psi]\!]\}$
4. $[\![\exists x \phi]\!] = \{g | \exists k : k[x]g \,\&\, k \in [\![\phi]\!]\}$

In addition, there are two pragmatically motivated operations that relate utterances to contexts. One adds the truth-conditional content of an utterance to the context, and is formally modeled as set intersection. Intersective updates apply to every utterance: they take the truth-conditional content of an utterance (the set of truth-supporting assignment functions) and intersect them with the input context,

thereby outputting only those assignment functions that are compatible with the information conveyed by the utterance. I did not give a pragmatic motivation for intersection in the present paper, but it is already recognized by dynamic and static semanticists alike that this is easily done.[45] In brief, the explanation is as follows. One of the central goals of a conversation is to convey information. The conversational context encodes the mutual presumptions of the conversational participants, leaving as open possibilities anything compatible with the presumed information. Once an assertion has been accepted, it is only rational to eliminate all the possibilities that conflict with the information conveyed by the assertion. Given our assumptions about the nature of contexts and contents, the formal operation that models this is intersection.

The other pragmatic update adds a new file card to the context. Adding a new file card is modeled as a function that takes all the assignment functions in the input context and returns all the x-variants (or, more generally, variants for the appropriate variable) that assign x to an object with the properties predicated of x in the logical form of the utterance. As the reader can probably already guess, this update is pragmatically triggered when a sentence containing an indefinite in an introductory context is uttered. The function that returns the appropriate x-variants of assignment functions is the formal operation that represents what we've been informally glossing as adding a new file card. If we are modeling information about objects under discussion in terms of assignment functions, then updates regarding the objects under discussion must be modeled as changes to the assignment functions.

PL_{+D} pragmatic updates

Where C is the input context, P is the content of an arbitrary assertion, and H is an arbitrary predicate, the 2 pragmatic updates are:

1. $C[P] = C \cap P$ (Truth-conditional update)
2. $C[\exists xHx]$[46] $= \bigcup_{g \in C} \{h|\ h[x]g \,\&\, h(x) \in F(H)\}$ (New file card update)

Let's see how PL_{+D} explains (32). Conversational participants always update the context with the informational content of an assertion (by intersection). The content of (32a) is the set of assignment functions such that for each assignment function, there exists at least one x-variant that assigns x to a woman who walked in. If there is such an object in the model, this will amount to the denotation being the set of all assignment functions. The intersective update, therefore, has no effect on the input context. (32a), as an existential, also triggers the second pragmatic update. The function for adding a new file card takes the input context and returns all the x-variants of each assignment function that assign x to a woman who walked in. So after semantically and pragmatically processing (32a), the context is in the same state as after the semantic processing of the same sentence according to DPL. (32b)

[45] See in particular Stalnaker (1978) for more discussion on the pragmatic motivations of this update.

[46] I present the rule based on the simple case for purposes of perspicuity. The general rule for $\exists\, x\phi$ is a little more complicated to define.

will trigger the normal intersective update. Its content is the set of assignment functions that assign x to something that ordered lunch. Since all the the assignment functions in its input context assign x to a woman who walked in, the resulting output context will include only assignments that assign x to a woman who walked in and ordered lunch. The final output context is therefore the same one DPL predicts, and intuitively the correct one; the assignment functions record the information that there is at least one object that is a woman, and walked in, and ordered lunch.

I have been glossing over one complication of PL_{+D}. The problem with treating pronouns as straightforward free variables in a static semantics is that formulas containing free variables do not get the truth conditions we really want. The set of assignment functions (32b) determines, for instance, is the set of *all* assignment functions that assign x to an object that ordered lunch, and so the sentence would be true so long as something in the model ordered lunch. This may or may not be a bad consequence, depending on one's take of the data. This may just be the correct way of modeling the truth conditions of a pronominal sentence out of context. But within a context, we want a sentence containing a pronoun to pick out a subset of these assignment functions. In (32), this is the set of assignment functions that assign x to a woman who walked in. This is easily fixed, however, by making the semantic clause for pronouns sensitive to the assignment functions in the input context. This is an intuitive way to model the fact that pronouns are anaphoric—they must look to something prior for their value. Moreover, dynamic and static semanticists alike agree that it is perfectly acceptable to have context-sensitive expressions in a static semantics. It is the context-*affecting* nature of certain expressions that make some people think we need a dynamic semantics. If we want to leave the semantic clause for free variables alone (say, because we want to employ free variables for something other than the translation of pronouns), we can always introduce a new sort of symbol for translating the pronoun into the logic that acts as a context-sensitive free variable. I propose we translate pronouns as lower case p's (loosely following Dekker (2004)) with alphabetical subscripts that connect them to the variable associated with their antecedents. For example, (32b) would be translated as follows:

(33) $\mathit{ordered.lunch}(\mathrm{p}_x)$

And the semantic clause for such phrases would be:

$$[\![Rp_x \ldots p_z]\!] = \{g \in C | \langle [\![x]\!]_g \ldots [\![z]\!]_g \rangle \in F(R)\} \text{(where C is the input context)}$$

PL_{+D} offers analogous explanations of novelty and licensing to DPL, though pragmatically instead of semantically explained. Adding a new file card resets the x-values of the assignment functions in the context (like the semantic value of an existential in DPL); this accounts for novelty. It also outputs all the x-variants of the input assignment functions; it is this new context relative to which the anaphoric pronoun is interpreted, thus accounting for licensing. Whereas in DPL we saw something we thought of as semantic binding, in PL_{+D}, we can think of the relationship between indefinite and pronoun as *pragmatic* binding. A sentence containing an indefinite pragmatically triggers the assignment functions in the

context to be set up in a way such that when the information conveyed by a pronominal sentence is added to the context, the pronoun receives the correct interpretation.

References

Asher, N., & Lascarides, A. (2003). *Logics of conversation*. Cambridge University Press.

Bratman, M. (1990). *What is intention? In intentions in communication*. MIT Press.

Crimmins, M. (1992). *Talk about beliefs*. MIT Press.

Davis, W. A. (1998). *Implicature: Intention, convention, and principle in the failure of gricean theory*. Cambridge University Press.

Dekker, P. (2004). Grounding dynamic semantics. In M. Reimer & A. Bezuiden- hout (Eds.), *Descriptions and beyond*. Oxford University Press.

Evans, G. (1977). Pronouns, quantifiers, and relative clauses. *Canadian Journal of Philosophy, 7*(3), 467–536.

Farkas, D. (2002). Varieties of indefinites. In: *Proceedings of SALT* 12.

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

Groenendijk, J., & Stokhof, M. (1990) Two theories of dynamic semantics. In *Lecture Notes in Computer Science*.

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

Groenendijk, J., & Stokhof, M. (1999) Meaning in motion. In: K. von Huesinger & U. Egli (Eds.), *Reference and anaphoric relations* (pp. 47–76). Dordrecht: Kluwer.

Grosz, B., & Sidner, C. (1990). Plans for discourse. In: *Intentions in communication* (pp. 417–444). MIT Press.

Harnish, R. (1976). Logical form and implicature. In J. J. Katz, T. G. Bever, & T. Langedoen (Eds.), *An integrated theory of liguistic ability*. Thomas Y. Crowell.

Heim, I. (1982). *The semantics of definite and indefinite noun phrases*. University of Massachusetts Dissertation.

Kamp, H. (1981). A theory of truth and semantic representation. In *Formal methods in the study of language*, Mathematisch Centrum.

Kamp, H. (1988) Comments on stalnaker, belief attribution and context. In *Contents of thought*. University of Arizona Press.

Kamp, H., & Reyle, U. (1993). *From discourse to logic*. Kluwer Academic Publishers.

King, J. (1994). Anaphora and operators. *Philosophical Perspectives, 8*, 221–250.

King, J. (2007). *The nature and structure of content*. Oxford University Press.

Krifka, M. (2001). Non-novel indefinites in adverbial quantification. In C. Condoravdi, & G. R. der Lavalette (Eds.), Logical perspec- tives on language and information, stanford. CSLI Press.

Leech, G. (1983). *Principles of pragmatics*. Longmans.

Levinson, S. (1983). *Pragmatics*. Cambridge University Press.

Levinson, S. C. (2000). *Presumptive meanings: The theory of general- ized conversational implicature*. MIT Press.

Lewis, D. (1979). Scorekeeping in a language game. *Journal of Philo- sophical Logic, 8*(3), 339–359.

Lewis, K. (2011). *Understanding dynamic discourse*. Rutgers University Dissertation.

Litman, D. & Allen, J. (1990). Discourse processing and commonsense plans. In *Intentions in communication*. MIT Press.

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

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

Roberts, C. (2003). Uniqueness in definite noun phrases. *Linguistics and Philosophy, 26*, 287–350.

Roberts, C. (ms). Retrievability and incomplete descriptions.

Rooth, M. (1995). Indefinites, adverbs of quantification and focus seman- tics. In G. N. Carlson & F. J. Pelletier (Eds.), *The generic book*. Chicago University Press.

Stalnaker, R. (1978). Assertion. *Syntax and Semantics, 9*.

Stalnaker, R. (1998). On the representation of context. *Journal of Logic, Language, and Information, 7*, 3–19.

Szabo, Z. (2000). Descriptions and uniqueness. *Philosophical Studies, 101*, 29–57.

Thomason, R. (1990). Accommodation, meaning, and implicature: Interdisciplinary foundations for pragmatics. In P. R. Cohen, J. Morgan & M. E. Pollack (Eds.), *Intentions in communication*. MIT Press.

von Fintel, K. (1994). *Restrictions on quantifier domains*. University of Massachusetts at Amherst Dissertation.

von Fintel, K., & Gillies, T. (2008). CIA leaks. *Philosophical Review, 117*, 77–98.