

**This is the penultimate draft.
For citations please consult the published version.**

Do we need dynamic semantics?*

Karen S. Lewis
Barnard College, Columbia University
Department of Philosophy
klewis@barnard.edu

1 Introduction

I suspect the answer to the question in the title of this paper is *no*. But the scope of my paper will be considerably more limited: I will be concerned with whether certain types of considerations that are commonly cited in favor of dynamic semantics do in fact push us towards a dynamic semantics. Ultimately, I will argue that the evidence points to a dynamics of discourse that is best treated pragmatically, rather than as part of the semantics.

There is no doubt a lot to be learned from the work being done in dynamic semantics. My worry is that the move to a dynamic semantic framework is often insufficiently motivated, that the argument for dynamic semantics too often involves a quick move from the observation of dynamic phenomena in discourse to the conclusion that we need dynamic *semantics*. I think that traditional semantic and pragmatic theories have the resources to explain the same phenomena in similar ways. Furthermore, I think the latter are the *right* explanations.

*Thanks are due to Sam Cumming, Andy Egan, Thony Gillies, Michael Glanzberg, and Jeff King for helpful comments and discussion. Earlier drafts of this work benefitted from the comments and questions of the New York Philosophy of Language Workshop, and the members of Andy Egan's graduate seminar at Rutgers University. Extra thanks to the editors of this volume, Alexis Burgess and Brett Sherman, for their helpful comments and patience. All mistakes are my own.

The contrast between dynamic and traditional, static semantics is roughly as follows. On many traditional views of semantics, a sentence conveys information about what the world is or could be like. On this view, to understand a sentence is to know the conditions under which it is true. The semantic values of sub-sentential expressions are their contribution to the truth-conditions of the whole. On dynamic semantic theories, a sentence, as well as expressions smaller and larger than sentences, convey *context change potentials*, where a context is a representation of the state of the conversation, including things like the information mutually presumed by the conversational participants, the question(s) currently under discussion, or the objects under discussion (which I will refer to as *discourse referents*).¹ In other words, the semantic values of expressions are their contribution to the discourse. On this view semantic content is like instructions on how to update a context. To understand a sentence is to know how to change the context.

One of the central roles of a semantic theory is to explain what gets communicated by linguistic expressions (be it sentences, sub-sentential expressions, or larger chunks of discourse).² This is not to propose the naive view that what gets communicated is always and only the semantic value of expressions, but that these values play some important role in the explanation. I'm interested in the question of what, if anything, adopting a dynamic semantic framework buys us in terms of explaining communication with language. There are lots of interesting questions regarding dynamic semantics that I will not address in this paper, for reasons of space:³

A good reason to adopt a new framework over an old one is that it is in a position to play this role in a better fashion than the old framework.⁴ In the first half of this paper, I will argue that dynamic semantics doesn't obviously put us in a better position. The central difference between dynamic and static semantics is whether semantics or pragmatics accounts for certain features of communication. While static semantics plus pragmatics

¹I am remaining neutral on the metaphysics of conversational contexts. They might be subjective representations in the minds of the conversational participants, or abstract, objective objects. For now, nothing rests on this distinction.

²See, for example, Yalcin (2012).

³Another interesting question is the question of whether dynamic semantics buys us something in terms of explaining semantic entailment relations. I will put this question aside for the purposes of this paper.

⁴By *framework* I mean the very general idea behind all dynamic semantic theories that meanings are context change potentials. There are differing semantic theories within this framework.

seems to face in principle problems in accounting for discourse dynamics, I will argue that these problems are merely the result of a narrow conception of the resources that static semantics and pragmatics have available. In particular, I will argue that much of the case for dynamic semantics relies either explicitly or tacitly on a conception of the static semantics of sentences as propositions conceived as sets of truth-supporting worlds, and of pragmatics as a) working off the contents of entire sentences and b) something that cannot deliver the sort of systematic explanations warranted by the linguistic data. I will argue that once we give up some of these assumptions, it becomes apparent that dynamic semantics does not give us new expressive resources when it comes to explaining discourse dynamics.⁵ In other words, we can accept that contexts look much like the dynamic semanticist describes them, and that updates happen in much the same way as dynamic semantics describes, but deny that these updates are part of *semantics*. Rather, we can keep traditional, static contents and adopt a dynamic *pragmatics*. In the second half of the paper, I will argue that there are good reasons to think pragmatic explanations are more explanatorily fruitful than their dynamic semantic counterparts. Dynamic semantics treats discourse dynamics as a systematic phenomenon — this is supposed to be an advantage — but I will argue the data is not nearly as systematic as the treatment supposes. The static semantic/dynamic pragmatic account naturally explain the variances in the relevant data that dynamic semantics cannot naturally account for. Furthermore, many of the discourse updates are the sorts of updates we perform on non-linguistic information as well. Pragmatic principles, which I understands in a broadly neo-Gricean way, are general principles of rational, co-operative activity and not special to conversation.⁶ These naturally explain this similarity between non-linguistically and linguistically motivated updates to the context. We simply miss important generalizations if we treat information flow through a discourse as part of the semantics.

⁵Again, this is not to deny that dynamic semantics might give us more expressive resources when it comes to some other phenomena or that there is some other advantage to the dynamic framework.

⁶In fact, I think of neo-Gricean pragmatics in an extremely broad way, in that any explanation that appeals to general principles governing rational activities counts as pragmatic, even if we give up the co-operativity idea. In this broad sense, even certain kinds of game-theoretic explanations count as what I mean by *pragmatic*. Nor do I mean to take on board any of Grice's specific views.

2 Background

As I said, dynamic semantics takes the semantic value of a sentence to be its context change potential (CCP), and the meaning of sub-sentential expressions to be their contribution to the context change potential. On a traditional, static semantics, the semantic focus is on the sentence (relative to a context), whereas on a dynamic semantics the semantic focus is on the sentence (or sub-sentential expression) as part of discourse, relative to both an input and output context. Consequently, adopting such a framework is supposed to put us in a better position to capture certain data about the dynamics of discourse, certain robust and seemingly systematic phenomena that involve the flow of information throughout a discourse. Static semantics is often portrayed as being unable to account for: the two-way interaction between context and content, the fact that the order of sentences and sub-sentential expressions in a discourse matters for what is communicated, and that sentences often communicate more than merely truth-conditional content. Here are some examples that have commonly been taken to be suggestive of the need for dynamic semantics.

1. CROSS-SENTENTIAL ANAPHORA

- (1)
 - a. Jodie dropped ten marbles and found all of them, except for one.
 - b. It is probably under the couch.
- (2)
 - a. Jodie dropped ten marbles and found only nine of them.
 - b. ?? It is probably under the couch.⁷

There are two (related) challenges that discourses involving cross-sentential anaphora allegedly present for the static semanticist. First, the first sentences of each of the above discourses, though truth-conditionally equivalent, differ in their *licensing* abilities — (1a) licenses a singular pronoun (but not a plural one) and (2a) licenses a plural pronoun (but not a singular one). Second, the examples involve *unbound* pronouns. The pronoun in (1b) gets its value, somehow, from *one* in the preceding sentence, even though it is neither syntactically bound by it, nor can

⁷This is a slight variation on an example originally due to Barbara Partee, and discussed in Heim (1982), Kamp (1988), and Groenendijk & Stokhof (1999), among others.

it co-refer with it in any traditional sense (since *one* is not a referring term).

Considering the first problem first, the idea is that somehow (1a) and (2a) must change the context in different ways. This would explain why we are in different positions when it comes to the felicity of (1b) and (2b). The problem is not that there is no way to update a context in a static framework. On the contrary, it is commonly agreed that it is easy to define a pragmatically motivated operation on static propositions that updates the conversational context. Following Stalnaker (1978) (among others), we can think of the information mutually presumed by the conversational participants at any given point in the conversation as the *common ground*, and we can usefully model the common ground using a context set — the set of worlds that are epistemically (or doxastically) open given the common ground. When an assertion has been accepted in a conversation, rational participants won't keep worlds open that are ruled out by the newly accepted information. This is naturally modelled as intersection: intersect the proposition (conceived as a set of truth-supporting worlds) with the context set. The new context set will retain all the worlds from the input context that do not conflict with the content of what was asserted. Now it should be clear what the complaint about (1) and (2) is: (1a) and (2a) are truth-conditionally equivalent. If we take content to be truth-supporting worlds and update the context as just described (supposing that the input context is the same), the output context will be the same. But when we come to (1b), the context licenses the singular pronoun but not so for (2b). The dynamic semanticist's challenge to static semantics is that some other change to the context must be incurred by processing (1a) and (2a), respectively. Though not generally explicitly noted, this challenge must include the assumption that this change can't be defined by a pragmatically motivated operation that acts on static propositions.

For example, Gillies (2004) writes that this example provides *prima facie* evidence that meaning is more than merely truth-conditional, that “semantics is richer than propositional content” (p.597), which in turn provides evidence for a framework like dynamic semantics which takes meaning to be this richer sort of thing. Heim (1982) argues that traditional, static theories of cross-sentential anaphora cannot account for

this example without appealing to the wording of the two sentences to explain the difference. She claims that the phenomenon in question is a systematic one and a dynamic account can give a more systematic explanation than one that appeals to the wording (p.21-22). Groenendijk & Stokhof (1999) claim that this example leads to the conclusion that “the two opening sentences differ in meaning, and that hence truth-conditional content and meaning cannot be equated” and that therefore this example shows that “truth-conditional content is not the basic notion that oils the wheels of the interpretation engine” (p.52).

On a dynamic semantic framework one can say that (1a) encodes an instruction to update the context not only with the truth-conditional information, but with a *discourse referent* for the one missing marble (this instruction is encoded in the semantics for *one*).⁸ When we get to (1b), the singular pronoun has a discourse referent on which to pick up. On the other hand, (2a) introduces no such discourse referent for the missing marble, and thus when we get to (2b), there is no discourse referent for *it* to pick up on. Discourse referents represent information about the discourse itself — the objects that are under discussion in the discourse — rather than information directly about the world. (A discourse referent is an abstract representation of an object under discussion; it is neither an individual in the world/model nor a linguistic expression.) Hence, it seems that dynamic semantics is in a better position to explain the communication of *discourse information*, information about the discourse itself.

The very same aspect of the dynamic theory provides an answer to the second challenge as well (how to account for unbound anaphoric pronouns). Most static theories treat unbound pronouns as e-type or d-type—some form of disguised definite description.⁹ This is not the place to argue against that sort of theory, but just to point out that

⁸Here I am glossing over some important differences between different dynamic semantic theories of anaphora. For example, theories like Heim’s File Change Semantics (FCS) treat indefinites as contributing a variable, as does the pronoun, both of which are bound by a higher quantifier. On Groenendijk and Stokhof’s Dynamic Predicate Logic (DPL), indefinites are treated as existential quantifiers, which change the assignment functions in the context such that they can bind variables to the right outside of their syntactic scope.

⁹Another static theory — Context Dependent Quantifier theory — also treats pronouns as disguised quantifiers, but the type of quantifier varies with the context. In this example the quantifier would be an indefinite description.

dynamic semantics opens up new strategies for dealing with unbound pronouns. The context change potentials of certain kinds of expressions (such as indefinites) add information about the objects under discussion to the context; this updated information acts as the input context to subsequent sentences that may contain pronouns. Proponents of dynamic semantics (rightly, I think) point out that this captures an intuitive idea about the dynamics of discourse: sentences in a discourse do not provide isolated information; rather, information flows throughout a discourse.

2. INTRA-SENTENTIAL ANAPHORA

- (3) a. A pregnant woman got on the subway and I gave her my seat.
 b. ?? I gave her my seat and a pregnant woman got on the subway.
- (4) If a nun owns a book, she reads it.

On any dynamic theory, *and* is not commutative; rather *and* encodes an instruction to update first with the first conjunct and then with the second.¹⁰ In example (3a), the first conjunct introduces a new discourse referent for a pregnant woman; by the time we get to processing the second conjunct, there is consequently a discourse referent on which the pronoun *her* can pick up. By contrast, when the order of the conjuncts is switched in (3b), we first have to update with the first conjunct, and there is no appropriate discourse referent for the pronoun.¹¹ This explains both the infelicity of (3b), and why it is hard, if not impossible, to hear *her* as being coreferential with *a pregnant woman*. On the other

¹⁰Some claim that the non-commutativity of *and* is the hallmark of a dynamic theory. For example, Groenendijk & Stokhof (1989) write: “A semantics is dynamic if and only if its notion of conjunction is dynamic, and hence non-commutative.” (p.13). This can’t be quite right – for there are theories that take content to be ordinary propositions, but *and* is not commutative (e.g. Jeff King (2007)) and there could be theories that take content to be context change potentials but define the CCP of *and* as commutative. It is in fact true that all existing dynamic semantics, as far as I know, define conjunction non-commutatively.

¹¹Of course, there might be some other discourse referent around if this sentence doesn’t occur discourse-initially, but it won’t be a discourse referent for the pregnant woman.

hand, on a static framework, *and* is generally commutative. One might therefore expect (3a) and (3b) to express the same truth-conditional content and be equally felicitous. Again, this is an example of where adopting a dynamic framework seems to have the advantage in explaining what gets communicated in terms of discourse information and the effect that order has on what gets communicated in a discourse.

In (4), an example of *donkey anaphora*, the pronouns in the consequent of the conditional are also unbound. In this case, on a dynamic semantic framework, *a nun* and *a book* in the antecedent introduce discourse referents into the context that are picked up by the pronouns in the consequent. Such a view presents no *prima facie* problem for dynamic semantics (this is not to say that specific accounts of donkey anaphora are not problematic) because sub-sentential expressions can also convey CCPs. On the other hand, if the output of the semantic theory is the truth-conditional content of entire sentences, it remains to be explained both how the context is updated with discourse information, and how this happens sub-sententially.¹²

3. COUNTERFACTUALS

- (5) a. If Sophie had gone to the parade, she would have seen Pedro dance; but of course,
 b. if Sophie had gone to the parade and been stuck behind someone tall, she would not have seen Pedro dance.
- (6) a. If Sophie had gone to the parade and been stuck behind someone tall, she would not have seen Pedro dance; but of course
 b. ?? if Sophie had gone to the parade, she would have seen Pedro dance.¹³

Roughly, the sequence in (5) seems acceptable because (5a) is not evaluated relative to any possibilities that conflict with its truth; (5b) sub-

¹²Donkey anaphora also presents the additional problem of explaining how the indefinite gets a universal interpretation, but like dynamic semantics, static semantics should appeal to the interaction between the semantics of the indefinite and conditional. As in the case of cross-sentential unbound anaphora, most static theories of donkey anaphora appeal to e-type or d-type theories of pronouns.

¹³Example is from Gillies (2007)

sequently introduces into the conversation the more remote possibility that Sophie gets stuck behind someone tall, and so, it too, seems true. By contrast, in the reverse discourse, (6a) introduces the possibility that Sophie gets stuck behind someone tall, and when we reach (6b) we cannot ignore this possibility since it has been introduced into the conversation already. Thus (6b) is unacceptable. Again, the fact that the order of the utterances makes a difference in interpretation or acceptability is taken as evidence that the relationship between sentences in a counterfactual discourse is relevant to their semantics. Since dynamic semantic contents are context change potentials, it seems natural to represent these observations in the framework: a counterfactual is always evaluated relative to the context in which it is uttered and in turn changes that very context.¹⁴

These are not the only examples of linguistic phenomena that motivate the move to a dynamic framework, but they are representative of the central motivating data. I think the same sorts of considerations that demonstrate that pragmatic explanations have better resources to explain these phenomena extend to the phenomena I have not mentioned (such as presupposition, questions, and epistemic modals).

3 Discourse dynamics and communication

I want to challenge the idea that dynamic semantics is in a better position than static semantics to explain the sort of phenomena outlined above. There are many detailed dynamic semantic accounts out there, and it is not the task of the current paper to evaluate them or offer rival accounts. Rather, I want to address the question of whether discourse phenomena like the ones outlined in the previous section should motivate us to adopt a theory within the dynamic semantic framework or not. The question is whether these phenomena exert pressure on us to change our basic semantic notion to that of context change potential. I will argue that they don't; static semantics plus a dynamic *pragmatics* has the resources to account for information flow through a discourse: the two-way interaction between context and content, the communication of discourse information, and the effect of the order on

¹⁴Roughly this argument is made in von Stechow (2001) and Gillies (2007), though I will look at Gillies' more nuanced argument in more detail in §3.3.

acceptability, interpretation, and truth. As I said before, on the static semantic/dynamic pragmatic view, contexts, conceived as representations of the state of the conversation, are the same as those in the dynamic semantic theory. Likewise, utterances do have the effect of changing the context. But on this view, context change is not a semantic notion.¹⁵

I think the case for dynamic semantics depends on a narrow conception of the range of semantic theories available to the static semanticist as well as a narrow conception of what pragmatics can do. Specifically, the arguments for dynamic semantics ubiquitously assume that the content of static sentences are propositions conceived as sets of truth-supporting worlds. But there are at least two other kinds of semantic objects that could be the content of sentences within a static framework: structured propositions or unstructured sets that contain more than just worlds (e.g. sets that also contain assignment functions or information states). In the case of pragmatics, arguments for dynamic semantics often assume that pragmatics does not operate sub-sententially and that pragmatics can't deliver explanations of (seemingly) systematic phenomena. I argue that neither of these assumptions is right.

The question I want to consider is not merely whether a static semantics can account for the linguistic data that motivates dynamic semantics — there are already static theories of the various phenomena out there (such as the e-type theory of anaphora), and the present project is not to evaluate their merits or flaws. Rather, I want to take seriously the dynamic semantic notion that information flows throughout a discourse, that content and context have a two-way interaction, that discourse information is part of what gets communicated, and argue that static semantics plus pragmatics has the resources to give this dynamic style of explanation of what's going on.

3.1 Case study 1: Cross-sentential Anaphora

I said before that the notion that dynamic semantics is in a better position to account for discourse dynamics relies on the assumption that static semantic contents are sets of truth-supporting worlds and that pragmatics works off only these contents. If the only semantic value we have to work with is a set of worlds (or similar object, such as a set of situations), and we maintain

¹⁵A good example of work that takes dynamic pragmatics seriously is the work of Craige Roberts; see for example Roberts (2004, ms). Of course, David Lewis (1979) and Robert Stalnaker (1978, 1998) are predecessors of this view. Philippe Schlenker (2008) also exemplifies what I have in mind.

the idea that these are the only input to pragmatics, then the pragmatically motivated update operations available to be defined are limited by what we can do with such a set. If semantic content provides only truth-conditional information, then it is mysterious as to where updates about other sorts of information come from. But if we have different sorts of semantic values acting as the input to pragmatics, different sorts of updates become available, because the semantic values provide additional information. And if we give up that pragmatics acts only on the content of sentences, different sorts of updates also become available. I illustrate this point here by suggesting some ways in which these assumptions can (indeed, should) be given up.

Many philosophers who work in the static, truth-conditional tradition hold that propositions are *structured*.¹⁶ There are many different views of structured propositions, but the basic idea is that structured propositions are complex entities made up of parts, and these parts reflect the semantic values of the elements of the relevant sentence. Furthermore, the constituents of a structured proposition are also generally thought to be bound together in a way that reflects the structure of the sentence. To return to the marble example, on any well-known view of structured propositions (1a) and (2a) do not express the same proposition. The proposition expressed by (1a) will contain the semantic value of *one* (or, a constituent that is determined, at least in part, by the semantic value of *one*), while the proposition expressed by (2a) will contain the semantic value of *nine* (or a constituent determined at least in part by it).

In fact, the problem raised for the possible worlds view by the marble example is not all that different from the sort of problems proponents of structured propositions have been raising against the view for decades. The possible worlds view predicts that truth-conditionally equivalent sentences express the same proposition, and people have objected that this is not a fine-grained enough account of propositions. It becomes particularly problematic when propositions are embedded under attitudes, since it seems not everyone who believes one proposition believes everything truth-conditionally equivalent to it. In the same way as sets of worlds propositions aren't fine grained enough to play their roles when it comes to belief, they may also not be fine grained enough to play their roles in discourse.

Not only do (1a) and (2a) have different contents on the structured propo-

¹⁶For example, King (2007), Soames (1985, 1987, 1989), Salmon (1986a,b, 1989a,b), Lewis (1972), and Creswell (1985), not to mention Russell and Frege.

sitions view, but they are different in just the sort of way we'd expect, given the difference in anaphora licensing in the examples. The structured proposition expressed by (1a) actually contains a constituent reflecting the fact that one marble was mentioned in the utterance. The structured proposition expressed by (2a) contains a constituent reflecting the fact that nine marbles were mentioned in the utterance. Thus, given some additional assumptions about the nature of contexts and how they are updated, it is unsurprising that the output context of (1a) contains some sort of representation that there is one missing marble under discussion, while the the output context of (2a) represents nine found marbles as being under discussion. After all, these are respectively represented in the semantic content of what was said. A structured proposition not only represents the way the world is, but also represents the way in which the information was presented. It reflects the structure of the sentence, which is a fact about the discourse. In this way, structured propositions can be said to encode discourse information.

Of course, structured propositions themselves don't encode instructions on how to update the context with the relevant discourse information. But unstructured propositions don't encode instructions on how to update the context with truth-conditional information. Rather, based on the information the unstructured proposition give us, we can give a pragmatically motivated explanation of the appropriate update (recall the Stalnakerian updates on the context set described earlier). The problem with unstructured propositions is that they gave us truth-conditional information and nothing more — we couldn't look inside them, so to speak, and see the structure of the sentence behind them. But structured propositions give us more information, and in this case, just the right sort of information we need to explain the difference in the discourses. I haven't yet given you a pragmatic update rule (I will below), but it is no longer mysterious where such a rule might come from, since the relevant information is right there in the content.

I don't mean to argue that we *need* give up semantic contents as sets of truth-supporting worlds in order to account for the dynamics of discourse; we can give up some assumptions about pragmatics instead. Nor do I mean to imply that the above is the only revision a static semanticist can make to static contents.¹⁷ The above merely illustrates one good option.

¹⁷See Yalcin (2011) for such an example. Yalcin treats epistemic modals in a way that mimics the dynamic account without making use of dynamic semantics; rather he relies entirely on static semantics conceived as sets of worlds *and* information states.

Heim thought it a downside of a potential pragmatic account of anaphora that it would have to rely on the words used and not just semantic content to account for cases like the marble example. But I don't see this as a particularly undesirable aspect of a pragmatic account. After all, one of Grice's four maxims, the maxim of manner, deals not in content but in form: be clear (avoid obscurity of expression, avoid ambiguity, be brief, be orderly). But regardless of whether Grice himself thought so or not, it is reasonable to suppose that rational agents engaged in co-operative communicative activity reason based not only on the content of a sentence but also on the particular words used. As I argued in Lewis (2012), if we think of wording as reflecting a speaker's discourse plans, the explanation is also systematic and predictive. If I'm right about this, then even if one maintains that semantic contents are sets of truth-supporting worlds, pragmatics has the resources to explain the difference between the marble cases.

If we maintain that the content of a sentence is a set of truth-supporting worlds, one already has to accept that certain things are going to depend on the words used rather than the content, since this way of thinking of content just doesn't cut the grain fine enough. For example, belief reports will have to rely on modes of presentation or some such thing, either to explain their differing truth values (on a semantic view) or explain why it seems to the agent, speaker, or listener that there is a difference in truth value even if there really is none (on a pragmatic view). While the marble cases is different because nearly everyone (except those really bad at math) know that the sentences are truth-conditionally equivalent, we can still appeal to the wording (or modes of presentation) to explain why it is clear to the discourse participants that the conversation is headed in different directions. Even for those who know that Mark Twain and Samuel Clemens name the same person, using one name indicates a plan to talk about the individual qua writer while the other doesn't. Similarly, choosing to mention the missing marble reveals a different discourse plan than choosing to mention the nine found ones. I myself prefer the structured propositions view, but think this is also a theory with enough resources to be worth pursuing.

Whether or not we choose to adopt structured propositions or the notion that pragmatics works off both content and wording, the pragmatic update rule for adding a new discourse referent can be explained in essentially the same way. Recall the lesson from the Stalnakerian pragmatic update on propositional content: it explains why informational content has the effect it does on the context by appealing to general features of conversation qua

rational activity with the purpose of gaining information. Similarly, we can explain why indefinites, in many contexts, affect the context by adding a new discourse referent by appealing to general features of conversation qua rational, (generally) co-operative activity. The pragmatic explanation of the update rule requires taking a few assumptions on board, but these are entirely assumptions I think we should make anyway, and nothing specific to the case of indefinite descriptions. First, to follow a conversation, participants must keep track of the objects under discussion (this should be fairly obvious, in that many conversations involve objects under discussion over the course of multiple utterances, though it is further required by the fact that language is anaphoric and some subsequent utterances will simply be unintelligible if the participants are not keeping track of the objects under discussion). Objects under discussion are just what I've been referring to by the technical term *discourse referent*. So in a well-run conversation, conversational participants must keep track of discourse referents as well as informational content. Second, rational, co-operative activity involves *planning*. Even in the most casual of conversations, speakers have local (i.e. short term) discourse plans (or at least plans in the sense of constraining the possible ways the discourse can go). These plans have to be recognizable — if the speaker's discourse plans are not recognizable to the participants, the discourse plans for the conversation will not be cohesive. (For example, one person might introduce a question under discussion, but no one will answer it because the plan wasn't recognizable.) Recognizing the speaker's plans supports expectations of where the conversation will or can go. (For example, recognizing that a speaker has introduced a question under discussion supports the expectations that someone in the conversation will go on to try to answer it.)

To better understand the pragmatic derivation of the update rule *add discourse referent*, consider a simple case of cross-sentential anaphora as in:

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

Take the semantics of the indefinite to be that of the ordinary existential quantifier. By assumption, then, the indefinite doesn't pick out any particular discourse referent or object in the world. Assuming that a speaker is being co-operative and making relevant contributions, the existential claim must in some way relate to the conversation. If we think of the state of the conversation as represented by the conversational context, then another way of

stating this is that the utterance has to relate to the evolving conversational context in some way. Most of the time, the merely existential information is not all that informative. The more relevant plan involves the objects under discussion. Since the speaker did not use a definite expression in (7a), which commonly pick up on familiar discourse referents (I'm leaving open the question of whether the latter is a semantic or pragmatic phenomenon), the recognizable plan is to introduce a novel discourse referent. Furthermore, the speaker did not use an existentially entailing expression such as *The bar is not devoid of women*, but one which explicitly invoked the term *a woman* (or, on the structure propositions account, expressed a proposition that has an indefinite description, or its components, as a constituent). This is again most readily understood as a plan to introduce a woman under discussion (or at least a plan to constrain the conversation such that a woman could be under discussion). Recognizing such a plan raises the participants' discourse expectations that the speaker or someone else will go on to say something more about the woman under discussion, and so the appropriate discourse referent is added to the context. This is not to say that all indefinite expressions introduce discourse referents that are later picked up on. Of course we can make existential claims in discourse without intending to go on and say more about the object in question. But it would be a very weird conversation if the speaker(s) made many existential claims and never picked up on any of them.

Similarly, in the marble example the fact the speaker utters the words *except for one* in (1a) (or expresses a proposition with the relevant constituent) indicates a plan to potentially go on and talk about the missing marble. This creates a discourse expectation among the conversational participants that there is a higher probability that the speaker (or someone else) will now go on to say something about the missing marble, and therefore participants update the context with a discourse referent for the missing marble. In (2a) the speaker chooses different words: she explicitly says the phrase *found only nine of them*. This creates the expectation that she is likely to go on and talk about the nine found marbles, not the missing one, and so the context is updated with a discourse referent for the nine found marbles, but not the missing marble.¹⁸

¹⁸See Lewis (2012) for much more detail on the pragmatic explanation of adding discourse referents to the context and Lewis (MS) for how to extend this basic story to embedded cases.

Notice that in these cases the update to the context is exactly the same as the one in the dynamic explanation; the difference is just that the dynamic update is part of the semantics of indefinite descriptions, whereas the static update is a pragmatically motivated operation that typically acts on assertions of sentences containing indefinite descriptions. The point is that broadly Gricean pragmatics has more resources than the marble argument for dynamic semantics allows, and these resources can be exploited to create pragmatic update systems that account for discourse data in the same style as dynamic semantics. It remains to be determined which system better captures the right sort of explanation, and which one accounts for the right entailments, but the point here is that (despite claims to the contrary) the need for updates to the context regarding discourse information is not reason enough to adopt the dynamic semantic framework. Static semantics and traditional pragmatics have the resources to account for these sorts of updates.¹⁹

It might be unsurprising to some that combining structured propositions with pragmatic update rules (or pragmatic update rules that can effectively work off the structure of sentences) yields functionally equivalent (or close to functionally equivalent) results to a dynamic semantics. After all, Discourse Representation Theory (DRT) is functionally equivalent to dynamic semantics. And DRT, at its core, involves static structured entities (DRSs) and construction rules, which not only describe how to construct the DRSs, but encode the information about context change (such as adding a discourse referent). DRT doesn't consider their DRSs to be propositions or the construction rules to be pragmatically motivated, but structurally the theory is very similar to the one I am proposing.²⁰ Another way of running the argument I've been making is: DRT is functionally equivalent to dynamic semantics. Structured propositions plus pragmatically motivated update rules are functionally equivalent to DRT. Therefore, structured propositions plus pragmatic update rules are functionally equivalent to dynamic semantics.

¹⁹There are other pragmatic responses to the marble argument that maintain that content is a set of truth-supporting rules. I have in mind particularly Stalnaker (1998), who appeals to the referential intentions of the speaker to explain the difference in output context. For arguments against this account, see Lewis (Forthcoming).

²⁰There are also idiosyncratic features to Kamp & Reyle (1993)'s version of DRT involving the specific semantics they want to give to specific constructions (e.g. indefinite descriptions are not quantifiers on their view), but these are not intrinsic features of the DRT framework.

3.2 Case Study 2: Intra-sentential Anaphora

In the intra-sentential anaphora cases, the purported problem for static semantics doesn't rely only on the assumption that truth-conditional content acts as input to pragmatics, but also on the assumption that pragmatics doesn't work off the value of sub-sentential constituents. Recall that in each example, it seems the context is updated with a novel discourse referent before the entire sentence is processed: in (4) the context is updated before we get to the consequent and in (3) before we get to the second conjunct. Dynamic semantics has the resources to account for the sub-sentential data in the same way it accounts for the cross-sentential data because sub-sentential expressions also express context change potentials. This is not to endorse any particular dynamic semantic account of anaphora; particular accounts all have their own problems. The point is simply that this sort of sub-sentential data is something that the framework is in just as good a position to account for as the cross-sentential data. On the other hand, it might seem that static semantics is off the bat in a worse position. This has nothing to do with particular theories, but with the scope of the framework. If the semantic content of an entire sentence is the input to pragmatics, then even if the appropriate updates to the contexts can be pragmatically motivated when it comes to sentences, the sub-sentential data presents an additional obstacle since the pragmatic updates don't apply. But this is no obstacle for the static framework *in general*. This is again only an obstacle for one, narrow conception of the static framework — one in which pragmatics works only off of sentence contents. Anyone who wants to join the ranks of the many philosophers who already think that pragmatics works at a sub-sentential level has no in principle obstacle here at all. For example, Levinson (2000), Taylor (2001), and Simons (2010) all argue, for different reasons, that a broadly Gricean conception of pragmatics applies at the sub-sentential level. Furthermore, the reasons people give for pragmatics acting at the sub-sentential level generally have nothing to do with discourse dynamics. Thus these resources that enable the static semanticist to deal with sub-sentential updates are also independently motivated.

There are many arguments for the existence of sub-sentential pragmatics that are beyond the scope of the present paper, but I think the important question is to ask whether sub-sentential expressions have a conversational purpose. The basic idea in Gricean pragmatics is the co-operative principle, that conversational participants will make contributions appropriate to the

purpose of the conversation, or stage of the conversation. If we accept that (at least certain) sub-sentential expressions make conversational contributions, then the same sort of reasoning that can be performed on the contents of full sentences can also be performed on the contents of sub-sentential expressions. Sub-sentential expressions do seem like the sorts of things that interlocutors can grasp and reason about, and I think that this gives us good reason to think that pragmatic effects occur based on sub-sentential expressions.²¹

To name but a few, Simons (2010) and Taylor (2001) offer good arguments to this effect; they give very different evidence that, taken together, strongly point to pragmatic reasoning on sub-sentential expressions. For example, Simons points out that sub-sentential clauses can at times express the main point of a particular utterance and be addressed by interlocutors. As an illustration of these phenomena she offers examples like the following:

- (8) a. Where did Jane go last week?
 b. Henry believes she was interviewing for a job at Princeton.
 c. Henry thinks she was interviewing for a job at Princeton.
 d. Henry said she was interviewing for a job at Princeton.
 e. Henry hinted she was interviewing for a job at Princeton. (p.146)
- (9) A: If Jane comes later, we can fill her in.
 B: She won't be coming. (p.147)

In (8), the clause embedded under the propositional attitude verbs is what answers the initial question. In (9), speaker B — quite naturally — addresses the claim embedded in the antecedent of the conditional. This is good evidence that sub-sentential clauses have conversational purposes, that they are the sorts of the things interlocutors grasp and reason about.

Taylor (2001) has a very different argument that points in the same direction: he argues that the sort of reasoning that goes on in completing propositional content is the exact same sort of reasoning involved in calculating implicatures — but the former reasoning involves partial, sub-propositional content. For example, consider the contrast in what is communicated by the following two sentences (given that they are uttered in 2013 in the United States):

²¹What I take to be the important feature is weaker than what some think is necessary for motivating pragmatic updates on the context, that is that uttering a sub-sentential expression involves performing a self-standing speech act.

(10) I've had breakfast.

(11) I've had eel.

In a typical scenario (10) conveys that the speaker has had breakfast *already that morning* while (11) conveys that the speaker has had eel *at some point in her life*. Succeeding in communicating these contents, respectively, relies on the speaker's intentions and the mutual world knowledge of speaker and interlocutors: that we live in a society in which people eat breakfast every day, but not one in which we eat eel every day. As Taylor points out, the interpretations can easily be flipped if we change facts about the context: (10) can get the *at some point in the past* interpretation if we're talking to a dinner-eater in a society of people who each eat only one meal a day (i.e. there are breakfast eaters, lunch eaters, and dinner eaters). (11) can get the *already today* interpretation if we imagine we're at a sushi tasting event and everyone at the event is trying to insure that everyone gets at least one taste of each kind of sushi. Since this reasoning is required for figuring out the proposition expressed by the speaker, it has to work off partial content. So again, partial contents seem to be the sorts of things interlocutors can grasp and reason about.

Finally, Stalnaker (1978) points out that interlocutors reason (at least about certain things) as the speaker is speaking – not afterwards on the entire content. For example, as he says, “when I speak, I presuppose that others know I am speaking, even if I do not assume that anyone knew I was going to speak before I did.” (p.86) Rather, when I begin speaking, I expect them to update with the commonplace fact that I am speaking, and this can furthermore be useful in conversation, as this is how they know who is tired when I say “I am tired” — and they don't have to wait until the end of the sentence to figure that out. Again, all this evidence points to the conclusion that sub-sentential or sub-propositional pragmatic reason is all too common. Thus static semantics has no *prima facie* problem with donkey anaphora either, since updates can occur mid-sentence. The same sort of reasoning that I argued for in the previous section regarding discourse referent update based on entire sentences can be applied at the sub-sentential level.²²

²²See Lewis (MS) for a more detailed discussion of these arguments and account of donkey anaphora on a static semantic/dynamic pragmatic framework.

3.3 Case study 3: Counterfactual conditionals

To return to the case of counterfactuals, the problem once again stems from assumptions about the limits of a pragmatic explanation. Recall that the two counterfactuals in (5) (the *Sobel sequence*) seem true, and in fact according to most semantics for counterfactuals both come out as true.²³ If the context gets updated with nothing but truth-conditional content, reversing the order of the two counterfactuals shouldn't make a difference to their truth — updating with one truth generally doesn't make another, compatible truth go false. The upshot is that a counterfactual doesn't just affect the context in terms of a truth-conditional contribution, but also in terms of which worlds are in the domain for evaluating a subsequent counterfactual. To use the terminology from the standard semantics for counterfactuals, the observation in this case is that one of the effects of a counterfactual on the context is to determine which worlds are part of the domain of closest worlds. (6a) expands the domain of closest worlds to include ones in which Sophie is stuck behind someone tall at the parade. Given this context, (6b) is false. (Gillies and von Fintel argue for treating counterfactuals as strict conditionals with contextually determined domains, but this is not the only way the same result can be accomplished.)

Both Gillies and von Fintel argue that this calls for a dynamic semantics. For example, von Fintel writes:

There are essentially dynamic facts concerning the way the order of counterfactuals in a sequence matters to the coherence of the sequences and to the plausibility of arguments. These facts demonstrate the need for context-change in the semantics of counterfactuals. (von Fintel, 2001, p.7)

Gillies also begins by jumping from a concern about the effect a counterfactual has on the context to talking in terms of CCPs as part of the semantics of counterfactuals, though ultimately he has a more nuanced argument for why the semantics of counterfactuals have to be CCPs. But before I turn to Gillies' argument, which involves looking at *might*-counterfactuals as well as *would*-counterfactuals (*mights* and *woulds* for short), let's examine how a pragmatic account could explain the data from Sobel and reverse Sobel sequences.

²³Gillies and von Fintel use the data as an argument against the Lewis-Stalnaker semantics for counterfactuals, but really it is a more general argument than this.

Gillies and von Stechow both argue that counterfactuals have presuppositions on their antecedent, namely a counterfactual *If P, would Q* presupposes that there are some P-worlds in the relevant domain. They both want to treat this presupposition in terms of a CCP: processing the counterfactual includes changing the relevant domain to add P-worlds if there are none, otherwise the process idles. Taking counterfactuals to be sensitive to the contextually determined domain, which can be updated as just described, explains the difference between discourses like (5) and (6). In (5a), the relevant domain is expanded to include the closest worlds in which Sophie goes to the parade;²⁴ these are all worlds in which she sees Pedro dance, and so the counterfactual comes out true. The presupposition of (5b) expands the domain to include the closest worlds in which Sophie goes to the parade and gets stuck behind someone tall; these are all worlds in which Sophie does not see Pedro dance, and so this counterfactual too comes out true. But reverse the order, and things are very different. (6a) presupposes there are worlds in the relevant domain in which Sophie goes to the parade and gets stuck behind someone tall; these are worlds in which Sophie does not see Pedro dance and so (6a) comes out true. But when we get to (6b), there are lots of worlds in which Sophie goes to the parade in the domain already, and so the CCP has no effect. But these worlds include ones in which Sophie does not see Pedro dance (since she gets stuck behind someone tall), and so (6b) is false.

There are three elements that make this explanation work, and none of them require treating the semantics of counterfactuals as involving CCPs. First, the semantics of counterfactuals have to be sensitive to a set of worlds determined by the conversational context. (This just requires context-sensitivity in the semantics of counterfactuals, which everyone agrees does not require a dynamic semantics.) Second, counterfactuals have to update the context with the presuppositions introduced by their antecedents. Third, these updates have to be the sort that survive to be the relevant input context for a subsequent counterfactuals in the same discourse (i.e. the presupposition is accommodated globally rather than locally). It's these latter two that von Stechow and Gillies think the CCP machinery is needed for, and for which I think a pragmatic explanation can be naturally given.²⁵

²⁴This assumes, for the sake of argument, an ordering on worlds in which Sophie getting stuck behind someone tall at the parade are farther away than worlds in which she goes to the parade and does not get stuck behind someone tall.

²⁵Moss (2012) also argues that the relevant data can be accounted for pragmatically, though her view is significantly different from the one I propose. She argues that coun-

If we accept von Fintel and Gillies' claim that counterfactuals have these presuppositions (and I'll question whether we should accept that in a moment), we don't need to adopt the dynamic semantic framework to account for how or why the presuppositions get accommodated, since there are widely accepted pragmatic notions of accommodation, following David Lewis and Robert Stalnaker. If asserting something like (6a) really presupposes that there are some Parade-and-Tall-worlds, and if it is uncontroversial that there are such worlds, then the conversational participants accommodate, i.e. they adjust the context minimally to include Parade-and-Tall worlds. We need not appeal to something in the semantics of the counterfactual presuppositions to explain this.²⁶

Subordination explains why the addition of Parade-and-Tall worlds is a change that sticks around for the duration of a string of counterfactual utterances. In general, a context change induced by a modal, quantified or conditional sentence will stick around until the modal, quantificational, or conditional talk ends, since subsequent sentences are understood as subordinated to previous ones as in:

- (12) If there are non-specialists at the meeting, be sure to make your work accessible to them. If they ask you a confused question, politely clarify for them.
- (13) A wolf might walk it. It would eat you first.

Unless the speaker can rule out the Parade-and-Tall worlds as irrelevant, once introduced, they cannot be ignored. Notice that this is essentially the same account for what is going on in (5) and (6) without appealing to CCPs.

terfactuals of the form *If P & Q, would R* raise to salience the might-counterfactual *If P, might Q*, which pragmatically clashes with the relevant would-counterfactuals in these cases (making them unassertable, not false). Again, my aim is to show that we can give a theory that is essentially the same as the dynamic semantic account except that the updates are pragmatic. Like on Gillies and von Fintel's theories, the account I sketch predicts that the problem in the reverse Sobel sequence is semantic rather than pragmatic.

²⁶Of course, there is a longer argument that may be in the background: the idea that all presuppositions should be treated as CCPs to explain presupposition projection. But this is not the argument either of them make, and Gillies even says that the presuppositions of counterfactuals are not really presuppositions at all, since they don't pass the relevant test. So even if one were motivated by the above reasoning for presuppositions in general (which I am not), it would be hard for at least Gillies to have this in the background, since his "presuppositions" aren't real presuppositions in the first place.

There are also other resources available for the static semantic/dynamic pragmatic account. Since the so-called presuppositions on counterfactuals don't act much like presuppositions anyway, we can give up the idea that there are such presuppositions, and stick with the idea that the entire semantics of a counterfactual *If P, would Q* consists in the claim that the P-worlds in the minimal domain are Q worlds. Like in the Lewis-Stalnaker analysis, how worlds are ordered (and thus what ends up in the minimal domain) is contextually determined. Unlike the Lewis-Stalnaker account, one way the context can change is that the minimal domain gets bigger when salient possibilities that can't be ruled out are introduced. (6a) raises to salience the possibility that Sophie goes to the parade and gets stuck behind someone tall. Salient possibilities, unless they can be ruled out, get added to the relevant set of possibilities for evaluating counterfactuals.²⁷

I've suggested two ways in which a static semantic/dynamic pragmatic theorist could account for the discourse dynamics displayed by the difference between (5) and (6), giving an account that makes the same predictions as that of von Stechow and Gillies in terms of which counterfactuals are true and which are false, and why. But I would be remiss if I didn't address Gillies's central reason for adopting CCPs: might-counterfactuals (von Stechow does not talk about these). Gillies points out that there are sequences very similar to the Sobel sequences that involve might-counterfactuals. The contrast between these sequences cannot be explained by appealing to the presuppositions of the antecedents:

- (14) a. If Sophie had gone to the parade, she would have seen Pedro dance; but, of course,
 b. if Sophie had gone to the parade, she might have been stuck behind someone tall and then wouldn't have seen Pedro dance.
- (15) a. If Sophie had gone to the parade, she might have been stuck behind someone tall and then wouldn't have seen Pedro dance; but, of course,
 b. ??if Sophie had gone to the parade, she would have seen Pedro dance.

²⁷Again this is close to, but different from, Moss (2012)'s account, since the possibility raised to salience here is just $\diamond(P \& Q)$, rather than *If P, might Q*.

Gillies argues that might-counterfactuals presuppose the possibility that there is an antecedent and consequent world, i.e. a might counterfactual *If P, might R* presupposes $\diamond(P\&R)$ (which is also its truth-conditional content). On his view, might-counterfactuals amount to instructions to add a previously unconsidered possibility to the context, i.e., add the closest P&R-worlds. Gillies argues that given the might data, we can't treat content statically with dynamic pragmatic update operations defined on top of static contents. His argument goes as follows. If we take presuppositions to be requirements on the definedness of a sentence, that is, the sentence expresses no proposition unless the presupposition is satisfied, then we cannot treat mights both statically and as the duals of woulds without encountering unwanted falsehood or gappiness. Our first option is to say that might-counterfactuals are true at a context if they are true when that context is updated with their presupposition. In this case, might-counterfactuals are pretty much everywhere true, including the input context for (14a), making the latter false (where the intuition we want to capture is that it is true) — hence the unwanted falsehood. Or, we could say that might-counterfactuals are true in the updated context and get no truth value at the input context (where the presupposition is unsatisfied). In this case, a lot of would-counterfactuals that we thought were defined go from defined to undefined. Since *If P, might R* is equivalent to \neg (*If P, would $\neg R$*), and a negation sign never takes anything from defined to undefined, whenever *If P, might R* is undefined, so is *If P, would $\neg R$* , and Gillies is right to point out that this is a highly undesirable result. Since Gillies's semantics is CCPs all the way down instead of defined in terms of truth, he avoids both these results. But we need not adopt a dynamic semantics to avoid these results. Here are three ways to avoid this problem while maintaining the ability to account for the contrast between (14) and (15) in a similar way to Gillies.

All solutions require giving up the idea that $\diamond(P\&R)$ is a definedness condition on *If P, might R*, but I think there is good reason to give this up. The first two solutions involve giving up $\diamond(P\&R)$ as a presupposition. The third solution involves giving up the idea that these are the sort of presuppositions that dictate definedness conditions. All the options have the central solution in common: might-counterfactuals, as well as woulds, are defined in contexts in which the alleged presuppositions are not satisfied. In such contexts, would-counterfactuals are vacuously true, and might-counterfactuals false. However, like most cases of accommodation, the context that is important for evaluating the counterfactual is one in which the so-called presupposition

is satisfied.

One reason to give up $\diamond(P\&R)$ as a presupposition is to be consistent with the option I presented above for the Sobel sequences that involved giving up any sort of explanation in terms of presuppositions. Since the so-called presuppositions don't act much like presuppositions anyway, this path looks like a natural one. Another reason to give it up is that $\diamond(P\&R)$ looks even less like a presupposition of might-counterfactuals than the presuppositions on the antecedents of would-counterfactuals do. Gillies argues that the cases are parallel, and should be treated as one phenomenon. But consider the following contrast. Given certain facts about the conversation and the world (such as the fact that I am an average basketball player and Kobe Bryant doesn't currently have casts on both arms), conversational participants are likely to accept the first of these counterfactuals but not the second:

- (16) If I had played one-on-one with Kobe Bryant yesterday and he had casts on both arms, I would have won.
- (17) If I had played one-on-one with Kobe Bryant yesterday, he might have had casts on both arms, and then I would have won.

If Gillies is right, both of these counterfactuals have the same CCP; they instruct the hearer to (permanently) add some one-on-one-and-cast-worlds to the context and test to see if all those worlds are worlds in which I win the match. Each is true if and only if this update is successful. But the first is clearly true, while the second is not. So it seems that the roles that the antecedents of woulds and the antecedents-plus-consequents of mights play are actually different. (I will return to this in the next section, when I discuss why the pragmatic view can better account for this contrast.)

So how to account for the contrast between (14) and (15)? I agree with Gillies and Lewis (1979) that mights (whether counterfactual or on their own) can serve to introduce new possibilities into the conversation that were previously ignored. But I think this can be explained entirely by the semantics of might-counterfactuals plus pragmatics. Take (14) for example. When the conversation begins, the participants are considering as relevant only the worlds in which Sophie goes to the parade and has a clear view. When a speaker utters (14b), she must want to expand the set of worlds that are relevant to the conversation, for at least two reasons. Since worlds in which Sophie goes to the parade and gets stuck behind someone tall were not in

the realm of relevant possibilities up until this point, if she doesn't want to expand the realm, then her utterance is irrelevant. Worse still, it is false, and so she must intend to expand the context to make it true.²⁸ The rest of the conversational participants know this, and so the set of contextually relevant worlds (in the minimal domain) is expanded.

The third way to avoid Gillies' dilemma is to maintain that everything he calls a presupposition is a presupposition, but deny that they are the sort of presuppositions that are required for definedness. For example, the presupposition associated with *too* is likely such an example:

(18) John went to a hockey game tonight, too.

(18) presupposes that some other conversationally salient individual went to a hockey game on the night in question, but even if this presupposition fails, we can still assign a truth value to the sentence based on whether John went to a hockey game on the night in question. One might think that the same holds for existence presuppositions on quantifiers.

(19) Every lawyer in Dawson City is a smoker.

(19) presupposes that there is at least one lawyer in Dawson City, but we can still assign a truth value to it if the presupposition fails; it is vacuously true since the set of lawyers minus the set of smokers is the empty set.

The upshot of this section and the last two is that the dynamic semantic framework does look to be in a better position to explain communication (at least about discourse dynamics) when we compare it to a static framework that takes sentence contents to be sets of truth-supporting worlds and pragmatics to work only off of sentence contents. But when we compare the dynamic framework to the full range of static options, neither framework is in an obviously better position than the other. The examples outlined in the introduction of this paper don't seem to offer any evidence (be it *prima facie* or not) in favor of dynamic semantics when we allow that the static semanticist has more tools in her toolbox than unstructured sets of truth-supporting worlds that get fed to pragmatic machinery. Both frameworks seem to be in

²⁸Or, at least she must either assume there is already such a possibility in the context or that the context will be expanded to accommodate such a possibility. In this case, if she assented to (14a), she has to assume the latter, but there are other cases, such as when the might-counterfactual comes discourse initially, that it could be the former.

an equally good position to offer the same style of explanation of how speakers communicate changes that should be made to the context. How, then, to decide between the two frameworks? In the following section, I explore some ideas that speak in favor of the static semantic/dynamic pragmatic approach.

4 Some considerations against dynamic semantics

Up until now, I've been arguing that the dynamic semantic framework doesn't put us in a better position than the static semantic framework in terms of explaining the relationship between semantic value and communication, that static semantics and neo-Gricean pragmatics (broadly construed) together have the resources to give an account of discourse dynamics, including the communication of discourse information, the effect of order on interpretation, and the two-way interaction between context and content. Now I want to suggest a few considerations as to why dynamic semantics puts us in a worse position.

Dynamic semantics presents update rules as part of the recursive compositional semantic calculus (in fact, as the central semantic notion). This suggests a systematicity in context change dictated by these rules. For example, in Heim (1983) updating by intersection (the truth-conditional effect of a full sentence or sentential clause) happens invariably, since it is just part of the meaning of a sentence. As Beaver (2001) points out, "Heim provides a method of stating semantics... in such a way that admittance conditions²⁹ can be read off from the semantic definitions without having to be stipulated separately" (p. 85). In virtually all dynamic semantic accounts of unbound anaphora, the explanation involves something that is semantically introduced by the antecedent which then provides a value for the pronoun (either by binding it, putting constraints on it, or being bound by the same higher quantifier). On the subject of counterfactuals, Gillies and von Stechow argue that the worlds in the presupposition invariably expand the domain of worlds — this is just the semantic effect of processing and accepting a counterfactual.³⁰ In fact, this systematicity is often cited as a central reason

²⁹A formula is admitted in a context iff the context satisfies the presuppositions of the formula.

³⁰These are all provided the sentence is accepted. The update need not occur if the

in favor of dynamic semantics. For example, Hardt (1999) writes:

Since Montague, a primary focus of semantics has been to describe a compositional method for constructing the logical representation of a sentence meaning, and then evaluating that representation with respect to a given context. A major insight of dynamic semantics is that sentences have a systematic relation to context in two ways: not only are they evaluated with respect to the current context, but they also systematically change that context.(p.187)

In the same vein, one of Heim's central arguments against static views of anaphora is that such theories can't differentiate between the felicity of minimal pairs like the following:

- (20) a. John has a spouse. She is nice.
 b. John is married. ??She is nice.
- (21) a. John has a bicycle. He rides it daily.
 b. John is a bicycle-owner. ??He rides it daily.³¹

Though the second discourse in each pair is (more or less) equivalent to the first, and each raises to salience John's spouse and John's bicycle, respectively, the anaphoric pronoun is not felicitous (or at least significantly less felicitous than in the first discourse). This is essentially to reiterate the same point the marble case made.

I've been arguing that static semantics together with pragmatics can account for these sorts of updates (including these apparently systematic differences). But I think it can account for it better, because there are important generalizations about context change that are best captured by a pragmatic account. By putting the updates into the semantics, dynamic semantics misses out on these generalizations. Semantically encoded aspects of meaning are the conventional ones. But updates to the context seem like the sort of things that are amenable to explanation from more basic principles of conversation, or rational co-operative activity in general, and not convention. In particular, the pragmatic account has a natural explanation of why updates

sentence is not accepted by conversational participants.

³¹Heim (1982) p. 24

to the context can vary with factors such as conversational purpose. And the data is considerably more unsystematic than the arguments for dynamic semantics would have us believe. While I am not arguing that dynamic semantics couldn't account for the variations in updates I'm about to point out, I am arguing that there is no *natural* explanation that already falls out of the explanation of the update rules, as in the pragmatic case. This feature of the static semantic/dynamic pragmatic account taken together with the expressive resources I argued for above create a powerful reason to doubt that dynamic semantics is the right semantics.

I have mentioned three sorts of updates in this paper: updates with truth-conditional information (i.e. the elimination of worlds that conflict with asserted information), updates with a new discourse referent, and updates with new worlds relevant for counterfactual discourse. I will consider them one at a time. Whereas updating the context with the truth-conditional information conveyed by an assertion was a part of the semantics for a dynamic system like Heim's, as I mentioned in §2, Stalnaker explains the same update by appealing to the fact that rational agents who are engaged in an activity with the purpose of gaining new information won't keep around epistemic possibilities that conflict with information that they have accepted. One important advantage of the Stalnakerian explanation is that we can use the same sort of reasoning to predict what sorts of updates to the context there are in conversations with different purposes (i.e. ones in which the central purpose is not inquiry into which world is actual). As we might expect, there are different updates in a conversation where the purpose is, say, to get a lot of suggestions on the table (e.g. a brainstorming session). In this sort of conversation it is rational to keep around incompatible possibilities upon accepting assertions, since the point of the conversation is to collect ideas rather than rule them out.³² Another such example is court proceedings: if one is on a jury, one does not want to update by intersection with the content of what each witness says, but one still wants to update with information *that she said it*. Given the different nature of a discourse in a courtroom,

³²Some people have suggested that things like brainstorming sessions have implicit modals that take scope over the whole discourse (or over each idea). Of course, there are some conversations in which people say things like "we might do x" and "we might do y", and I'm sympathetic to the idea that this is implicit in brainstorming session where it is not explicit. But notice that the static semantic/dynamic pragmatic theorist doesn't have to posit a tacit modal where there is no overt one — pragmatics alone explains what's going on in terms of the conversational updates.

this difference in updating behavior is expected on the pragmatic account.³³ Likewise, there are other sorts of conversations in which one expects one's interlocutor to contradict herself — for example, in speaking to politicians, who are likely to lie or contradict themselves. In this case, update by intersection is not the best strategy since one knows inquiry is not the purpose of the conversation and update by intersection will likely get us into the absurd context — one where all worlds have been ruled out.

Another advantage of the pragmatic explanation of truth-conditional update is that it is a natural extension of the same principles that guide our updating based on non-linguistic events. To borrow an example of Stalnaker's, if a goat walks into the room while we're having a conversation, we're both going to update with that information. That is, we're both going to eliminate worlds in which there are no goats that just walked into the room. This is just a rational thing to do given that we take our experience at face value (e.g., we do not believe we're hallucinating). These are the same principles of rationality that appear to be at work in conversation, so there seems to be no need for divergent explanations of the phenomena.

The same holds for other sorts of updates. In §3.1, I suggested that updates to the context about which objects are under discussion could be explained in terms of the discourse plans and discourse expectations of rational agents engaged in co-operative discourse. One of the advantages of this sort of explanation is that it accounts for deviations from the standard data. For example, in the marble case, it is generally agreed that the infelicitous nine marble example becomes felicitous (or at least greatly improved) by a long pause or a change of speakers in between the two sentences. Dynamic semantics, in giving the phenomenon in question such a rigid, systematic treatment, doesn't have a natural explanation of this data. But the pragmatic planning theory does: it is a lot more natural to expect a change in discourse plans after a long pause or a change in speakers. Long pauses indicate thinking, which is no doubt part of changing one's mind about discourse plans. And interpersonal discourse plans, especially on small things like what object is currently under discussion, are a lot more likely to disagree than

³³It is true that Grice's own account is often criticized as not being able to account for these sort of courtroom cases. But this is based on idiosyncracies of Grice's actual account, that is, that inducing belief in the hearer is the aim of assertion and that all cases of conversation are cases of co-operative activity. As I said earlier, I am understanding pragmatics broadly. In fact, it is precisely *because* such situations are not co-operative that we expect different pragmatic effects in such cases.

intrapersonal discourse plans.

For example, consider the following case from the movie *When Harry met Sally*:

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

This example is perfectly felicitous, which is surprising from a dynamic semantic viewpoint that treats the relationship between the introduction of discourse referent and anaphoric pronoun as semantic. Of course, a dynamic semanticist could say in these cases that the discourse referent is added to the context by accommodation, but this explanation seems *ad hoc*. By contrast, on a pragmatic story, the explanation is a natural extension of the tools already used in the standard case. In this example, it is clear that Harry has a discourse plan to tell Sally about his impending marriage. Marriage is already a topic of conversation at this point, since Sally has just told Harry that she has no interest in marriage at all. But Sally finds Harry extremely annoying and is shocked that anyone would want to marry him, and therefore clearly has a plan to switch the conversation to talk of his fiancée. The long pause indicating Sally thinking, as well as the change in speakers, make the conversational participants (and anyone watching the movie) primed to accept a change in conversational plan. And since it is utterly unambiguous about whom Sally is asking, the pronoun sounds perfectly fine.

In general, in circumstances in which the speaker's intentions are perfectly clear, and it is obvious why the speaker chose the circumlocutory route, examples like the so-called infelicitous marble example are perfectly acceptable. For example, the following is from a real life conversation between myself and a friend of mine. We were at a conference and both of us expected a certain person to be there, but did not want to let on to other people around us who we were talking about. Observing to herself that this person was not present at the conference, we had the following exchange:

- (23) a. Her: I noticed that not everyone is here.
 b. Me: I know. I heard he was going to be late.

Again, these sort of common exceptions to the norm are unexpected on a dynamic account of context update with discourse referents, but not surprising on a view in which context update in general has to do with recognizing

speaker's intentions. Like in the truth-conditional update case, there are similar mechanisms at work in the non-linguistic arena. If a goat walks into the room, we might update the conversational context with a discourse referent for the goat, upon the reasonable assumption that someone in the conversation is going to want to talk about the goat.

The same holds for the case of counterfactual conditionals. As Gillies is well aware, the update mechanism expands the relevant domain in the context, but there is no mechanism for shrinking the domain. The problem is the domain can shrink, and it can shrink a lot more easily and a lot more often than the dynamic semanticist should be comfortable with. For instance, the basketball example I gave above is a case in which considering P-and-Q-worlds does not have the effect of adding P-and-Q-worlds to the domain in any permanent way. As a reminder: we are considering a case in which I actually have average basketball skills and Kobe Bryant did not have casts on his arms yesterday. In this case, the reverse Sobel sequence sounds just fine:

- (24) a. If I had played one-on-one with Kobe Bryant yesterday and he had casts on both arms, I would have won.
 b. But of course, if I had played one-on-one with Kobe Bryant yesterday, I would have lost.

If Gillies wants to account for this, he has to introduce some sort of shrinking mechanism — and it's not clear how to do so in a non-*ad hoc* way on a dynamic semantic account. On the pragmatic story, however, the felicity of the above sequence is completely expected. The pragmatic explanation for *adding* worlds to the global context (i.e. the sort of update that sticks around) is that once possibilities are made salient by the antecedent (or once the presuppositions are accommodated) they cannot be ignored *unless they can be ruled out*. Since Kobe Bryant did not have casts on his arms yesterday, we can safely rule out the worlds in which he does as relevant possibilities.³⁴ Thus the reverse Sobel sequence sounds fine. Since we can't rule out Sophie getting stuck behind a tall person at the parade, those worlds stick around as part of the input context to subsequent (subordinate) counterfactuals. When we can rule out such worlds, reverse Sobel sequences are not infelicitous.

³⁴Moss (2012) makes a similar proposal about shrinking contexts.

Again, this is natural on the pragmatic account, and would be unexpected if the dynamic semantic account was right.³⁵

The examples of the 3 kinds of updates — involving truth conditions, discourse referents, and the presuppositions of counterfactuals — all face similar problems when they are put into the recursive compositional calculus in a dynamic semantic account. They are *too* systematic, resistant to the many exceptions to these update rules. Again, this is not to say that the dynamic semanticist can't offer some mechanism to explain the phenomena I've pointed out. But the dynamic semantic framework does not offer a natural explanation, and the existence of these "exceptions" undermines the motivation for including updates in the recursive compositional calculus. Instead, the data is suggestive of a pragmatic account, which not only has the ability to naturally explain the variance in updates to the context, but actually predicts the precise variance in many cases. This doesn't rule out that there are other good arguments for dynamic semantics out there. I haven't addressed motivations stemming from entailment relations, discourse relations (as in SDRT), or the fact that dynamic semantics can treat declaratives, interrogatives, and imperatives as a single kind of semantic object (i.e. CCPs). But I have undermined one central motivation for dynamic semantics – one that looks at discourse dynamics – and I think the pragmatic nature of context change should generally give us pause when considering whether to adopt a dynamic semantic framework.

References

- Beaver, David. 2001. *Presupposition and Assertion in Dynamic Semantics*. CSLI Publications.
- Creswell, M.J. 1985. *Structured Meanings*. MIT Press.
- von Stechow, Kai. 2001. Counterfactuals in a Dynamic Context. In *Ken Hale: A Life in Language*, MIT Press.
- Gillies, Thony. 2004. Epistemic Conditionals and Conditional Epistemics. *Nous* 38. 585–616.

³⁵There are some cases in which it is not clear whether certain possibilities can be ruled out or not. In general, I think certain kinds of modal and counterfactual discourse involve negotiation between conversational participants as to the relevant worlds.

- Gillies, Thony. 2007. Counterfactual Scorekeeping. *Linguistics and Philosophy* 30. 329–360.
- Groenendijk, Jeroen & Martin Stokhof. 1989. Context and Information in Dynamic Semantics. In H. Bouma & B.A.G. Elsendoorm (eds.), *Working Models of Human Perception*, 457–486. Academic Press.
- Groenendijk, Jeroen & Martin Stokhof. 1999. Meaning in Motion. In K. von Huesinger & U. Egli (eds.), *Reference and Anaphoric Relations*, 47–76. Dordrecht Kluwer.
- Hardt, Daniel. 1999. Dynamic Interpretation of Verb Phrase Ellipsis. *Linguistics and Philosophy* 22. 185–219.
- Heim, Irene. 1982. *The Semantics of Definite and Indefinite Noun Phrases*: University of Massachusetts dissertation.
- Heim, Irene. 1983. On the Projection Problem for Presuppositions. In *Second Annual West Coast Conference on Formal Linguistics*, .
- Kamp, Hans. 1988. Comments on Stalnaker, Belief Attribution and Context. In *Contents of Thought*, University of Arizona Press.
- Kamp, Hans & Uwe Reyle. 1993. *From Discourse to Logic*. Kluwer Academic Publishers.
- King, Jeffrey. 2007. *The Nature and Structure of Content*. Oxford University Press.
- Levinson, Stephen C. 2000. *Presumptive Meanings: The Theory of Generalized Conversational Implicature*. MIT Press.
- Lewis, David. 1972. General Semantics. In Davidson & Harman (eds.), *Semantics of natural language*, Dordrecht Reidel.
- Lewis, David. 1979. Scorekeeping in a Language Game. *Journal of Philosophical Logic* 8(3). 339–359.
- Lewis, Karen. 2012. Discourse Dynamics, Pragmatics, and Indefinites. *Philosophical Studies* 158. 313–342.

- Lewis, Karen. Forthcoming. Speaker's Reference and Anaphoric Pronouns. *Philosophical Perspectives* .
- Lewis, Karen. MS. A pragmatic account of cross-sentential anaphora on embedded indefinites.
- Moss, Sarah. 2012. On the Pragmatics of Counterfactuals. *Noûs* 46. 561–86.
- Roberts, Craige. 2004. Context in Dynamic Interpretation. In *Handbook of Pragmatics*, Blackwell.
- Roberts, Craige. ms. Retrievalability and Incomplete Descriptions.
- Salmon, Nathan. 1986a. *Frege's Puzzle*. MIT Press/Bradford Books.
- Salmon, Nathan. 1986b. Reflexivity. *Notre Dame Journal of Formal Logic* 27. 401–429.
- Salmon, Nathan. 1989a. Illogical Belief. *Philosophical Perspectives* 3.
- Salmon, Nathan. 1989b. Tense and Singular Propositions. In Perry Almog, Wettstein (ed.), *Themes from Kaplan*, Oxford University Press.
- Schlenker, Philippe. 2008. Be Articulate: A Pragmatic Theory of Presupposition Projection. *Theoretical Linguistics* 34(3). 157–212.
- Simons, Mandy. 2010. A Gricean View on Intrusive Implicatures. In Klaus Petrus (ed.), *Meaning and Analysis: New Essays on Grice*, Macmillan.
- Soames, Scott. 1985. Lost Innocence. *Linguistics and Philosophy* 8. 59–71.
- Soames, Scott. 1987. Direct Reference, Propositional Attitudes, and Semantic Content. In *Propositions and Attitudes*, Oxford University Press.
- Soames, Scott. 1989. Semantics and Semantic Competence. *Philosophical Perspectives* 3.
- Stalnaker, Robert. 1978. Assertion. *Syntax and Semantics* 9.
- Stalnaker, Robert. 1998. On the Representation of Context. *Journal of Logic, Language, and Information* 7. 3–19.

- Taylor, Kenneth. 2001. Sex, Breakfast, and Descriptus Interruptus. *Synthese* 128. 45–61.
- Yalcin, Seth. 2011. Nonfactualism about Epistemic Modality. In *Epistemic Modality*, Oxford University Press.
- Yalcin, Seth. 2012. Dynamic Semantics. In *The Routledge Companion to Philosophy of Language*, Routledge.