Felicity and Presupposition Triggers

Michael Glanzberg
michael.glanzberg@utoronto.ca

For Michigan workshop—still work in progress—please do not circulate

Many current pragmatic approaches to presupposition, beginning at least with Stalnaker (1974), are based on a notion of *felicity*. Presuppositions are requirements a sentence places upon a context for an utterance of that sentence to be felicitous. These requirements in turn are induced by the presence of certain *presupposition triggers*, such as lexical items like factive verbs, or constructions like clefts. In this framework, the notion of felicity is left unanalyzed, relying upon speakers' abilities to provide felicity judgments. Likewise the mechanisms by which presupposition triggers lead to requirements on context are left largely unexplored (as most attention has been focused on the study of presupposition projection).

In this paper, I shall begin to explore the notion of felicity and how it relates to presupposition triggers. My discussion will center around the observation that the kind of felicity relevant to presupposition is not an entirely uniform phenomenon. I shall argue this by a closer examination of *in*felicity. Two important subcategories of infelicity will emerge, distinguished by their *discourse status*. Generally, infelicity induces an opportunity for conversational repair. In some cases, we will see, the repair is obligatory, while in others, it is merely optional. I shall argue that this distinction correlates to a division among presupposition triggers. One category, which we may call *strong presupposition*, leads to obligatory repair upon failure. The other, which we may call *weak presupposition*, leads only to repair optional status.

I shall then attempt to explain the distinction between strong and weak presuppositions, by offering an analysis of presupposition triggers in a framework of dynamic intensional logic. I shall argue that a systematic treatment of a wide range of triggers can be given in terms of a single dynamic operator. I shall also show how the dynamic properties of this operator can help to explain the different discourse properties of presuppositions. In dynamic logics, two sorts of *failure* can be distinguished. Update instructions can fail by resulting in a fail state; but instructions can also fail by being undefined for a given input. I shall propose that weak presuppositions result in the former kind of failure, while

strong presuppositions result in the latter.

I shall hence maintain that a more discriminating analysis of presupposition triggers can explain the distinction between strong and weak presuppositions. More generally, I hope this sheds some light on the basic pragmatic notions of presupposition and felicity.

I Presupposition and Felicity

In many discussions, presupposition is explained in terms of a family of embeddings. Consider:

(1) a. It was John who broke the vase.

b. It was not John who broke the vase.

c. If it was John who broke the vase, then Mary will be angry.

d. Someone broke the vase.

It is striking that (1d) is judged to be an implication of each of (1a–1c). This sort of implication under embedding provides a test, the *embedding test*, which identifies (1d) as a presupposition of (1a). [1]

The intuitive picture is that the embedding test tests for a kind of 'backgroundedness'. The reason (1d) is implicated through these embeddings is that it is somehow held in the background, and so is not really affected by the presence of negation or conditionals. The *pre* in presupposition is that it is content which is in the background, or taken for granted, in an utterance of a sentence.

This sort of test, though it is helpful in identifying presuppositions, offers little explanation. Presuppositions are no ordinary implications: implications should hardly be preserved under negation. And the idea that presuppositions are 'in the background' is inviting, but not very precise. We still would like to understand what a presupposition is, and what produces them, even if we have a way to spot them.[2]

[1] Embedding under negation is certainly crucial to identifying presupposition, but different authors often opt for slightly different embeddings. See, for instance, Beaver (1997, 2001), Chierchia and McConnell-Ginet (1990), Kadmon (2001), or Soames (1989).

[2] As Chierchia and McConnell-Ginet (1990) point out, the embedding test can overgenerate. It counts non-restrictive relative clauses as presuppositions, as in:

(i) a. Mary, who was late, took the train.

b. Mary, who was late, did not take the train.

c. If Mary, who was late, took the train, she had a long day.

d. Mary was late.

Intuitively, the content of the non-restrictive relative is not 'in the background'. This has lead some to ask if there is a further 'presupposition property' which could be tested for.

One of the leading theories, deriving from Stalnaker (1974) and Karttunen (1974), explains presupposition as a *pragmatic* phenomenon. Presuppositions arise from what Stalnaker calls *presuppositional requirements*. These are requirements a sentence places upon a *context* for the use of the sentence to be *felicitous* in the context. According to this view, presupposition is pragmatic in that it is a matter of the behavior of sentences in contexts.

The pragmatic theory can be fleshed out a little more by saying something further about context. We think of a context as an information state. At first pass, we can think of it as a set of propositions taken to be common ground among speakers in a conversation. If we think of these propositions as sets of worlds, we can think of the context as their intersection: the set of worlds compatible with what is common ground in the conversation. With this aspect of context in place, we can generally think of requirements placed upon context as requirements that the context *entails* some proposition. If a context c and proposition p are sets of worlds, we require that $c \subseteq p$. (Below, we shall complicate this picture somewhat, by adding discourse referents as well as worlds, but the basic idea will remain.) Presuppositional requirements are thus relations between a sentence s, a context c, and a proposition p such that for s to be felicitous in c, c must entail p.[3]

This kind of requirement generates both the behavior we see in the embedding test, and offers an explanation of the 'backgroundedness' of presuppositions. If a sentence carries a presuppositional requirement p, any context in which it us felicitous will already make p true. Hence, we have an implication from s to p. But such requirements are not part of the asserted content, and so are not sensitive to the embeddings of the embedding test. For instance, if the use of a demonstrative requires the context to entail that there is a demonstrated individual, so does the use of a demonstrative under negation. Presuppositions are background as they are not content asserted by the use of a sentence *per se*; rather, they are the way the context is already required to be for the use of the sentence to be successful. In this way, they are taken to be 'background'.

This picture is in some ways quite elegant. It offers an explanation of the basic features of presupposition, and many researchers have attempted to use it as a basis for explaining presupposition projection. It is striking, though, that it makes crucial use of an unanalyzed notion of *felicity*.

To some extent, this may not be such a surprise. The basic data for a great deal of research in pragmatics are felicity judgments, understood as judgments of the acceptability or unacceptability of

[3] I am suppressing a number of issues here that I do not believe matter for what follows. Perhaps most important, identifying the relevant aspect of context for presupposition as common ground is controversial. Stalnaker himself analyzes this aspect of context in terms of an attitude of speaker presupposition.

a sentence for use in a given context. And indeed, some cases of presupposition give us quite clear judgments of infelicity. Consider:

(2) a. That palm tree is about to fall.

b. Context: no palm tree available.

(3) a. It was John who solved the problem.

b. Context: it is known that the problem has not been solved.

(4) a. John regrets voting for Bush.

b. Context: it is known that John did not vote.

Speakers will find these sentences markedly unacceptable in the contexts described.

Even so, we should be aware that in working with presupposition, we are as much identifying a theoretically important notion of felicity as we are simply gathering data. The notion of felicity itself, and judgments of felicity or infelicity, can be highly plastic. Originally, in the hand of Austin (1975), *felicity* meant something much more broad than what is relevant for presupposition: a very general category of inappropriateness reflecting practical or social concerns as much as linguistic ones.[4] Likewise, many speakers would have as firm judgments that it is unacceptable to say *you old cow* to the Queen as it is to say *that palm tree* when there is none.

In the case of presupposition, it is a persistent idea that there is a single, uniform phenomenon at issue. Hence, if we pursue the pragmatic view, there must be a single, uniform notion of felicity which goes with it. Though I shall not question that presupposition is a basic linguistic concept, I shall argue below that the notion of felicity that goes with it is not uniform. Presuppositions and the notion of felicity that goes with them come in significantly different kinds.

II Infelicity and Discourse

Even if we accept Stalnaker's basic pragmatic account of presupposition, two issues remain. One, as I noted, is that presupposition appears to be correlated with a specific notion of felicity. The other is that Stalnaker's framework is highly general, and does not really strive to explain how presuppositional requirements are introduced. These issues go together. Understanding how known presuppositional requirements are introduced can tell us something about what sort of infelicity we might encounter

[4]The *Oxford English Dictionary* gives as its second entry for *infelicity* "2. The quality of not being happily suited to the occasion or circumstances; unlucky inaptness or inappropriateness; with pl. an unhappily inappropriate expression or detail of style."

when they are not met, while understanding what sorts of infelicity are involved in presupposition could tell us something about the range of genuine presuppositional requirements.

In this section, I shall point out that the infelicities that go with presupposition failure are not uniform after all. I shall suggest they fall into two categories, and I shall offer some diagnostics which help us spot which presuppositions go with which. In the following sections, I shall look more closely at some of these presuppositions, and try to explain why they result in the categories of infelicity they do.

To illustrate the two categories into which presuppositions fall, let us look at two typical presuppositions. One is a demonstrative like (2). The other is the presupposition triggered by *even* in a sentence like *Even JOHN solved the problem*. This carries a presupposition: roughly, that someone other than John solved the problem, and that it was unlikely or unexpected that John did (by some contextually salient scale). Both the embedding test and the intuitive idea of 'background' support that *even* triggers some such presupposition. Now, consider cases where these fail.

(5) a. i. That palm tree is about to fall.

ii. Context: no salient palm tree.

b. i. Even JOHN solved the problem.

ii. Context: assumed that John was most likely to solve the problem.

I think we judge both (5a) and (5b) to be infelicitous. But not in the same way. The problem in (5a) is what philosophers have often described as failure to express a proposition, or failure to have a truth value. Though we have infelicity in (5b), we are not inclined to describe it in this way. It is easy to give a truth-value judgment. It is true (if John really did solve the problem).

This already shows that one long tradition in thinking about presupposition is not adequate to the range of empirical phenomena the embedding test (and related intuitions about backgroundedness) bring to light. It is often proposed that presupposition failure leads to some kind of truth-value gap, or to failure to express a truth-evaluable proposition. Likewise, it is held that presupposition failure is to be detected by the absence of truth-value judgments.[5] Examples like (5b) cast doubt on this analysis.

Actually, there was already reason to doubt this analysis. It is true that in cases like (5a), we often see a marked unwillingness to provide truth-value judgments. (As Strawson (1964) put it, we are "squeamish" about offering truth-value judgments in these cases.) But it has been fairly well

[5] This idea goes back at least to Frege (1892) and Strawson (1950, 1952)

established that there are a number of ways that assessments of assertions can be given for cases like these; assessments which certainly look like truth-value judgments. One fairly typical one is a pattern of what is sometimes called 'presupposition-canceling negation'. Consider a case of a definite description. As has been much discussed (e.g. Horn, 1989; von Fintel, forthcoming), one can reply to (6a) with (6b) in:

(6) a. The King of Russia is coming to visit.

b. The King of Russia is NOT coming to visit—there is no King of Russia!

Though it is somewhat more difficult to hear, this pattern extends to other presuppositions. Consider:

(7) a. That palm tree is going to fall. George said so.

b. That palm tree is NOT going to fall—there in no palm tree.

With the right intonation pattern, with stress on the negation, and a kind of stress on each each word of the original sentence, it does seem we can say this.[6]

Another problematic case is what we might call 'repair-to-negation':

(8) a. Is that palm tree going to fall?

b. Er . . . no . . . , there is no palm tree.

In cases like this, we note a defect in an utterance, and reach a negative judgment on the basis of it. In (8), it may be as much that we are rejecting some implication of the attempted claim, such as that we are in danger of being hit. But we can also get rejection of what may *look like* what is said itself. For instance, we have:

(9) a. This pen [demonstrating a pen] is in danger of being struck by that palm tree.

b. Well, there is no palm tree, so I guess not.

The speaker here manages to reason from the defect in the utterance to its somehow not being able to be true, and gives a negative assessment.[7]

Initially I offered the contrast in (5), and suggested it showed a subdivision with presuppositional phenomena. Cases like (5a) appear to go with lack of truth value, whereas those like (5b) do not. But

[6] I am inclined, with Horn (1989), to see these as cases of metalinguistic negation. But regardless, they undermine the simple claim that in cases of presupposition failure, we will find speakers unwilling to offer anything that looks like a truth-value judgment. For this point, we need not decide if the phenomenon here is one of metalinguistic negation or something else.

[7] Some of the issues surrounding cases like (9) are discussed in von Fintel (forthcoming), Lasersohn (1993), Reinhart (1981), and in my (2002).

in casting doubt on the account of infelicity by lack of truth-value for judgments for (5a), have I cast doubt on their being a genuine distinction at all?

I believe there is something that can be salvaged from the idea that some presuppositions lead to lack of truth-value judgments, and it will help us to get a better understanding of the difference between (5a) and (5b).

To see what can be salvaged, it will be useful to turn out attention to the notion of a *conversational repair*. A repair is a conversational move of fixing a problem. When we see willingness to give truth-value assessments in cases like (6–9), we see it in conjunction with the initiation of a conversational repair. This is most evident in cases like the repair-to-negation cases (8) and (9). As we see there, repairs have some more or less typical markers, such as certain discourse particles (I put *er* in (8)). This is characteristic of what conversational analysts call 'other-initiated repair', as are pattern of pausing between turns, and other markers.[8] It appears that the intonation pattern in (6) and (7) also marks repair.

The characteristic mark of cases like (5a) is that in certain settings, they induce to an *obligatory* repair. In contrast, cases like (5b) only induce an *optional* repair. The right sort of obligatory repair is what often appears to us as lack of truth value.

We may think of truth-value judgments as occurring in discourse as responses to yes-no questions like (10a), or corresponding queries about assertions like (10b):

(10) a. Is that palm tree going to fall?

b. i. That palm tree is going to fall.

ii. Is that right?

The Strawsonsian idea that cases like (5a) lead to lack of truth value can be put that speakers cannot answer such queries. We have seen this is not quite right. They can, but in doing so, they must initiate a repair.

The sense in which speakers *must* initiate a repair is somewhat elusive. Discourses are messy things, and speakers are capable of putting up with all kinds of defects in them. It is often easier to talk around a problem than fix it. The trick for spotting the infelicity which goes with (5a) is to construct discourse settings that limit these options. Considering such settings will give us some rough-and-ready tests for the sort of infelicity which goes with cases like (5a).

[8] See Schegloff *et al.* (1977) and Levinson (1983).

One such test is what I shall call the *echo-assessment* test. In this, we ask speakers not only to make an assessment, but to attempt it with the same words as the initial assertion (or as close as we can come, modifying, for instance, occurrences of terms like *I* and *you*). So, for instance, we have no trouble with:

(11) a. Is Al Gore president?

b. No, Al Gore is not president.

On other hand, what is correct about the Strawsonsian observation is that speakers will not offer:[9]

(12) a. Is that palm tree about to fall?

b. # No, that palm tree is not about to fall.

In cases like this is speakers will only give an assessment by initiating a repair, as in (7) or (8). When they do, they typically tend to avoid echoing the defective construction, as in (8), or at least use it only in a marked way, as in (7). We do not have:

(13) # Er . . . there is not palm tree, so I guess that palm tree is not about to fall.

In defective cases, the echo-assessment test asks speakers to do something which would reproduce the defect. If repair of the defect is required, speakers will not do this. Rather, if they make any assessment, they will initiate a repair and work around the defective construction. Thus, we can use the echo-assessment test to spot cases where repair is necessary.

A second diagnostic for the need for repair is the *indirect speech report test*. In cases like (5a), speakers will be unwilling to provide *indirect* speech reports without initiating a repair. Normally, the repair is given by offering a direct quotation instead. Compare:

(14) a. # George said that that palm tree is going to fall.

b. George uttered 'That palm tree is going to fall', but there is no palm tree.

As with the echo-assessment test, this test asks speakers to do something which would reproduce a defect. If the defect requires repair, speakers will not be willing to do so, or at least not unless they can do so in a way that also makes the needed repair.[10] I shall group the echo-assessment and indirect speech report tests as the *repair tests*.

[9] I mark discourse unacceptability by '#'.

[10] The repairs made by shifting to direct quotation are quite weak, amounting to marking what is defective about an utterance, and moving on. A more elaborate repair might be provided by the original speaker in the next turn. They might say something like 'oh, I meant . . . '. But this does not appear to be required, once the initial repair is made.

The repair tests distinguish between (5a) and (5b). They show (5a) and not (5b) to be a case of obligatory repair. We get:

(15) a. Even John solved the problem.

b. Yes, John did . . . but why did you say 'even'?

c. # That's NOT SO. He would have solved it if anyone did.

The indirect speech report test supports this:

(16) a. Even I solved the problem. (Said by John.)

b. John said that even he solved the problem . . . but of course, that's a bid odd, as he would have if anyone did.

c. # John said 'Even I solved the problem', but that doesn't make sense, because he was most likely to have done it.

Repair might be initiated, depending on what other premiums there are in the conversation. Repair here is *optional*, whereas in cases like (5a) it is *obligatory*.[11]

Generally, any infelicity will correspond to an occasion where a repair is possible. We should expect presupposition failure to correspond to a certain sort of repair. Not a repair of syntax, or of social awkwardness, but of whatever led to a requirement on context not being met. I have not proposed any sort of account of what that category is. But I have noted that it is not entirely homogeneous. Within it, there is a significant subdivision between infelicities corresponding to obligatory repair, and those corresponding merely to optional repair. There are at least two different ways for a presupposition to go wrong.

III Strong and Weak Presuppositions

We have no seen a subdivision within infelicities based on their discourse status: some lead to obligatory repair, and some merely to optional repair. We have seen in (5) that failures of presuppositions can fall into either category.

[11] There are quite a few ways in which the repair tests are difficult to apply. If they are run in contexts where there is a great premium on not being misleading, they will over-generate, counting utterances with false implicatures as like the presupposition failure of (5a). The tests are thus at best rough-and-ready diagnostics. It is best to think of them as being applied in the 'courtroom' setting, where the premium is only on speaking correctly, with no regard for whether was is said is misleading or evasive. We also always need to assume that the tests are applied when there are no performance errors of syntax or phonology requiring a different sort of repair. Especially for the echo-assessment test, we need to assume speakers know enough about the relevant facts, and enough about the context to know, for instance, the referents of indexicals.

The division within infelicities thus results in a distinction between different sorts of presuppositions. Presuppositions have *triggers*: lexical items or constructions the presence of which induce particular presuppositional requirement for sentences.[12] The different categories of infelicity serve to classify presupposition triggers into two groups, depending on whether they lead to obligatory repair upon presupposition failure, or to merely optional repair. We have already looked at two triggers: demonstrative NPs and *even*. We have seen that they fall into different groups. Demonstratives fall into the former group, resulting in obligatory repair when their presuppositional requirements are not met; while *even* falls into the latter, resulting in optional repair. Let us call the former *strong presuppositions*, and the latter *weak presuppositions*. A strong presupposition trigger induces a requirement which leads to obligatory repair upon failure, while a weak presupposition induces a requirement leading only to optional repair.

Running the repair tests, we can list some examples of strong and weak presuppositions:

(17) **Strong presuppositions.**

a. Demonstrative NPs: That palm tree is about to fall.

- Presupposition: Contextually available value of *that palm tree*.
- Trigger: Demonstrative NP *that palm tree*.

b. Clefts: It was John who solved the problem

- Presupposition: Someone solved the problem.
- Trigger: Structure of cleft.

c. Factives: John regrets voting for Bush.

- Presupposition: John voted for Bush.
- Trigger: Lexical item *regret*.

(18) **Weak presuppositions.**

a. Focus-sensitive particles: Even John solved the problem.

- Presupposition: Someone other than John solved the problem, and it was unlikely or unexpected that John did.
- Trigger: Lexical item *even*.

[12] Of course, to really know what presuppositions a sentence carries based on what presupposition triggers it contains, one would have to know how presuppositions project. For the most part, I shall concentrate on elementary presuppositions—presuppositions of simple sentences—and so ignore the issue of projection.

b. Iteratives: John solved the problem too.

- Presupposition: Someone other than John solved the problem.
- Trigger, lexical item *too*.

This is by no means an exhaustive list, but it includes the presuppositions I shall discuss further, and is a reasonable sample of the range of presuppositional phenomena.[13]

We have already applied the repair tests to demonstratives and *even*. To check on more example, we can observe that clefts indeed are strong presuppositions. Consider a context in which the presupposition of (17b) cannot be accommodated, such as one in which we are discussing the fact that no one has solved the problem. In such a context, we find repair obligatory:

(19) a. The problem is unsolvable.

b. It was John who solved the problem.

c. Is that right?

d. Err . . . no one solved the problem (as we all know).

No assessment or reporting of content seems possible until we repair the utterance.

We now have a taxonomy of presupposition triggers into strong and weak. The taxonomy is based on the results of presupposition failure. Strong presuppositions lead to obligatory repair upon failure, while weak ones lead to optional repair.

IV Analyzing Presupposition Triggers

So far, I have pointed out some issues left open by the pragmatic view of presupposition in Section (I), and I have done some taxonomy that bears on it in Sections (II) and (III). The Stalnaker-inspired approach to presupposition leaves open just what the phenomenon of felicity is, and it also leaves

[13] The factive case raises a few additional issues. I follow Stalnaker (1974) and Chierchia and McConnell-Ginet (1990) in taking factive presuppositions to be conversationally generated, and hence contextually cancelable. I shall have very little to say about factives here. But let me note that the contextual cancelability of these presuppositions can create cases where repair is not obligatory, and does not even seem to be optional. As Simons (2001, ms) pointed out, certain sorts of ignorance contexts seem to cancel these sorts of presuppositions. Suppose we have:

(i) a. I wonder what's going on with John.
b. Well, he doesn't seem to regret voting for Bush.
c. That's right (he didn't vote for him at all).

I am inclined to think the factive presupposition is simply canceled in this context, so we do not have presupposition failure at all. The continuation of the last sentence seems to be entirely optional; but it does not appear to be a repair.

open the way that presuppositional requirements are induced. The taxonomy of strong and weak presuppositions simply observes that different presupposition triggers lead to different sorts of felicity properties. In this section, I shall begin to examine how we might explain the way presuppositional requirements arise. This will produce an explanation of why we see the distinction between strong and weak presuppositions.

My proposal will be there is a single sort of operator which induces presuppositional requirements. Using this operator, we will be able to analyze the presupposition triggers discussed in Section (III). I shall offer a more formal analysis of it in the next section. Here, I shall sketch the proposal in more informal terms.

The basic idea is that presuppositions arise from the presence of an operator which *tests* a context to see if it satisfies some requirements. We can put this idea naturally in terms of the view of assertion as context-update in Stalnaker (1978) (a precursor to the sort of dynamic semantics I shall employ below). We think of the goal of assertion as conveying information. Information is conveyed against a background of shared information. This information comprises the context—the set of worlds compatible with what is common ground in the conversation. Thus, to convey information is to add to what is common ground. This updates the context, by in effect eliminating worlds from it.

We can thus think of the semantics of a sentence as encoding *update instructions*. Conveying information linguistically amounts to processing these instructions. In the simplest cases, the update instruction is simply to throw out worlds not compatible with what the sentence says. Writing $c[\phi]$ for the update of c by instruction $[\phi]$, we have in the simplest cases $c[\phi] = \{w \in c \mid w \vDash \phi\}$. We will see more complex sorts of update instructions in a moment.[14]

One sort of update instruction is what is sometimes called a *test*. A test is an instruction to check and see if something holds. If it does, proceed, if not, fail. This stops an update procedure from proceeding if some condition is not met. One thing we might want to do in an update procedure is test a context to see if it meets some condition. If it does, we proceed, without modifying the context at all. If it does not, the computation fails.

Stalnaker's notion of pragmatic presupposition already suggests a way that presuppositions test contexts. They look to see if a context entails some proposition. For a given proposition p, we execute this sort of test simply by looking into the context c and checking that $c \subseteq p$. If it is, the output is

[14] I am trying to remain neutral on whether the basic semantic notions are in some fundamental way 'dynamic', as partisans of dynamic semantics sometimes hold. I only require that the dynamic picture give us a good explanation of what happens in successful assertion, and so be faithful to the phenomena of felicity and presupposition. This is compatible with a more 'static' view of the underlying semantics of a sentence.

c, if it is not, the output is a fail state. But the idea of looking into the context and checking to see if it contains something we may need is not really restricted to propositions. A demonstrative, for instance, may require us to first check a context so see if there is a salient demonstrated individual. If there is, we go on, using it in our computations. If not, we fail.

Let me introduce some notation. We can write the operator for testing a context by checking to see if it contains something satisfying some property F as $\downarrow xF(x)$ (to be read *find and* x *such that* $F(x)$ *in the context*). The update instruction of $\downarrow$ looks like:

(20)
$$c[\downarrow xF(x)] = \begin{cases} c & \text{if } c \text{ contains an } x \text{ such that } F(x) \\ fail & \text{otherwise} \end{cases}$$

I am being rather informal about what it is for a context to 'contain' something. This will be developed more later. For the moment we can observe that in the propositional case:

(21)
$$c[\downarrow x(x = p)] = \begin{cases} c & \text{if } c \subseteq p \\ fail & \text{otherwise} \end{cases}$$

My proposal is that instructions in the scope of a $\downarrow$ operator are presuppositions, and specific presupposition triggers can be analyzed in terms of $\downarrow$.

Many of the important aspects of presupposition follow nicely on this proposal. The 'backgroundedness' effect of presupposition is accounted for in very simple terms. $\downarrow$ has no effect on the context if it succeeds; rather, it only verifies that the context is already some way. This amounts to noting what is already in the background.[15]

Likewise the basic embedding properties of presuppositions also follow from the behavior of $\downarrow$. For instance, using the usual rule for dynamic negation, we see $c[\neg[\downarrow x(x = p)][\phi]] = c \setminus (c[\downarrow x(x = p)][\phi]) = c \setminus (c[\phi]) \subseteq p$ if $\downarrow$ instruction succeeds.

Most importantly, if we analyze the sources of presuppositions in terms of $\downarrow$, we get a direct explanation of why presuppositions factor into strong and weak. In the next section, I shall look

[15] I hasten to add that something needs to be said about so-called 'informative presuppositions'. I am not sure just what the right approach is, but my current inclination is to simply see it a the result of a Gricean process of accommodation. If a $\downarrow$ instruction fails, as when we encounter any conversational problem, it is open to attempt to fix the context to avoid the problem. This amounts to *accommodation*, which adds the needed information to the context 'on the fly', and proceeds as if the context had been acceptable. It would be possible to write this option directly into the semantics of $\downarrow$, but I am inclined to think it is already covered by general Gricean principles.

at a couple of presupposition triggers in more detail. But for the moment, let me sketch how the explanation will go.

Let us start with a factive presupposition like that of *regret*. Consider:

(22) John regrets voting for Bush.

We can think of *regret* as breaking down into a factive component and an 'attitudinal' component. j regrets p iff j takes a negative propositional attitude towards p (something like being sad, but note, *sad* is itself factive), and p is true. The latter component is the presupposition, and so should be in the scope of a ↓ operator. We thus have something like:

(23) a. c[John regrets voting for Bush]

b. $c[\downarrow x(x = p)][S(j, x)]$

Intuitively the update instructions are: check the context to see that it contains p; John is sad about *that*.[16]

In (23b), the occurrence of x in the second instruction is in effect 'bound' by ↓. In the jargon, ↓ is a dynamic binder. I shall work out the dynamic semantics for ↓ in the next section. Intuitively, a test like ↓ sets up a *discourse referent* which can be a target of anaphora. In checking the context to make sure it entails p, we thereby make p appropriately salient for discourse anaphora. We set it up as a discourse referent, which the later occurrence of x picks out. The first instruction tells us to find p in the context. The next says we are to update the context with the information that John is sad about *it*.

I thus need to be a little more careful about my claim that a presupposition tests the context, but does not change it. In checking to see that something is already available in a context, it can be raised to a higher degree of salience, making it a target for anaphoric reference. Checking a context to see that something is in it can set up a discourse referent. But we still do not want a presupposition to change the basic information state of the context. It should only make a target something we already had, and not change the information we had about it. In the fuller treatment to follow, where we think of information states as combinations of worlds and discourse referents, we will insist that presuppositions are propositional test in that they do not change the worlds of the context, even though they will often change its system of discourse referents.[17]

[16] As I hold that factive presuppositions are conversationally generated, I owe a story about how conversational processes can insert a ↓ operator. I have tried to tell this story in another place. For the moment, I just want to illustrate the ways ↓ can be applied.

[17] The relation between presupposition and anaphora is highlighted by van der Sandt (1992). I am unsure to what extent the idea I am proposing here fits in with his. Van der Sandt's basic idea is that a presupposition is an anaphor looking for an

It is the possibility this sort of dynamic behavior within an elementary presupposition, setting up a discourse referent which is used in the processing of a simple sentence, which is crucial for explaining the strong/weak contrast. To see why, let us look briefly at what the update instructions for *even* should look like. As I mentioned at (5b), *Even JOHN solved the problem* carries the presupposition that it was unlikely or unexpected that John solved the problem, and that someone other than John solved it. For the moment, just abbreviate this as p (I shall discuss this more thoroughly in the next section). The update instructions then look like:

(24) a. c[Even John solved the problem]

b. $c[\downarrow x(x = p)][R(j)]$

There is a crucial difference between this and the instructions in (23). In (23), the second instruction contains a variable bound by $\downarrow$. In (24), the second instruction contains no such variable. Indeed, nothing about the first instruction is relevant to the processing of the second.

This leads to an important difference if the presuppositions of the two fail. Suppose we have a presupposition failure for (24). The result of processing the first instruction is a fail state. There is a natural discourse strategy to apply in such a case. Reset to the prior information state, and try to go on to process the rest of the discourse. This is especially natural for failed $\downarrow$-instructions, as if our check of the context fails, we might still try to see what we were to do with it. With (24), this works perfectly well. After resetting the context after the failed $\downarrow$, we can perfectly easily process $[R(j)]$, updating to only worlds in which John solved the problem. We have an infelicity, as there was a fault in our update computation. But it is an infelicity we can work around in processing the utterance, and no further repair is required. We have optional repair, and a weak presupposition.

antecedent. Presupposition projection then looks like an anaphora resolution algorithm. This explains, for instance, why *If France has a King then the King of France is bald* places no requirement on the context. The presupposition of *the King of France* is in effect bound by the antecedent.

This may make my idea the opposite of van der Sandt's. My presuppositions set up discourse referents, whereas his are anaphors looking backwards for antecedents. But to some extent, he and I are after different things. Van der Sandt is primarily concerned with *projection*, whereas my concern is mostly with the behavior of elementary presuppositions. When it comes to satisfying elementary presuppositions in contexts, we could bring the two closer together by thinking of the presupposition as a van der Sandtian anaphor, and think of the pragmatics of satisfaction as incorporating a $\downarrow$-like operation. An elementary presupposition would look like $P(x)$, which would be satisfied if the context supports $\downarrow xP(x)$. (An extensive discussion of how to bring the theory of van der Sandt (1992) closer to the satisfaction theory of Heim (1983) is given in Zeevat (1992).) But I think the problem of strong vs. weak presuppositions, and the analysis of presupposition triggers I am offering to solve it, gives us good reason to put $\downarrow$ in the semantics. This leaves the issue of projection untouched, and the issue of whether some sort of local satisfaction would suffice, or if we would need to modify the account of $\downarrow$, *a la* van der Sandt, to allow first for resolution and then for accommodation, remains open.

I should mention one other respect in which my suggestion is perhaps more van der Sandtian than it might first appear. Though my basic model is one of presupposition satisfaction, it departs in one crucial respect from many of the current satisfactional theories. Many of these theories incorporate some partiality into their semantics, either by making the context-update function partial, or making the update operation a partial operation (e.g. Beaver, 2001; van Eijck, 1993, 1994; Heim, 1983; Krahmer, 1008); Examples like (5b) cast doubt on the adequacy of any theory which builds partiality into the semantics of elementary presupposition across the board, making a discourse-based theory far more attractive.

In contrast, suppose were to attempt this strategy for (23). Again, we find a failed ↓-instruction. We rest the context, and then attempt to process $[S(j,x)]$. But this does not work. The update $[S(j,x)]$ is not even defined, as the context fixes no discourse referent for x. The failure of the ↓-instruction is precisely a failure to provide such a value. So, when we attempt to process the second instruction, we get a strong form of failure. It is not merely that processing $[S(j,x)]$ results in a fail state, it is *undefined*. Thus, we have no way to compute a content, unless we make some repair. We have obligatory repair, and a strong presupposition.

V Formal Development

The explanation of the difference between strong and weak presuppositions I offered in Section (IV) needs to be fleshed out in two ways. We need a fuller explanation of how ↓ works, and we need a more thorough analysis of some presupposition triggers to be convinced it is really responsible for presuppositions. I shall do some of both in this section. I shall sketch the semantics for a dynamic logic suitable for modeling the kinds of context updated discussed above (basically that of Groenendijk *et al.* (1996)), and offer a definition of ↓ in it. I shall then look in detail at a couple of presupposition triggers which exemplify the strong/weak contrast. My primary example of a strong presupposition will be the *cleft* construction, while my primary example of a weak presupposition will be *too*. These can be treated in a first-order framework without additional modification. I will briefly comment on how to extend the system with propositional variables, allowing for treatment of *even* and factives. I will also comment briefly on the difference between the presuppositions of demonstratives and clefts (both strong, but importantly different in other ways).

With this, I hope to flesh out an explanation of the strong/weak distinction, but also to begin the more general project of analyzing presupposition triggers and felicity.

V.1 Dynamic Semantics

We need a semantics which handles discourse referents and models context-update by eliminating possibilities. Hence, we need a dynamic modal system. The one I shall sketch here is a modest simplification of the one from Groenendijk *et al.* (1996)).

First, we need to model information states—contexts. We roughly want to think of these states as providing a set of common ground propositions, and a domain of salient objects as well. We will model propositions as set of world, so we may begin with a set W of worlds, and a domain of discourse D. D

is the domain of contextually salient individuals, and we can pretend all worlds in W have universe D.

Possibilities, as usual in dynamic semantics, are going to be thought of as combinations of a world and an appropriate assignment of values in D to variables. So, we can define a *system of discourse referents* r to be a (finite) partial function from the set of variables to D. A *possibility* is a pair $\langle r, w\rangle$ where r is a system of discourse referents and $w \in W$.

A context is an instance of an information state. Initially, we thought of an information state as a set of possible worlds (representing a collection of propositions). We now have to be a little more careful about discourse referents. An information state should have a coherent system of discourse referents. Hence, an *information state* is a set of possibilities s such that if $i, i' \in s$, then i and i' have systems of discourse referents with the same domain.[18] (Observe, two possibilities in an information state can disagree on the value of a discourse anaphor, but they must agree on whether it is well-defined.)

Information states can change both by changing the worlds in them, and changing their system of discourse referents. To compare information states as discourse referents change, we need a few definitions. For possibilities $i = \langle r, w\rangle$, $i' = \langle r', w'\rangle$, define the relation of *extension* between possibilities:

$$i \leq i' \text{ iff } r \subseteq r' \text{ and } w = w'.$$

We can then define *extension* between information states:

$$s \leq s' \text{ iff } \forall i' \in s' \exists i \in s (i \leq i').$$

An extension of s can add discourse referents and can eliminate possibilities.

We will be especially interested in changes of information states that only affect discourse referents. Define s *subsists* in s' iff $s \leq s'$ and every possibility in s has an extension in s'. If s subsists in s', s' changes s only by adding discourse referents.

Most of the clauses for context-update are standard. Update operations are applied to information states. For instance, writing $s[F(t)]$ for the update of information state s by $F(t)$, we have $s[F(t)] = \{i \in s \mid i \vDash F(t)\}$.

Dynamic effects, as usual, are generated by the existential quantifier. Define $r[x/d]$ to be the

[18] Groenendijk *et al.* (1996) use a peg system, separating systems of discourse referents from the values they take in worlds. As I shall not make use of that extra complexity, I have simplified slightly.

referent system that agrees with r, except $dom(r[x]) = dom(r) \cup \{x\}$ and $r[x/d](x) = d$. Let $i[x/d] = \langle r[x/d], w\rangle$, and $s[x/d] = \{i[x/d] \mid i \in s\}$. The existential quantifier is then defined by:

$$s[\exists x\phi] = \cup_{d\in D}(s[x/d][\phi]).$$

$\exists$ can set up genuinely new discourse referents.

We now need to define the $\downarrow$-operator. For the moment, let us concentrate on the first-order version. $\downarrow xF(x)$ instructs us to find x in the context if it is there. If it is, we set up x as a discourse referent, but otherwise, do nothing to the context. We are testing the context to see if it contains x, and the success of the test introduces the discourse referent but otherwise adds no content. We thus expect $\downarrow$ to look in some ways like the dynamic $\exists$. But in not affecting the context in any other way, $\downarrow$ looks more like a modal update. Modal updates are given by:

$$s[\Diamond\phi] = \{i \in s \mid s[\phi] \neq \emptyset\}$$

$$s[\Box\phi] = \{i \in s \mid s \text{ subsists in } s[\phi]\}.$$

Modal updates do not change the information state if they are true, and return the empty set if they are false. They are tests of the information state.

The semantics of $\downarrow$ combines both aspects:

(25)
$$\downarrow x\phi(x) = \begin{cases} \cup_{d\in D}(s[x/d][\phi]) & \text{if } s \text{ subsists in } \cup_{d\in D}(s[x/d][\phi]) \\ \emptyset & \text{otherwise} \end{cases}$$

This works like $\exists$, setting up a discourse referent which can take as value in the context anything which is F, so long as the context guarantees that something is F in every open possibility.[19] This is the dynamic notion of *finding* an x in the context, and setting up a discourse referent for it, just as we wanted.[20]

[19]Recall the basic fact from dynamic logic that $\exists xF(x) \longrightarrow G(x)$ is equivalent to $\forall x(F(x) \longrightarrow G(x))$, which made its application to donkey anaphora so inviting.

[20]Intuitively, this seems a lot like testing $\Box\exists xF(x)$. It cannot quite be that, as in the kind of dynamic framework we have adopted, modal operators block dynamic binding. $\downarrow$ is thus something of a hybrid, combining aspects of a dynamic quantifier and a modal operator. I have some misgivings about this—it would be nice to have a framework in which it appears more natural. Perhaps the framework of Kohlhase *et al.* (MS), which separates dynamic and λ-binding might do a better job.

V.2 Fail States

The definition of $\downarrow$ outputs $\varnothing$ if the presupposition it introduces fails in a context.[21] It is the custom in dynamic logic to treat $\varnothing$ as a *fail* state, and I have followed suit.

It should be noted that thinking about the pragmatics of assertion provides independent grounds for counting $\varnothing$ as a fail state. It is not just a matter of the formalism of dynamic logic. One of the conditions on successful assertion in Stalnaker (1978) is that assertions should be informative, which amounts to eliminating worlds from a context, but not reducing the context to $\varnothing$. $\varnothing$ is a contradictory information state, and so cannot be an informative update.[22]

For this reason, we should expect updates resulting in the fail state $\varnothing$ to be infelicitous. They are inappropriate on Gricean or Stalnakerian grounds. We can see that a failed presupposition will at least have this result. Hence, we can expect failed (unaccommodated) presuppositions to yield infelicity, just as the Stalnakerian pragmatic picture holds.

As I mentioned in Section (IV), there is another sort of failure we can see with presuppositions (and with dynamic operators generally). Consider an information state s whose system of discourse referents is not defined on x. Then an update like $s[G(x)]$ is *undefined*. It is not that it runs to failure—it cannot start. As I mentioned a moment ago, failed $\downarrow$ instructions can lead to this situation, even if we apply the fault-tolerant strategy of resetting the context to its prior state if an update results in $\varnothing$. For instance, consider an update like:

(26) $s[\downarrow xF(x)][G(x)]$

If $\downarrow xF(x)$ fails, then $s[\downarrow xF(x)] = \varnothing$. But then the instruction $[G(x)]$ is *undefined*. So long as the system of discourse referents of s is not defined on x, it does not help to invoke the fault-tolerant strategy and reset to s. $s[G(x)]$ is still undefined. This is the situation of strong presupposition failure, as it is the situation in which we cannot process our update instructions at all without first initiating a repair.

V.3 Clefts

Now that we have a more rigorous definition of $\downarrow$, we should apply it in a few cases. My primary example of a strong presupposition trigger will be the cleft construction. This is, I believe, a clear

[21] I hasten again to mention the possibility that accommodation could repair this sort of failure.
[22] For more discussion of fail states and related notions, see van Eijck (1993).

example of a strong presupposition. The judgments for it are quite firm. And, it can be treated in the first-order system I just sketched, without adding additional apparatus for propositional variables.

Let us look at a typical cleft. Our example in (17b), repeated here, was:

(27) It was John who solved the problem.

(27) presupposes that someone solved the problem, and asserts that John and no other contextually salient individual did so. This is the semantics of what É. Kiss (1998) calls *identificational focus*. An identificational focus provides an exhaustive list of contextually salient elements of which a predicate holds.[23]

The exhaustivity of clefts is brought out by the following pair (É. Kiss, 1998):

(28) a. It was a hat and a coat that Mary picked for herself.

b. It was a hat that Mary picked for herself.

These two are incompatible. But exhaustivity is restricted to context. Suppose we are in a room with our colleagues, all of whom are distinguished scholars; but while all of them also strive to be musicians, only John plays well. Compare:

(29) a. It is John who is a good scholar.

b. It is John who is a good musician.

The first is false, as it fails to be correct in its exhaustive classification. The second is true, even though many people in the wider world are much better musicians than John.

We can represent the truth conditions of a singular cleft, for a contextually salient set of individuals D as:

(30) a. It was John who solved the problem.

b. $\forall x \in D(P(x) \longleftrightarrow x = j)$

To make room for plurals, we might re-write this:

(31) $\{x \in D \mid P(x)\} = \{x \in D \mid x = j\}$

[23] I agree with É. Kiss and with Rooth (1999) that the semantics of clefts—identificational focus—is distinct from that of intonational focus in English. It might be that clefts have a strictly stronger semantics, and so might need to incorporate something like the alternative semantics for focus (or some other semantics) over and above the semantics of exhaustive listing. However, many clefts occur with a pitch accent on the clefted constituent (the phrase after *it was*), which raises the possibility that the semantic contributions of clefting and intonational focus are separate (see Delin, 1992, 1995; Prince, 1978). For simplicity, I shall set this issue aside in what follows.

For a plural cleft, we can then have:

(32) a. It was the boys who ate the cake.

b. $\{x \in D \mid C(x)\} = \{x \in D \mid B(x)\}$

This give the truth conditions of a cleft. But of course, they carry presuppositions as well.

Most standard lists of presuppositions list clefts as carrying an existential presupposition. The usual embedding tests confirm this:

(33) a. It was John who solved the problem.

b. It was not John who solved the problem.

c. If it was John who solved the problem, he won the prize.

d. Someone solved the problem.

Given that the semantics of the cleft is that of exhaustive listing, it is tempting to ask if the presupposition is something stronger. Clefts might require context to give us a list of potential candidates, out of which an exhaustive list is then given. I am not entirely sure, but my current thought is that this is too strong. Consider the following dialog:

(34) a. Someone solved the problem, but I have no idea who it might be.

b. I know. It was John who solved it.

Context provides a range of potential candidates, in the usual way that it provides a contextually salient domain of individuals, but I do not see that this is part of the presupposition triggered by the cleft.

To implement semantics of the cleft in the update framework just sketched, we should say something about their structure. I am hesitant to say too much here, for several reasons. Our first-order update theory is not really built to capture the syntax of natural language. Far more importantly, the syntax of clefts remains controversial, and I am no expert in syntax at all. But nonetheless, I would like to say enough about the structure of clefts to get a set of update instructions for them which are not completely absurd by lights of compositional semantics.

To this end, I shall help myself the bare bones of a proposal by É. Kiss (1998), whose claim about the semantics of identificational focus I have already endorsed. In essence, she proposes that syntactically, the clefted constituent occupies the specifier position of a functional projection FP (focus phrase). Her analysis provides a structure like:

(i) $[_{IP}$ It was $[_{FP}$ John$_i$ $[_{CP}$ who$_i$ $[_{IP}$ t$_i$ solved the problem]]]]

But she also suggests that in the underlying semantics, the clefted constituent, occupying spec-FP in the syntax, expresses the *value* of an operator expressing exhaustive identification. Writing this operator EXH, we might have something like:

(35) a. [[John$_i$]$_{F^{ID}}$ [who$_i$ t$_i$ solved the problem]]

b. John = EXH_i[t$_i$ solved the problem]

Semantically, we have the same structure in:

(36) a. It was John who solved the problem.

b. John is the one who solved the problem.

This yields the truth conditions given in (30).[24]

We need to explain how this sort of structure results in a strong presupposition. The way it does, I suggest, is in how EXH maps to update instructions. The process of updating a context by computing $EXH_x(F(x)) = J$ can be broken up into two steps:

(37) a. $EXHx(F(x)) = J$

b. Find an x in the context. An exhaustive list of the values x can take is given by J.

c. $[\downarrow xFx] + [EXH(x) = J]$

This has the effect of breaking up the contribution of EXH into two parts. The first it the presuppositional contribution. The second is the asserted content, based on that presupposition.

The semantics of $\downarrow$ provides the presupposition, but we still need the semantics for EXH. To provide this, we need to be able to check what values a given variable takes in a world. For $i = \langle r, w\rangle$ and $i' = \langle r', w'\rangle$, say $i \sim i'$ iff $w = w'$. For $i \in s$:

(38) $i(EXH(x)) = \{d \mid \exists i' \in s(i \sim i' \wedge i'(x) = d)\}$

We can then compute $s[EXH(x) = J]$ in the usual way. Observe, the value of J can shift across worlds. If we have *It was the boys who broke the table*, we want a reading on which the boys involved differ in different possibilities.[25]

[24]For a recent discussion of some other options for the structure of clefts, see Hedberg (2000).

[25]There is also a rigid reading. This can be captured if we rigidify the description. But as is often the case with dynamic systems, rigidity must be implemented by brute force in the framework of Groenendijk *et al.* (1996). As they note, it is consistent that $\forall y \Diamond (x \neq y)$ in the framework we are using.

Let us check that we have the right truth conditions. Suppose $\downarrow xP(x)$ does not fail in s. Then we can compute:

(39)

$$\begin{aligned} s[\text{It was John who solved the problem}] &= s[\downarrow xP(x)][EXH(x) = \{j\}] \\ &= [\cup_d(s[x/d][P(x)])][EXH(x) = \{j\}] \\ &= \{i[x/d] \mid P(x) \wedge i[x/d](EXH(x)) = \{i[x/d](j)\}\} \end{aligned}$$

We start with a context in which in every world, someone solved the problem. We update to include only those worlds in which exhaustively John is the one who solved it, just as we expect.

The presupposition given by $\downarrow xF(x)$ turns out to be existential, not surprisingly, given it outputs $\exists xF(x)$ if the test succeeds. A context supports $\downarrow xF(x)$ just in case each world in it is one in which at least one individual is F.

As promised, this is a strong presupposition. It is for the same reasons we saw with the informal discussion of factives in Section (IV). The second instruction $[EXH(x) = \{j\}]$ is not defined if the presupposition fails.

V.4 *Too* and *Even*

We should now look at some weak presuppositions. A good starting place is *too*, which fits quite naturally into the framework I have developed.

Too is focus-sensitive. There are distinct presuppositions for:

(40) a. JOHN left too.

b. John LEFT too.

In the case of focus on *John*, we have:

(41) a. JOHN left too.

b. Presupposition: someone other than John left.

c. Asserted content: John left.

More generally, there must be a contextually salient element other than the value of the focused constituent which satisfies the predicate resulting from removing the focused constituent.

For (41), the update instructions are immediate:

(42) s[John left too] = $s[\downarrow x(x \neq j \wedge L(x))][L(j)]$

Here again, we have an existential presupposition. But this time, it is weak. In this case, the non-presuppositional update instruction—the one not in the scope of $\downarrow x$—has no occurrence of x. Hence, the fault-tolerant strategy I discussed above works, and we have only optional repair upon failure. If the context does not support the presupposition, the result of processing $s[\downarrow x(x \neq j \wedge L(x))] = \emptyset$. In that case, we can employ the strategy of reverting back to s and processing $s[L(j)]$. As we have no further occurrence of x in our update instructions, this update is well-defined. Hence, we do have a failure in the update algorithm, but a failure of a kind that a simple fault-tolerant procedure can overcome.

We get as similar result for *even*, though it will require some modification of the formal apparatus to provide it.

Let us look more closely at *even*, more or less following Rooth (1985). Examples like *Even JOHN solved the problem* appear to have the structure:

(43) [[$_{\text{NP}}$ Even JOHN] solved the problem]

However, it will simplify our presentation, and not seriously effect the content of what I shall propose, to treat it as a sentential operator:

(44) [[Even] [JOHN solved the problem]]

The interpretation of this is usually given in two parts as I mentioned at (5b):

(45) a. Asserted content: John solved the problem.

b. Presuppositions:

i. N solved the problem, for some contextually salient $N \neq$ John.

ii. For any contextually salient N such that N solved the problem, it is more likely/expected that N solved the problem than that John solved the problem.

The second presupposition, which builds in the contrast of *even*, is naturally put in terms of the alternative semantics for focus (in some appropriate form). For an embedded sentence S with a focused constituent, let $Alt(S)$ be the set of propositions that result from replacing the focused constituent of S with each of its contextually salient alternatives. The presuppositional clauses then tell us that $Alt(S)$ is nonempty, and for any $p \in Alt(S)$, p is more likely than $[\![S]\!]$.

The first presupposition is existential, but the second one is not. It appears to be propositional.[26] We already saw that factives require propositional presuppositions as well. More generally, it is well-know that presupposition is a cross-categorial phenomenon.

Putting aside the issue of how to define $\downarrow$ for propositional variables, the form of the update instructions for (44) is again not hard to see. Let us abbreviate the propositional presupposition of (44) as L. Then we want:

(46) a. s[Even John solved the problem]

b. $s[\downarrow x(x \neq j)][\downarrow q(q = L)][S(j, p)]$

Again, we see a weak presupposition, as the discourse tests of Section (II) predict.

There are a number of options for how propositional variables might be added to the formal apparatus to allow for this sort of update instruction. I am not entirely sure what the best route is, but about the most minimal modification that suffices is to add proposition variables ranging over sets of worlds (rather than, say, possibilities in some information state s). Propositions are then 'purely' propositions. We can add proposition constants as well, but I shall not bother with proposition predicates. Update of an information state by a proposition is essentially by intersection, though we have to be careful about reference systems. $s[a] = \{i = \langle r, w\rangle \in s \mid w \in a\}$. The natural definition of $\downarrow p(p = a)$ is then:

(47)

$$s[\downarrow p(p = a)] = \begin{cases} s[p/a][a] & \text{if } s \text{ subsists in } s[p/a][a] \\ \varnothing & \text{otherwise} \end{cases}$$

As there is only one element of $\wp(W)$ identical to a, we can re-write this to make it more in parallel with (25):

(48)

$$s[\downarrow p(p = a)] = \begin{cases} \cup_{q \in \wp(W)}(s[p/q][q = a][q]) & \text{if } s \text{ subsists in } \cup_{q \in \wp(W)}(s[p/q][q = a][q]) \\ \varnothing & \text{otherwise} \end{cases}$$

This allows us to use the update instructions given in (46) and (23).[27]

[26] Though I am sticking with the standard account of the presuppositions of *even*, my own judgments on the existential presupposition are quite weak. It would be fine with me if we left it out.

[27] The current definition of $\downarrow p$ still seems to me to be a bit forced (a kludge, as the programmers say). I would still like to

V.5 Anaphoric Potential

My analysis of presupposition triggers via ↓ makes them set up discourse referents. It should be noted that the anaphoric potentials of these referents differ. In some cases, we can use definite descriptions anaphoric on the value of ↓, but not pronouns:

(49) a. John solved the problem too. The other person did it first

b. # John solved the problem too. She did it first.

c. Even John solved the problem. The unlikeliness is astounding.

d. # Even John solved the problem. It not such a high degree of unlikeliness.

In some cases, we can use a pronoun:

(50) a. It was John who solved the problem. He did it easily.

b. John regrets voting for Bush. It has been bothering him for days.

My suspicion is that the availability of pronominal anaphora has something to do with the overt presence of the material referring to the discourse referent (even if it winds up being processed as within the scope of ↓). This is not the only place where some distinction like this seems to matter. It has been observed, for instance, that implicit and explicit arguments differ in anaphoric potential in much the same way:[28]

(51) a. I had Bill wash my car. He did a good job.

b. # I had my car washed. He did a good job.

c. I had my car washed. The guy did a good job.

V.6 Complex Demonstratives

We now have a fairly general account of the strong/weak distinction. It corresponds to whether we find anaphora on a ↓ in the semantics of a presupposition trigger. The informal sketch I gave in Section (IV) casts ↓ as an entirely general operator for introducing presuppositions. The formal work of this section certainly bears out that it can be applied quite widely, though more needs to be done to be sure it really can capture any presupposition we might want.

find a way to implement ↓ that captures its cross-categorial nature better. As I mentioned above, it might well be that one of the formalisms which more smoothly integrate λ-abstraction and dynamic binding, such as Kohlhase *et al.* (MS), might allow for a better definition. Of course, the cost would be more complicated machinery at the beginning.

[28] References—ter Meulen, Partee?

I shall close by pointing out that ↓ is flexible enough to capture some differences between presuppositions. Particularly, it can also capture the sorts of presuppositions which go with complex demonstratives. As I mentioned back at 5a, we have:

(52) a. That palm tree is about to fall.

b. Presupposition: contextually salient palm tree.

But there is more to say about this trigger. Unlike the cleft case, for instance, this does not seem to be merely an existential presupposition. Being a contextually salient palm tree for a demonstrative seems to require that there be a particular individual which is known in the context to be a palm tree. We can capture this sort of stronger presupposition using ↓ as well:

(53) a. s[that palm tree is about to fall]

b. $s[\downarrow x \Box P(x)][F(x)]$

s supports this presupposition if there is at least one object in D—one contextually salient object—which satisfies P in every world in s. Hence, it is already common ground that the object is a palm tree. This is stronger than the sort of existential presupposition we saw with the cleft. (Actually, the presupposition of the demonstrative is probably still stronger than this. We should require not just contextual salience, but the kind of super-salience that results from being demonstrated, or the object of a demonstrative intention.[29])

VI Conclusion

I began by noting that an account of presupposition along pragmatic lines leaves open how to understand the relevant notion of felicity, and how presuppositions are triggered. I then pointed out that presupposition failures fall into two categories, corresponding to the discourse statuses of repair obligatory and repair optional. Most of my effort was then directed at an analysis of some presupposition triggers which explains this difference. According to the ↓-based analyses, the difference stems from the different ways update instructions can fail, and whether a simple Gricean fault-tolerant strategy can bypass the failure.

This gives us the beginning of an idea of what the notion of felicity that goes with presupposition is like. In the sort of dynamic framework I have used here, felicity is analyzed as non-failing processing of

[29] No doubt the presuppositions of complex demonstratives is still a controversial subject. I have pursued it at more length in unpublished work with Susanna Siegel. All I want to do here is note that if we need these stronger sorts of presuppositions, ↓ can provide them.

update instructions. I am inclined to see this as more than an artifact of the theory. Computing update instructions offers a model of how speakers might come to convey information by use of language. Non-failing computation is the standard of successful communication. We have seen that indeed information can be conveyed even in the face of failure, by invoking some fault-tolerant strategy or another. But the judgments that go with weak presupposition support the idea that this still counts as infelicity. Felicity is simply the situation of successful processing of update instruction. It is successful communication directly by way of the semantics of the sentence used.

References

Austin, J. L., 1975. *How to Do Things with Words*. Cambridge: Harvard University Press, 2nd edn. Edited by J. O. Urmson and M. Sbisà.

Beaver, D. I., 1997. Presupposition. In J. van Benthem and A. ter Meulen (eds.), *Handbook of Logic and Language*, pp. 939–1008. Amsterdam: Elsevier.

———, 2001. *Presupposition and Assertion in Dynamic Semantics*. Stanford: CSLI Publications.

Chierchia, G. and S. McConnell-Ginet, 1990. *Meaning and Grammar: An Introduction to Semantics*. Cambridge: MIT Press.

Davis, S., 1991. *Pragmatics*. Oxford: Oxford University Press.

Delin, J., 1992. Properties of it-cleft presupposition. *Journal of Semantics* **9**:289–306.

———, 1995. Pesupposition and shared knowledge in *It*-clefts. *Language and Cognitive Processes* **10**:97–120.

É. Kiss, K., 1998. Identificational focus versus information focus. *Language* **74**:245–273.

Frege, G., 1892. Über Sinn und Bedeutung. *Zeitschrift für Philosophie und philosophische Kritik* **100**:25–50. References are to the translation as "On Sense and Meaning" by M. Black, reprinted in Frege (1984).

———, 1984. *Collected Papers on Mathematics, Logic, and Philosophy*. Oxford: Basil Blackwell. Edited by B. McGuinness.

Glanzberg, M., 2002. Context and discourse. *Mind and Language* **17**:333–375.

Groenendijk, J., M. Stokhof and F. Veltman, 1996. Coreference and modality. In S. Lappin (ed.), *Handbook of Contemporary Semantic Theory*, pp. 179–213. Oxford: Blackwell.

Hedberg, N., 2000. The referential status of clefts. *Language* **76**:891–920.

Heim, I., 1983. On the projection problem for presupposition. *Proceedings of the West Coast Conference on Formal Linguistics* **2**(114–125). Reprinted in Davis (1991).

Horn, L. R., 1989. *A Natural History of Negation*. Chicago: University of Chicago Press.

Kadmon, N., 2001. *Formal Pragmatics*. Oxford: Blackwell.

Karttunen, L., 1974. Presupposition and linguistic context. *Theoretical Linguistics* **1**:181–194. Reprinted in Davis (1991).

Kohlhase, M., S. Kuschert and M. Pinkal, MS. A type-theoretic semantics for λ-DRT. Unpublished manuscript.

Krahmer, E., 1008. *Presupposition and Anaphora*. Stanford: CSLI Publications.

Lasersohn, P., 1993. Existence presuppositions and background knowledge. *Journal of Semantics* **10**:113–122.

Levinson, S. C., 1983. *Pragmatics*. Cambridge: Cambridge University Press.

Prince, E. F., 1978. A comparison of wh-clefts and it-clefts in discourse. *Language* **54**:883–906.

Reinhart, T., 1981. Pragmatics and linguistics: An analysis of sentence topics. *Philosophica* **27**(1):53–94.

Rooth, M., 1985. *Association with Focus*. Ph.D. dissertation, University of Massachusetts at Amherst.

———, 1999. Association with focus or association with presupposition? In P. Bosch and R. A. van der Sandt (eds.), *Focus*, pp. 232–244. Cambridge: Cambridge University Press.

Schegloff, E. A., G. Jefferson and H. Sackes, 1977. The preference for self-correction in the oranization of repair in conversation. *Language* **53**:361–382.

Simons, M., 2001. On the conversational basis of some presuppositions. *Semantics and Linguistic Theory* **11**.

———, ms. Why some presuppositions are conversational implicatures. Manuscript.

Soames, S., 1989. Presupposition. In D. Gabbay and F. Guenthner (eds.), *Handbook of Philosophical Logic*, vol. IV, pp. 553–616. Dordrecht: Kluwer.

Stalnaker, R. C., 1974. Pragmatic presuppositions. In M. K. Munitz and P. K. Unger (eds.), *Semantics and Philosophy*, pp. 197–213. New York: New York University Press. Reprinted in Stalnaker (1999).

———, 1978. Assertion. In P. Cole (ed.), *Pragmatics*, vol. 9 of *Syntax and Semantics*, pp. 315–322. New York: Academic Press. Reprinted in Stalnaker (1999).

———, 1999. *Context and Content*. Oxford: Oxford University Press.

Strawson, P. F., 1950. On referring. *Mind* **59**. Reprinted in Strawson (1971).

———, 1952. *Introduction to Logical Theory*. London: Methuen.

———, 1964. Identifying reference and truth-values. *Theoria* **30**. Reprinted in Strawson (1971).

———, 1971. *Logico-Linguistic Papers*. London: Methuen.

van der Sandt, R. A., 1992. Presupposition projection as anaphora resolution. *Journal of Semantics* **9**:333–377.

van Eijck, J., 1993. The dynamics of descriptions. *Journal of Semantics* **10**:239–267.

———, 1994. Presupposition failure—A comedy of errors. *Formal Aspects of Computing* **6A**:766–787.

von Fintel, K., forthcoming. Would you believe it? The King of France is back! In A. Bezuidenhout and M. Reimer (eds.), *Descriptions and Beyond*. Oxford: Oxford University Press.

Zeevat, H., 1992. Presupposition and accomodation in update semantics. *Journal of Semantics* **9**:379–412.