Modals and a Compositional Account of Counterfactuals

Nicholas Asher
Department of Philosophy
The University of Texas@Austin
and
Eric McCready
Department of Linguistics
The University of Texas@Austin

January 17, 2005

1 Introduction

Over the past twenty years, there have been numerous discoveries and theoretical proposals for modals in Romance and Germanic languages, as well as for conditionals and counterfactuals. But there has been little attempt to produce a compositional account of counterfactuals from an account of the meanings for the conditional and the modals.[1] Here we present a detailed, dynamic semantics for modals that leads to a new semantics for counterfactuals.

For the modals *might* and *would*, there are two important sets of observations that pull any putative account of these constructions in different directions. Veltman's (1996) seminal paper on epistemic modals introduces the idea that epistemic possibilities conveyed with a modal like *might* interact dynamically with factual information introduced in a discourse. Factual information introduced into a discourse may rule out certain epistemic possibilities that might otherwise be present. Consider the following minimal pair.

(1)   a.   It might be sunny. It's not sunny.

     b.   #It is not sunny. It might be sunny.

Veltman's examples show that epistemic possibilities are subject to what information has already been introduced in the discourse; once it has been established in the discourse that it is not sunny, it is no longer permissible to introduce the epistemic possibility that it might be sunny, as in (1b). If you think of *mightp* as ambiguous between the epistemic reading that says "for all that is known given the discourse so far, it's

---

[1] A prominent exception to this generalization is Gillies (2003) who makes a good start on such a project, but there remains a lot of work to be done.

possible that p" and the metaphysical reading, (1b) only has a marginal metaphysical reading, whereas (1a) has both readings felicitously.[2] Further, while some have claimed that (1b) has a reading according to which the speaker revises his contribution, we claim that revisions need much more linguistic marking than is possible in (1 b). Like Veltman, we believe that (1b) remains very marginal even when one has the revision scenario in mind. Introducing as an epistemic possibility something that has already been established in the discourse as in (1c) is also infelicitous, though less so than (1b).

(1c)   ? It is sunny. It might be sunny.

This observation has largely been accounted for on pragmatic grounds: since p entails *mightp*, the assertion of *mightp* after an assertion of p conveys no new information and so violates the Gricean maxim that a speaker's contributions should be at least minimally informative.

Another observation that has motivated a considerable amount of work in dynamic semantics—e.g. the accounts of Roberts (1987) and Frank (1997) in Discourse Representation Theory—is that uses of modals may pick up and refine epistemic possibilities introduced by the use of modals in previous utterances. This phenomenon has come to be known as *modal subordination*. Consider for example the discourses in (2).

(2)   a.   A wolf might walk in. It would eat you first.

      b.   A wolf might walk in. # It will eat you first.

      c.   A wolf might walk in. It might eat you first. But then it might not.

There is a striking difference between (2a) and (2b). In (2a) the use of the epistemic modal *would* enables the pronoun *it* to find its intended antecedent, the wolf introduced under the scope of the modal in the first sentence, while futurate *will* in (2b) does not allow this possibility. (2c) shows that the modal *might* has the same effect as *would* in enabling the accessibility of the intended antecedentl Nevertheless, since *a wolf* occurs under the scope of the modal operator in that sentence, it is unavailable as an antecedent for the pronoun in nonmodal contexts, which is what standard dynamic semantics predicts. The accessibility of the antecedent under the scope of a modal to a pronoun also under the scope of a modal, however, was something that standard dynamic semantic accounts of anaphora as well as more traditional accounts could not predict; and the accounts of Roberts and Frank provided significant insights into the semantics of anaphoric expressions.

The natural way to account for Veltman's observations and the natural way to account for the data on modal subordination don't easily combine. They involve different techniques for evaluating formulas with respect to a set of indices or points of evaluation: one approach is to evaluate a formula relative to properties of sets of such evaluation points—a sort of "collective" approach (to use some terminology from another area of semantics, the interpretation of plurals); the other evaluates a formula relative to individual evaluation points—a "distributive" approach. The history of the efforts in

---

[2]We find that the metaphysical reading is much better for the following variant of (1b): *It is not sunny. But it might have been sunny.*

this area suggests that a good semantics for the modals has to be both collective and distributive with respect to points of evaluation.

But before we get to theory, let's look at a bit more data about modals. An interesting use of the epistemic modal *would* concerns its effects as an agreement marker in discourse. In (3), B's utterance conveys an agreement with A's assertion by conveying that in all or most of B's epistemic possibilities, Kim's teasing Pat was something to be expected.

(3) A:. Kim teased Pat.

    B:. Kim would do that.

This leads one to wonder about the relationship between *might* and *would*. Their semantics must be strongly interconnected, and they appear to be weak duals, in the sense that one cannot have *mightp* and *would¬p*. Similarly, most speakers baulk at *wouldp* but *might notp*. We will assume something slightly stronger–namely that it's the case that *wouldp* implies that *mightp*, which intuitions also appear support (witness the implication in (4c).

(4) a. #John might come to the party but John would not come to the party.

    b. #John would not come to the party but John might come to the party.

    c. John would come to the party $\longrightarrow$ John might come to the party.

The semantics of the modals is rich and intriguing. But especially intriguing is how so many modals combine productively together with a conditional whose antecedent is adjusted for the appropriate mood to create counterfactuals of various kinds.

(5) a. If I were not to sleep tonight, I would topple over tomorrow (I might topple over tomorrow).

    b. What should be your mascot if you were a school? (Google)

    c. If we were to get more serious, should I tell him my age? (Google)

    d. If you were a school, your mascot should be a fierce animal.

    e. If it were easy, anyone could do it (Google, *Converge Magazine*).

The ease with which modals and conditionals combine strongly suggests that we should try to build a compositional account of the semantics for counterfactuals. The slight differences in meaning between the different modals create different sorts of counterfactuals; for instance *should* and *would* don't have quite the same range of meanings; *should*, and nor do *could* and *might*. And not unexpectedly, there is a definite change in meaning between (5e) and the variant *if it were easy, anyone might do it*.

2   Previous Accounts

In order to set the stage for our analysis, we need first to sketch out the backdrop for most, if not all of the recent analyses of modals. That backdrop is that of dynamic semantics. Dynamic semantics was designed to account for how the meaning of an utterance in a particular discourse context might change or contribute to that very discourse

context. Accordingly, all dynamic semantic theories analyze the meaning of a sentence as a relation between contexts, an input context and an output context. Broadly speaking, contexts in dynamic semantics are sets of n-tuples consisting at least of a world, an assignment function, and often other things as well.[3] There are two broad parameters in the way dynamic semantics may describe the effects of an utterance on a context: first it may be *distrtibutive* and define the effects of a sentence on each element of the context or not (this is the collective view); secondly these effects may affect those elements of the context or not. If not the semantics is called *eliminative*. DRT of Kamp and Reyle (1993) as well as DPL (Groenendijk and Stokhof 1990) are examples of distributive semantics that are not eliminative, whereas Veltman's (1985, 1996) update semantics which he uses for his account of the modals, is eliminative but not distributive.

Let's take a look at Veltman's update semantics first. Contexts are understood as sets of worlds and the epistemic sense of *might* that is to be captured is one that surveys the possibilities common to the discourse participants who have accepted the information in the discourse so far. Thus, the picture is that as discourse proceeds the set of epistemic possibiilities left open to the participants gradually narrows as they build up a common ground of information between them. When updating a context with a formula, we either eliminate worlds from that context that do not satisfy it or the formula operates as a "test" on the context as a whole. Formulas of the form *mightϕ*, *wouldϕ* and conditionals function as tests on the context. Here is the semantic clause for formulas of the form *mightϕ*:

- Veltman's account of *might*:

- Let σ be a set of possible worlds and ϕ a formula of a propositional language with the usual interpretation at each world. Then:
  $\sigma + might\phi = \sigma$ if $\sigma \cap \phi \neq 0$
  $\qquad\qquad = 0$ otherwise.

The idea is that an information state will pass the test of *mightϕ* whenever ϕ is true in at least one world in the context. But updating contexts with factual information such as in (1c) will eliminate all those worlds that support the epistemic possibility that it is sunny, and so then attempting to update the information state with *it might be sunny* will yield an empty information state, a sign that something has gone wrong.

Veltman's theory is a theory that does not appeal to the speaker's private beliefs or intentions. An alternative is to adopt a much more plainly pragmatic story in line with the pragmatic story told for (1c): we could easily adapt Grice's maxim of quality to the effect that if a speaker utters (1b) he has violated the rule that you don't say something you don't know to be true. But even if the speaker is lying and so is already violating the maxim of quality, (1b) remains incoherent on an epistemic reading. That is the intuition that Veltman, we think, correctly captures.

Veltman's theory captures the data about the interactions between *might* sentences and non modal sentences effectively. But it suffers defects when we turn to other data

---

[3]In this paper we will use the notion of possible world in its ordinary sense as an evaluation point that gives a truth value to all well formed sentences of the formal language in a classical fashion. We do not investigate, here, a notion of epistemic possibility which may not be logically closed.

about the modals, in particular the data on modal subordination or how modal statements interact dynamically. *might* is just a test on the whole information state according to Veltman's semantics; so if the update with *might*ϕ is successful, then we get back the same context that we started with. The facts about modal subordination, however, require us to isolate the set of epistemic possibilities in which ϕ in order to further modify them; e.g., we must isolate those worlds in the context where the wolf walks in as in (2a), in order to further refine those possibilities as those where the hearer gets eaten first. This is simply not possible on the Veltman semantics: we can't isolate the relevant set of epistemic possibilities and we can't modify that set, unless we update with a nonmodal assertion. But the facts about modal subordination show that modification is possible with modals and that this modification is not equivalent to updating with a nonmodal assertion. Veltman's test semantics for *might* alone can't furnish us an appropriate context in which to evaluate subsequent modal claims.

There are other difficulties with update semantics as well that have to do with the interpretation of quantification and anaphoric links across sentences. Most research in dynamic semantics has adopted the distributive semantics for quantification (e.g., Groenendijk and Stokhof 1991, Kamp and Reyle 1993), and it is within this framework that the first accounts of the facts of modal subordination were developed. Roberts' (1987) account of modals, for instance, considers them to be two place operators, one argument of which is given by what is in the scope of the modal operator, the other given by some contextually available proposition. Thus, for (2a), the proposition under the scope of the *might* operator in the first sentence provides the first argument to the operator *would* in the second sentence, while the second argument is given by the material under the scope of the *would* in the second sentence. This furnishes the intuitively correct reading of second sentence with the following semantics of *would*(ϕ, ψ): every element of the context that satisfies ϕ also satisfies ψ; in other words, this semantics says that the formula *it would eat you first* is satisfied in a context just in case every element of the context in which a wolf walks in is also an element in which the wolf eats the hearer first. This is a distributive semantics for the modal operators.

Two place operator accounts like Roberts's (but see also Frank 1997) don't specify the first argument of a modal operator; that's left to pragmatics. So by itself these accounts don't get us the right interpretation of (2a); e.g., another possible interpretation of (2a) is that the first argument of *would* is not just the proposition under the scope of *might*, but any other contextually salient proposition, including propositions like the one expressed by *there's a wolf*. This of course is not an intended reading of the example. Further, the distributive semantics doesn't predict the data in (1ab); any number of contextually available antecedents might make (1b) acceptable. The two place operator accounts don't capture the interactions between modal and nonmodal information. We want to give a semantics for *might* that captures both the data about modal subordination and Veltman's observations.

3 Semantics

3.1 Combining Distributive and Non-Distributive Intuitions

Our goal, then, is to account both for the motivating data of the two place operator approach and of the non-distributive semantics for the modals. We begin with a standard dynamic semantic formalism like that of DPL, which has a very good account of quantification and intersentential anaphora. The elements of the context are world assignment function pairs. Let $\models$ be a standard Tarskian notion of static satisfaction and M a standard first order model.

- $(w, f)\|Pt_1, \ldots, t_n\|^{\mathcal{A}}(w', g)$ iff $g = f \wedge w = w' \wedge \mathcal{A}, w, f \models Pt_1, \ldots t_n$ iff $\langle \|t_1\|^{\mathcal{A}}_{(w,f)} \ldots \|t_n\|^{\mathcal{A}}_{(w,f)} \rangle \in R^{\mathcal{A}}_w$

- $(w, f)\|t_1 = t_2\|^{M}(w', g)$ iff $g = f \wedge w = w' \wedge \mathcal{A}, f \models t_1 = t_2$

- $(w, f)\|\phi \wedge \psi\|^{\mathcal{A}}(w', g)$ iff $\exists w'' h(w, f)\|\phi\|^{\mathcal{A}}(w'', h) \wedge (w'', h)\|\psi\|^{\mathcal{A}}(w', g)$
  in other words: $(w, f)\|\phi\|^{M} \circ \|\psi\|^{\mathcal{A}}(w', g)$

- $(w, f)\|\neg\phi\|^{\mathcal{A}}(w', g)$ iff $f = g \wedge w = w' \wedge \neg\exists w'', h(w, f)\|\phi\|^{\mathcal{A}}(w'', h)$

- $(w, f)\|\exists x\phi\|^{\mathcal{A}}(w', g)$ iff $w = w' \wedge \exists a \in A(w, f^{\frac{a}{x}})\|\phi\|^{\mathcal{A}}(w', g)$

In order to capture Veltman's examples in (1ab), we need to test the current information state when evaluating sentences of the form *mightϕ*. So we need to combine a DPL style semantics for the quantifiers with an update style semantics for modals. But we cannot simply use the DPL clauses for quantification with the operations for truth functional and modal operators defined by simple operations on sets. Because of the way existential quantifiers operate by resetting values of assignment functions, a standard approach to negation for example at the set level gets things wrong. Here's the standard and common sense clause for negation when trying to define operations on sets of world assignment pairs:

- $\sigma + \phi = \sigma - (\sigma + \phi)$.

However, now consider $\sigma + \neg\exists x\phi$. If the existential is to reset values of assignment functions as we expect in dynamic semantics then $\sigma + \exists x\phi$ may very well be disjoint from σ depending on the choice of resetting values. So $\sigma + \neg\exists x\phi = \sigma$ and it could very well be that we can then have a non-empty update result by updating with $\sigma + \exists x\phi$. We can solve this by using the standard DPL clauses and lifting the results to an information state level keeping everything distributive. As Fernando (1993) noted, we can make the non functional DPL semantics functional at the level of sets of world assignment pairs, or we could define a relational interpretation at the level of sets, mimicking DPL's nondeterministic behavior. But this alone won't give us the right semantics for the modals because it's still completely distributive and not collective..

The terminology we've used to talk about ways of interpreting modals suggests an analogy with the semantics of plurals. The literature on the semantics of plurals has demonstrated that not only are there sentences with collective interpretations and sentences with distributive interpretations but there are discourses combining both distributive and collective interpretations (the example is from Asher and Wang 2003):

(6)   Three boys worked tirelessly and mowed the whole meadow.

In this example we understand each boy working tirelessly individually and that collectively they act together so as to mow the whole meadow. Collective and distributive interpretations of plurals are analogous to non-distributive and distributive interpretations of modals: collective interpretations test sets of assignments within a context, while distributive interpretations are determined relative to individual assignments that comprise the individual world assignment pairs in the context. Similarly, we need to exploit sets of epistemic possibilities or sets of world assignment pairs to interpret some modals, while integrating this within a generally distributive semantics.

For a variety of reasons that we won't go into here (see Asher and Wang 2003 for details though), the only way we know to integrate the collective and distributive approach to assignments is to complicate the notion of a dynamic context by making each element in the context contain information about the assignments besides its own that form a natural group in the context. Here we're going to do the same for the evaluation points of basic dynamic states. Further, we'll need a non-deterministic relation between elements of evaluation; this non determinism is necessary at the level of assignments for plurals since, to put it somewhat roughly, plural indefinites may pick out many distinct candidate sets of individuals as values for the plural variable.

Accordingly, we extend our notion of information state and our structures $\mathcal{A}$. $\mathcal{A}$ will include a set of epistemic possibilities $\mathcal{E}$ each element of which is set of triples consisting of a world, an assignment and a set of epistemic possibilities. To define each set of epistemic possibilities properly, we proceed inductively. We begin with some choice α of some set of world assignment pairs and use that choice to inductively build up more complicated sets of epistemic possibilities.

- Set $E_{\alpha,0} \subseteq \mathcal{P}(W \times \$)$, where $\$$ is the set of all assignment functions.

- $E_{\alpha,n+1} \subseteq \mathcal{P}(W \times \$ \times E_{\alpha,n})$

- $\mathcal{E}_\alpha \subseteq \mathcal{P}(\bigcup_{n \in \omega} E_{\alpha,n})$

Every set of epistemic possibilities $\mathcal{E}_\alpha$ is thus well-founded. So elements σ of our dynamic contexts are triples of $\langle w, f, \mathcal{E}_\alpha \rangle$ for some α. We'll refer to the third element of an information state σ as $3(\sigma)$, where 3 is a projection function from σ to its third element . We will make use of the projection functions $1, 2, 3$ to pick out the world, assignment function or set of epistemic possibilities of a context element respectively. We adopt the constraint that epistemic possibilities at the outset include the actual world; namely, we stipulate that for the intial state σ_0, $\exists \sigma' \in \bigcup(3(\sigma_0))$ $(1(\sigma_0) = 1(\sigma') \wedge 2(\sigma_0) = 2(\sigma'))$.

We now restate our distributive dynamic DPL style semantics in terms of our new context elements.

- $\sigma \| R t_1, \ldots, t_n \|^{\mathcal{A}} \sigma'$ iff $\sigma = \sigma' \wedge \langle \|t_1\|^{\mathcal{A}}_{(1(\sigma),2(\sigma))} \cdots \|t_n\|^{\mathcal{A}}_{(1(\sigma),2(\sigma))} \rangle \in R^{\mathcal{A}}_{1(\sigma)}$

- $\sigma \| t_1 = t_2 \|^{\mathcal{A}} \sigma'$ iff $\sigma = \sigma' \wedge \|t_1\|^{\mathcal{A}}_{(1(\sigma),2(\sigma))} = \|t_2\|^{\mathcal{A}}_{(1(\sigma),2(\sigma))}$

- $\sigma \| \phi \wedge \psi \|^{\mathcal{A}} \sigma'$ iff $\sigma \| \phi \|^{\mathcal{A}} \circ \| \psi \|^{\mathcal{A}} \sigma'$

7

- $\sigma\|\neg\phi\|^{\mathcal{A}}\sigma')$ iff $\sigma = \sigma' \wedge \neg\exists w'', h\ \ \sigma\|\phi\|^{\mathcal{A}}\sigma^{\frac{1(\sigma)}{w''},\frac{2(\sigma}{h}}$

- $\sigma\|\exists x\phi\|^{\mathcal{A}}\sigma'$ iff $\exists a \in A\ \ \sigma^{\frac{a}{x}}\|\phi\|^{\mathcal{A}}\sigma'$, where $\sigma^{\frac{a}{x}}$ is the result of replacing $2(\sigma)$ with $2(\sigma)^{\frac{a}{x}}$

So far nothing here is out of the ordinary. But we're missing something. As discourse proceeds we learn things and so refine and indeed revise our epistemic possibilities in light of what has been learned. Let us call the *discourse context* that set of triples that are the result our output of the evaluation of successive sentence-tokens in a discourse. A discourse context is the same sort of animal as an epistemic possibility—- a set of world assignment, epistemic possibility triples; and it contains the information of what has been said up to this point. Simplifying matters considerably, we will take what has been said in discourse has having been established and accepted as part of the common ground (thus passing over all the problems of correction, denial and disagreement— but see Asher and Lascarides 2003 or Asher and Gillies 2004 for discussions of these phenomena—in any case, we think that revision involves necessarily some sort of discourse indications that this is taking place, indications that are not present in these simple examples we consider here). Thus, whatever is true or supported in such a discourse context should be reflected in the set of epistemic possibilities of those triples σ that are part of the discourse context.

To define this constraint, we'll introduce a new notion of a *discourse update*, and auxilliary notions of *descendant satisfaction*, written $\models_d$ and revision. The notion of descendant satisfaction lifts our distributive semantics over context elements to sets of such elements (the term is due to Groenendijk, Stokhof and Veltman 1994, though we develop the notion somewhat differently). Thanks to the work of Lewis, Spohn and others, it is straightforward to define a revision function $\star$ on epistemic possibilities if we assume a partial ordering on the elements of epistemic possibilities. This partial ordering forms a system of spheres centered around each element σ. A set of such elements can also have a system of spheres $S(\epsilon) = \{\cup(S_n(\sigma)) : \sigma \in \epsilon\}$.

- Definition of Descendance and Satisfaction by epistemic possibilities

- σ has a ϕ descendant σ' iff $\sigma[\phi]\sigma'$

- $\langle\epsilon, \epsilon'\rangle \models_d \phi$ iff every $\sigma \in \epsilon$ has a ϕ descendant in ϵ'.

- $\|\phi\| = \{\langle\sigma, \sigma'\rangle : \sigma'$ is a ϕ descendant of $\sigma\}$

- Let $S_n(\epsilon)$ be the smallest sphere around ϵ such that elements in $S_n(\epsilon)$ have ϕ descendants. Then $\epsilon \star \|\phi\| = \{\sigma : \exists \sigma' \in S_n(\epsilon)\ \ \sigma$ is a ϕ descendant of $\sigma'\}$.
  $\mathcal{E} \star \|\phi\| = \{\epsilon \star \|\phi\| : \epsilon \in \mathcal{E}\}$.

With these notions we can now turn to the central notion of discourse update.

- Definition of Discourse Update:

- Let ϕ be a modal free formula. Then σ is a ϕ discourse update of σ' iff $\exists\sigma_1$ such that $(\sigma'[\phi]^{\mathcal{A}}\sigma_1 \wedge 1(\sigma) = 1(\sigma_1) \wedge 2(\sigma) = 2(\sigma_1)$ and
  $3(\sigma) = 3(\sigma') \star \|\phi\| \wedge \forall\epsilon \in 3(\sigma)\forall\sigma_2 \in \epsilon\ \ 2(\sigma_2) =_x 2(\sigma)$ for all x free in ϕ)

- Let ϕ be a formula of the form *mightϕ*, *wouldϕ* or $\phi \Rightarrow \psi$. Then σ is a ϕ discourse update of σ' iff $\sigma'[\phi]^{\mathcal{A}}\sigma$

The notion of discourse update contains the idea that nonmodal information introduced into the discourse must be reflected in the updated epistemic possibilities (hence the need for the revision operator $\star$). Further, the bindings of variables that occur free in a formula carry their already established values. Discourse update helps us to evaluate sequences of formulae that are translations of our examples. For instance in evaluating (1ab), we check whether the translations of those formulae give us a sequence of coherent discourse updates, where a coherent discourse update is one where for some input σ there is a non-empty output. Discourse update is also the notion that we need to define logical consequence.

A feature of discourse update that might strike one as odd is that modal formulas do not affect the epistemic possibilities in discourse update. But they don't need to, since they already do so in their basic semantics–that is, in how they affect dynamic transitions over σ. Here is the basic semantics for *might* and *would*. The semantics of the epistemic modals can either test or change the second element.

- $\sigma[\textit{might}\phi]_{\mathcal{A}}\langle 1(\sigma), 2(\sigma), \mathcal{E}'\rangle$, where $\mathcal{E}' = \{\epsilon' : \exists \epsilon \in 3(\sigma)\langle \epsilon, \epsilon'\rangle \models_d \phi\}$,
  if there is such an ϵ;
  $\sigma[\textit{might}\phi]_{\mathcal{A}}\emptyset$ otherwise.

- $\sigma[\textit{would}\phi]_{\mathcal{A}}\langle 1(\sigma), 2(\sigma), \{\epsilon' : \exists \epsilon \in 3(\sigma)\langle \epsilon, \epsilon'\rangle \models_d \phi\}\rangle$,
  if $\forall \epsilon \in 3(\sigma)\exists * \langle \epsilon, \epsilon*\rangle \models_d \phi$;
  $\sigma[\textit{would}\phi]_{\mathcal{A}}\emptyset$ otherwise.

might intuitively involves an existential quantification over epistemic possibilities. And like all existentials in dynamic semantics, it has a special status—that of resetting, in this case, epistemic possibilities. But this resetting is dependent on a test of the input; if the previous epistemic possibilities admit an update with the proposition under the scope of the *might*, then the resetting proceeds—if not, the update fails in the sense of producing no descendants for the input. This semantics incorporates the test idea of Veltman's semantics but it is not itself a simple test; it allows information under the scope of the *might* operator to transform the epistemic possibilities in the input. It also differs from "accommodation" views of *might* like von Fintel's (2002), according to which *might* always enlarges the epistemic possibilities under consideration. On this view it rather refines certain epistemic possibilities that must be already in place.

Our epistemic semantics for the simple modals is now complete. Unlike Veltman or Groenendijk, Stokhof and Veltman (1994), we separate out a set of epistemic possibilities associated with each world assignment pair, on which epistemic formulas operate. So our notions of logical consequence and validity can remain those familiar from DPL with one important change: we replace the basic notion of a dynamic transition with our notion of discourse update. This is needed to ensure that new factual information affects the epistemic possibilities in the relevant way.

- Logical Consequence: Let Γ be a sequence of formulae. Then $\Gamma \models \phi$ iff for all $\mathcal{L}$ models $\mathcal{A}$ for all information states σ, σ' such that σ' is a Γ discourse update of σ there is a σ'' such that σ'' is a ϕ discourse update of σ''

- Validity: $\models \phi$ iff $0 \models \phi$

We can verify some logical consequences of our definitions.

Fact 1 *Facts about* would *and* might*:*

- would$\phi \not\models \phi$

- $\phi \models$ mightϕ

- might$(\phi \wedge \psi) \models$ might$(\phi) \wedge$ might(ψ).

- would$\phi \models$ mightϕ

- $\phi \models$ wouldϕ

This semantics for *would* is sensitive to the fact that *would* ϕ is not equivalent to ϕ. This is a virtue. The fact that the discourse context supports *wouldϕ* and hence that all of the epistemic possibilities of any element of the context support ϕ does not guarantee that the discourse context supports ϕ, which is what our notion of entailment requires. Ther epistemic possibilities might not cover all the *logical* or *semantic* possibilities compatible with the information presented in the discourse or otherwise available from the context. So in sum: *wouldϕ* $\not\rightarrow$ ϕ. Our semantics, like Veltman's original semantics, verifies $\phi \models$ *mightϕ* by the definition of *might* and the definition of discourse update. On the other hand, *mightϕ* $\not\models$ ϕ. Resetting the epistemic possibilities to reflect ϕ doesn't necessarily affect the actual world of evaluation; although the actual world and assignment of evaluation are elements of the epistemic possibilities associated with them in the empty information state. Updating with modal formulas may make our epistemic possibility set go counterfactual.

These clauses also verify the desired entailment from *wouldϕ* to *mightϕ*. But we also have: $\phi \models$ *wouldϕ*. This might seem to be a potential problem. But clearly it's pretty terrible to have ϕ and $\neg$*wouldϕ*:

(7)   a.   ??John is at the party, but he wouldn't be at the party.

       b.   ?John is at the party; so he would be at the party.

(7b) may be pragmatically difficult, but semantically OK. (7a) is nonsense unless we understand some sort of suppressed antecedent of a counterfactual as occurring there, for which there is little if any evidence.

Further, this semantics predicts that because the epistemic possibilities must always verify what has already been established in the discourse, Veltman's examples (1ab) immediately fall out as predicted. *might* ϕ resets the epistemic possibilities of an element of the discourse context σ to those where ϕ holds as long as ϕ was an epistemic possibility in σ. So updating an empty information state with *it might be sunny* simply resets the possibilities and we can then update with the factual information that it's not sunny which will revise the epistemic possibilites to reflect the fact that we have now learned that it's sunny. However, updating first with *it's not sunny* makes a subsequent update with *it might be sunny* fail, because the input information state does not contain it's being sunny as an epistemic possibility. Further, our semantics predicts that examples like (4ab) should not yield any coherent output for any given input state.

Finally, this semantics makes some intuitive predictions about the cases of stand alone *would*. Consider again (3). When attached to *A*'s assertion it's natural to understand *B* as saying that in all of his epistemic possibilities, what *A* says turns out to be true. It marks a form of agreement, which is intuitively what is going on in (3b). The anaphoric account would make *B*'s assertion with stand alone *would* some sort of logical truth on the standard semantics for *would* and so should be ruled out on pragmatic grounds of informativeness.

3.1.1 A Problem?

A problematic inference in our semantics as its stands is the inference from *might φ* to *would φ*. Without further modification, this inference is valid, and it also looks valid (depending on how you construe validity in the anaphoric framework) in the analyses of Frank and Roberts. Now some speakers have the feeling that the assertion of *would φ* after an assertion of *might φ* (by the same speaker) is redundant. And such redundancy is a mark of entailment. But still we have qualms about the validity of this inference. For one thing, if you hear a sequence like *might φ would φ*, it sounds like some sort of a Correction. It's clearly not a valid inference. What stops it from being a valid inference, we hypothesize, is an evidential presupposition of *would φ*, which we'll write as $\partial(would\phi)$. There's a lot of evidence that modals in many languages have evidential presuppositions, something that we investigate in detail in Asher and McCready (2004b). Such presuppositions help us explain our understanding of a *might φ would φ* sequence as a Correction, since $\partial(would\phi)$ conflicts with an implicature of *mightφ*, which is that a stronger modality with a stronger evidential presupposition *doesn't* hold. Asher and McCready (2004b) also provide evidence for an important observation about *how* evidential presuppositions are to be accommodated. It follows the basic rule that "just saying it doesn't make it so." We accommodate an evidential presupposition *before* any modal updates. This is because the evidential requirement for *would* is required to be *actual* not just based on some epistemic possibility. Now if we attend to the inference from *might φ* to *would φ* we see immediately that in certain structures and at certain elements ϵ, $\mathcal{E}(\epsilon)$ as well as ϵ itself will not support $\partial(would\phi)$ and so the attempted accomodation will not yield a non-empty discourse update. But if the presupposition is not bound or accommodated, then the inference cannot go through. Thus, the evidential presupposition of *would* plays an important role in distinguishing *would* from *might*.

3.2 Modal Subordination and Discourse Structure

One of our main goals was to account for the basic facts of modal subordination. The classic (2a) works as expected. The possibility introduced by the *might* modality is picked up straight away and modified by a *would* sentence. A *might* sentence can also felicitously follow another *might* sentence; our semantics predicts modal subordination phenomena in that case as well. The updating of epistemic possibilities with non modal information, which may include binding information, allows us to quantify into epistemic possibilities as in:

(8)   A student just walked in. Pat might grant him an interview.

To see what happens, assume that the first sentence is of the form $\exists x W x$. So the existential quantifier will reset the value of x and the input context for $W x$ will be $\sigma \frac{a}{x}$. Suppose $\sigma \frac{a}{x}$ supports $W x$. Then the output context's epistemic possibilities will also support $W x$ with the assignment of a to x. The second sentence in (3.2) resets the epistemic possibilities, but since each one of these is a descendant of one in the input set of epistemic possibilities, this means that subsequent free occurrences of the variable x introduced by pronouns in the second sentence of (3.2) will carry the same value a.

On the other hand we cannot bind variables outside a modal context with quantifiers introduced inside a *might* or *would* operator.

(9)   # A wolf might have walked in. He bit you first.

Although the utterance of the *might* sentence resets the epistemic possibilities on the proposed semantics, it doesn't affect that actual world of evaluation or the actual assignment function. So the value that the variable is reassigned by the existential quantifier cannot be passed onto the occurrence of the variable introduced by the pronoun *he*.[4]

One of the salient features of this approach is that the possible interpretations of modal subordination sentences is quite restricted. Consider again (2c) and some variants (2d-e):

(2)   c.   A wolf might walk in. It might eat you first. But then it might not.

       d.   A tiger might walk in. Then a wolf might walk in. They might eat you first.

       e.   A wolf might walk in. It probably wouldn't eat you. But a tiger might walk in, and it definitely might eat you.

(2d) is an easy case for the present view. Since *might* modalities refine existing epistemic possibilities, this example shows how one might utterance can affect the interpretation of the second, and the result in this case is that the epistemic possibilities to be considered in interpreting the third clause are those where a wolf *and* a tiger walk in. A simple two place operator account like Roberts's cannot account for such examples, because it requires using both propositions under the scope of the *might* operators in the first and second clauses to interpret the third; Frank (1997) can but at the cost of stipulating an operation of propositional fusion on propositional anaphors that threaten to generate too many possible antecedents. As one example, consider (10). It looks as though the two place operator approach makes such examples acceptable when they are not, because the proposition that a wolf walks in is plainly available as an antecedent.

(10)   A wolf$_i$ might walk in. But then again there might not be any wolf. it$_i$ would scare you.

The present approach requires the interpretation of *would* on the standard dynamic semantics view of discourse interpretation to be constrained by the epistemic possibilities

---

[4]If we had developed our semantics with partial functions and had forced the pronoun to introduce an occurrence of the same variable as that introduced by the indefinite inside the modal, the discourse would be uninterpretable.

introduced by the *might* clause in the previous sentence. But these possibilities deny the existence of any wolf and this makes impossible the intended interpretation of the pronoun.

Another problem for something like Frank's account of modal subordination with the possibility of fusing propositional antecedents is (2e). The modal in the fourth clause in (2e) could take any of the propositions expressed by the last three sentences on the two place operator approach, whereas on the current approach the modal is constrained to take the last updated set of epistemic possibilities, predicting that there will be both a wolf and a tiger present. Salience concerns dictate that the tiger is the preferred eater. Now one could get the wolf to be the cause of the addressee's fright, if the context contained the information that wolves normally don't eat people unless there's a tiger present, in which case the wolves become very competitive and aggressive—and tend to eat people present.[5] We note that this interpretation comes for free on our approach, since all this information will be built into the context and can easily reverse the salience of the tiger in favor of the wolf. The two place operator approach would have to somehow include all this information in the first argument of the operator leading to complex merge operations over material both inside and outside modal operators, again seriously threatening overgeneration as well as being very *ad hoc*.

Let's now turn to an example that's difficult for our approach. Examples like (2c) are easily treated by 2 place operator approaches to these modalities where the first unspecified argument can be filled in by any anaphorically available proposition introduced in the discourse. On the standard dynamic semantic approach, utterances in a discourse are interpreted serially, and an input to the interpretation of the n+1st utterance in a discourse is an output from the interpretation of the nth utterance. So on this view of discourse interpretation, our proposal renders (2b) unsatisfiable, since the epistemic possibility that a wolf eats you first cannot be refined into one where it does not. So the present account is much more constrained as to which sequences of *might* and *would* sentences it countes as coherent. On the one hand, this is a good thing. As we've already argued, the two place operator approach along the dynamic semantic view that any previously introduced proposition can serve as an antecedent seriously overgenerates the possible readings for any modal subordination sentence.

On the other hand, the present account of modals has a problem since discourses like (2c) are perfectly interpretable. As the first author has argued at length elsewhere (Asher 1993, Asher and Lascarides 2003), what needs to change is the standard view of discourse interpretation in dynamic semantics: a more adequate view of discourse interpretation requires taking account of a much more elaborate notion of discourse structure, according to which new utterances may attach to several previous utterances in a discourse, depending on what the structure of the discourse is. The main constraint on attachment is called the *Right Frontier Constraint*: roughly only the constituents on the right hand side of the discourse structure (or more precisely those constituents including the last bit of information introduced into the structure and those consituents that dominate it[6]) are available for attachment. An example of where discourse struc-

---

[5]Thanks to Ken Safir for coming up with this example.

[6]See Asher (1993) and Asher and Lascarides (2003) for an extensive discussion of this constraint.

ture matters is (2c) in which a contrast is developed between the second and third clause, each of which elaborates on the first clause. Thus, we might imagine the discourse structure as an acyclic graph with the first clause at the root and the second and third clauses offering daughters to the root clause. the connections between the root and the daughters would be a discourse relation, the discourse relation of Elaboration. To represent the contrast we might link the two daughter nodes via a relation of contrast. Without going into details (see Asher 1993, Asher and Lascarides 2003), the nature of these connecting relations would dictate that the epistemic possibilities defined by the elaborating clauses would depend on (in fact "elaborate") the epistemic possibilities introduced at the root, while the semantics of the correction link would suggest that epistemic possibilities defined by one contrasting clause do not modify those of the other, precisely because they are contrasting, presenting distinct and different epistemic possibilities. Similarly in (10) the discourse structure is one where we have an elaboration of two distinct and contrasting possibilities. But because of the configuration of the discourse structure, the first constituent is not on the right frontier and no longer available for attachment. In the present theory, independently motivated facts about discourse structure do the work of Kratzer's modal restrictors, making for a more general and intuitive account.

Discourse structure also offers us the possibility of accounting for negated examples of modal subordination, which feature prominently in Frank's (1997) account of modal subordination.

(11) a.   I didn't buy a refrigerator. It would have taken up too much room.

  b.   I didn't buy a refrigerator. It would have cost too much.

These also don't follow straightforwardly on our account. But there are several options. One is to postulate a special discourse connection here according to which there is a causal link between the possibility described under the negation and the second clause containing the modal, which we gloss as $\neg p$, but if p then the result would have been q—or, p would have caused q. Furthermore, q in turn is a reason for not doing p. The pattern of a negated sentence discourse linked to a sentence with an epistemic modal suggests a type of elliptical explanation involving a conditional, a particular type in other words of discourse relation between two constituents in a discourse. This discourse link differs truth conditionally from the sort of narrative link one gets with the (2) examples. This kind of account is supported by the fact that negated modal subordination is not universally possible, as shown by the examples in (12), which lack the discourse connection we describe.

(12) a.   # I didn't buy a refrigerator. I would like it.

  b.   I didn't buy a refrigerator. I didn't like it.   (*de re* only)

Of course, the two place operator approach can also get the truth conditions right for examples like (11), but it fails to explain why contents under negation are special in a way that disjunctions are not:

(13) Either a wolf is walking in that door in ten minutes or a tiger is. It #would (will) eat you first.

14

The explanation for the special status of negation in these cases seems not to stem from the meaning of the operator itself but rather from the rhetorical structure that its uses in discourse suggest.

Finally, the flexibility of the anaphoric approach in fact makes it more difficult to give a compositional semantics for counterfactual conditionals. In counterfactuals, the "anaphoric antecedent" for the two place approach is in fact "grammatically" specified as the antecedent of the conditional; and this would have to be stipulated on a two place account. On the other hand, this falls out for free on the present one place operator account. That is, the grammatical connection between the antecedent and the consequent of a counterfactual furnish a discourse link between them; and so on this semantics, it will be the epistemic possibilities conditioned by the interpretation of the antecedent that the epistemic operator in the consequent clause will range over. But what epistemic possibilities are these? Before we can answer that question, we need to make another detour to get another two pieces of the semantics in place.

3.3 Conditionals

The idea of a compositional semantics for counterfactuals is to assign semantics to the various modals and to the *if ... then* expression within them and to see how those all combine together. We've now given the basic semantics for the modals *would* and *might* which figure typically in the consequent of counterfactuals. So now what sort of semantics ought the conditional itself to have? This is of course a question that many have worried over; but on the present view, the semantics of the conditional in counterfactuals should be exactly the same as that in so called "indicative" conditionals, in which the antecedent is in the indicative (as opposed to subjunctive) mood.

What features ought standard indicative conditionals to have? Gillies (2004) shows that they ought and can have the following features:

- they have Ramsey test like behavior; they should support *modus ponens*.

- they should support exportation

- they have a modal flavor: more specifically, they should obey the following equivalence: $\neg(\phi \Rightarrow \psi) \leftrightarrow (\Diamond(\phi \wedge \neg\psi)$

While the first two properties are obvious enough, perhaps we should provide some motivation for the modal properties of ordinary conditionals. Consider the following interaction from Gillies (2004) between two detectives, A and B, who are investigating a crime at a mansion for which grounds staff consists of a driver and a grounds keeper.

(14) a.   A: If one of the grounds staff committed the crime, then it was the driver.

      b.   B: No, it might have been the grounds keeper.

B's response *it might have been the grounds keeper* is a way of elaborating on his denial of A's claim; clearly we take B's response to explain *why* he thinks A's assertion is false. But if this is right, then the truth conditions of a conditional must involve

the semantics of modals in some way, and that in turn means that conditionals involve epistemic possibilities.

In Gillies's (2004) definition conditionals introduce tests on (our first component of) information states: a state σ will pass $\phi \Rightarrow \psi$ iff σ obeys the Ramsey test for this conditional—i.e. $\sigma + \phi + \psi = \sigma + \phi$. For us this doesn't reflect the fact that ψ can actually alter the information state—e.g., by having existential quantifiers in it. And we need to test epistemic possibilities in a given context not the discourse context itself. But we can get something equivalent by extending our notion of a descendant to a sequence of formulas. To say for example that σ has a ϕ, ψ descendant is just to say that σ has a ϕ descendant σ' and σ' has a ψ descendant σ''. The definition below ensures that all the epistemic possibilities of a given element of the discourse context together with the (actual) world and assignment support the conditional.

- $\sigma[\phi \Rightarrow \psi]_{\mathcal{A}}\sigma$ iff every ϕ descendant of σ has a ψ descendant and $\forall \epsilon \in \exists(\sigma)\forall \epsilon'$ such that $\langle \epsilon, \epsilon' \rangle \models_d \phi, \exists \epsilon'' \langle \epsilon', \epsilon'' \rangle \models_d \psi$.

Though the interpretation of the conditional is close to Gillies's or Veltman's, there are some advantages to the current approach. With Gillies's or Veltman's notion of a test, conditionals can't contribute to information growth, not unless we take a more complicated picture of what it is to learn new information (where we might distinguish between different information states and so consider sets of information states as inputs). By separating out elements of discourse contexts from epistemic possibilities we can "learn" (in our simplified sense here) conditionals, eliminating those worlds whose epistemic possibilities don't support the conditionals newly introduced into the discourse context. Because of the semantics of conditionals, updating the discourse context with a conditional will automatically be reflected within the epistemic possibilities permitted by the discourse context.

We can straightforwardly prove some desirable properties of our conditionals.

Fact 2 *Facts about conditionals.*

- $\phi \Rightarrow \psi, \phi \models \psi$

- $\phi \Rightarrow (\psi \Rightarrow \chi) \models (\phi \wedge \psi) \Rightarrow \chi$ *and vice-versa*

- *The deduction theorem holds for* $\Rightarrow$

- $\Diamond(\phi \wedge \neg\psi) \rightarrow \neg(\phi \Rightarrow \psi)$, *as long as* ψ *is not an epistemic formula.*

- $\neg(\phi \Rightarrow \psi) \rightarrow \Diamond(\phi \wedge \neg\psi)$

$\Rightarrow$ immediately verifies MP. Suppose that for an arbitrary σ, $\sigma[\phi \Rightarrow \psi]_{\mathcal{A}}\sigma$ and $\sigma[\phi]_{\mathcal{A}}\sigma'$. So σ has a ϕ descendant σ'; by the semantics of the conditional, every ϕ descendant of σ has a ψ descendant. So for some σ'', $\sigma'[\psi]_{\mathcal{A}}\sigma''$, which suffices to validate MP by our definition of logical consequence.

Considering exportation, suppose $\sigma[\phi \Rightarrow (\psi \Rightarrow \chi)]\sigma$ and suppose further that $\sigma[\phi \wedge \psi]\sigma'$, for some σ'. Now every ϕ descendant of σ must also have a $\psi \Rightarrow \chi$ descendant. And this implies that every ψ descendant of any ϕ descendant of σ has a χ descendant. Now σ' is a ψ descendant of a ϕ descendant of σ. So σ' has a χ

descendant. Since σ' was arbitrary, we've shown that every $\phi \wedge \psi$ descendant of σ has a χ descendant and thus $(\phi \wedge \psi) \Rightarrow \chi$. For the reverse direction, suppose we have $\sigma[(\phi \wedge \psi) \Rightarrow \chi]\sigma$ and suppose we have $\sigma[\phi]\sigma'$. We have to show that every ψ descendant of σ' has a χ descendant. Suppose not for some ψ descendant of σ', call it $\sigma*$. But then $\sigma*$ is a $\phi \wedge \psi$ descendant of σ and it must have a χ descendant.

Consider now the deduction theorem: $\Gamma, \phi \models \psi$ iff $\Gamma \models \phi \Rightarrow \psi$. From right to left is just an application of modus ponens, which we've already shown holds for $\Rightarrow$. Now from left to right, assume $\Gamma, \phi \models \psi$, which means that every Γ, ϕ descendant of some input σ has a ψ descendant. Now consider any Γ descendant of σ, call it σ' and consider any ϕ descendant of σ'; it also must have a ψ descendant. And since σ and σ' were arbitrary, we've shown $\Gamma \models \phi \Rightarrow \psi$.

To see how the modal connection works, suppose that $\sigma[\neg(\phi \Rightarrow \psi)]\sigma$. Then either some ϕ descendant of σ fails to have a ψ descendant or some $\epsilon \in \mathcal{E}(\sigma)$ has a ϕ descendant but no ψ descendant—in effect this is what $\models_d$ tells us. In the latter case, clearly σ supports $\Diamond(\phi \wedge \neg\psi)$. So suppose now σ has a ϕ descendant that has no ψ descendant. If ϕ is also non-modal then σ supports $\phi \wedge \neg\psi$ and eventually going through the recursion we will establish conditions in the epistemic possibilites that ensure that $\phi \wedge \neg\psi$, which entails $\Diamond(\phi \wedge \neg\psi)$. Now suppose that ϕ is modal but not a test—i.e. a *might* sentence. ϕ resets the epistemic possibilities in the output state and that output state supports $\neg\psi$, which since it is nonmodal must be reflected in the epistemic possibilities of the output context elements. So once again we see that $\sigma[\Diamond(\phi \wedge \neg\psi]\sigma$.

Discourse structure was important in our analysis of the simple modals *might* and *would*. It is also important in the analysis of some discourse patterns involving conditionals. For one thing it helps us explain the important connection between conditionals and disjunction. Zimmerman (2002) argues that disjunctions, at least when read with a falling intonation, present an exhaustive series of epistemic possibilities. Although that seems not be part of the "hard semantics" of disjunction (Asher and Bonevac (2003) present several counterexamples), it does appear to be a robust generalization about the discourse role of disjunctions. That is, disjunctions typically have, and in such contexts as (15) with the proper intonation must have, a role as presenting an exhaustive set of epistemic possibilities that includes the actual context of evaluation (in our sense the given world and assignment of some context element σ). Thus when (15a) is understood as having that particular discourse function, the inference from (15a) to (15b) is immediate, given the fact that the actual world and assignment are always part of the epistemic possibilities.

(15) a. Either one grounds keeper did it or the driver did.

      b. If the grounds keeper didn't do it, the driver did.

3.4 Irrealis Markers

With an account of conditionals in place, we're ready to tackle the last bit needed to define counterfactuals. The antecedent of a counterfactual in Romance and Germanic languages has some sort of "irrealis" operator—either introduced by the subjunctive mood or a special tense morpheme like the imparfait in French.[7] The irrealis operator

---

[7]See Iatridou 2000 for a discussion of the connection between tense and counterfactuality.

in the antecedents of counterfactuals introduces, we hypothesize, an imaging operator $\bullet$ on context elements, which is successful (i.e. every $\sigma \in \bullet(\sigma, \|\phi\|)$ has a ϕ descendant) and preserves as much information about what normally goes on in σ when ϕ holds as possible. Let $\mathcal{E}(\bullet(\sigma, \|\phi\|))$ denote a set of epistemic possibilities based on $\bullet(\sigma, \|\phi\|)$. Then

- The irrealis operator (first pass):
  $\sigma[\text{irr}\phi]_{\mathcal{A}}\sigma'$ iff $\sigma' = \langle 1(\sigma), 2(\sigma)\mathcal{E}(\bullet(\sigma, \|\phi\|))\rangle$

The irrealis operator moves the set of epistemic possibilities in an input element to those that are based on a set of normal worlds, in particular those normal worlds that satisfy the formula in its scope. The irrealis operator also carries with it a presupposition or perhaps more accurately the defeasible implicature that the proposition under its scope is not true at the actual world assignment pair—in other words it has a " contrary to fact" implicature. The semantics of *irr* is quite weak. But note that the irrealis operator doesn't have a stand alone use, that is a use outside the scope of some conditional or other non truth functional operator:

(16)   # Chris were to come to the party.

This isn't really explained on the present account, since we have other operators like *might* that reset epistemic possibilities that do have a stand alone use. However, it appears that when we look to the use of irrealis operators at least in conditionals and perhaps elsewhere that they also impose a requirement on the epistemic possibilities associated with the input state: not only does the irrealis operator modify the input state's epistemic possibilities, but it also requires that the input state's unmodified epistemic possibilities have a similar modal character with respect to the normal worlds that the irrealis operator makes use of. That is, the irrealis operator shifts the actual world's epistemic possibilities, but it also requires that all the actual world's unaltered epistemic possibilities would shift in the same way. Such semantic behavior makes the irrealis operator suited to go with an epistemic operator like the conditional that holds not only in the actual world but in the actual world's epistemic possibilities. It also means that the irrealis operator brings with it a condition of lawlikeness: if $\bullet$ gives us the normal dispositions and regularities with respect to some some proposition $\|\phi\|$ and some world σ, *irr* ϕ constrains the input state σ's epistemic possibilities to have the same dispositions and regularities with respect to $\|\phi\|$ as they do in σ.

Accordingly, we'll revise the semantics of the irrealis operator. Let
$\bullet(\epsilon, \|\phi\|) = \bigcup\{\bullet(\sigma', \|\phi\|) : \sigma' \in \epsilon\}$
Then:

- The irrealis operator (revised):

- $\sigma[\text{irr}\phi]_{\mathcal{A}}\sigma'$ if $\sigma' = \langle 1(\sigma), 2(\sigma)\mathcal{E}(\bullet(\sigma, \|\phi\|))\rangle \wedge$
  $\qquad \forall \epsilon \in 3(\sigma) \; \bullet(\epsilon, \|\phi\|) \subseteq \bullet(\sigma, \|\phi\|)$

- $\sigma[\text{irr}\phi]_{\mathcal{A}}0$ otherwise.

18

4 Counterfactuals Compositionally

We are finally now in a position to treat "present" counterfactuals. Their semantics is the result of applying the conditional operator to two arguments, the first modified by the irrealis operator, the second modified by an epistemic modal—either *might* or *would*—or some other modal operator like *should*. (17a), for instance, is analyzed below as (17b):

(17) a. If I were not to sleep tonight, I would topple over tomorrow.

b. *irr*(I not sleep tonight) $\Rightarrow$ *would* (I topple over tomorrow)

Such a sentence is satisfied in a structure at an input context or information state just in case revising the input context with my not sleeping tonight makes true that I would topple over tomorrow, and the latter will be true in the adjusted descendant of the input just in case all of the epistemic possibilities of that descendant make my toppling over true.

Let's first compare how this account's interpretation of the unary modals improves the prospects for a compositional semantics for counterfactuals. The anaphoric two place operator account of *would* must stipulate that the first argument of *would* is in fact "grammatically" specified in a counterfactual. The fact that the semantics for *would* must take the antecedent into account falls out for free on the present one place operator account. Further, our account shows why a *would* or *might* is necessary in a counterfactual construction: the irrealis shifts the epistemic possibilities but not the actual world or assignment, and in order for the proposition under the scope of the irrealis operator in the antecedent of the conditional not to be irrelevant to the evaluation of the consequent, we must use an epistemic modal to test the consequent in the relevant set of worlds. As far as we know, no other account of the counterfactual explains why modals must occur in the consequent of a counterfactual.

Let's turn to some logical properties of the counterfactual as we've defined it. Our counterfactual obviously licenses MP in the following sense: if we update with a counterfactual ($irr(\phi) \Rightarrow would(\psi)$) and irr($\phi$), the semantics of $\Rightarrow$ immediately yields *wouldψ*. But we've already noted that this would be an odd discourse sequence, since an assertion of *irrϕ* is infelicitous. It doesn't license MP with a simple update of ϕ, because the revision operation goes to the normal antecedent worlds and the actual world of evaluation may not be such a normal world. Further, an inference to the consequent of $\Rightarrow$ yields only the conclusion that ψ is true in all the resulting epistemic possibilities—not that it is in fact true. Recall that *wouldϕ* does not imply ϕ, though it clearly doesn't imply $\neg\phi$ either. The same thing goes with exportation and other inferences that the conditional $\Rightarrow$ licenses. The inferences break down because of the interference of the semantics of the irrealis and *would* operators.

Our counterfactual might license MP defeasibly if we adapted something like Asher and Morreau's (1991) or Morreaus's (1998) notion of defeasible entailment. In fact, this raises the question of the link between counterfactuals and normality conditionals.

Normality conditionals furnish a distinct reading of *if ..., then* for indicative, or epistemic conditionals. They are also associated with generic readings of bare plurals or other types of noun phrases (Krifka et al. 1992, Pelletier and Asher 1997).

Like counterfactuals, they exploit a notion of normality. But they are different as well. The normality conditional > is defined directly in terms of the operator •. Adapting to the dynamic account here, Asher and Morreau's (1991) basic semantics for > is $\sigma[\phi > \psi]_{\mathcal{A}}\sigma$ iff $\forall \sigma \in \bullet(\sigma, \|\phi\|) \exists \sigma'\langle \sigma, \sigma'\rangle \in \|\psi\|$ or equivalently every element in•$(\sigma, \|\phi\|)$ has a ψ descendant . Nonmonotonic inference using > essentially relies on a transformation of > statements into → statements provided this operation is consistent with all the information given in the premises for the inference. The proceedure developed in Morreau 1995 and refined in Asher and Mao 1997 works on the outermost > in the sentences of a set of premises and attempts to convert to a material implication → all the > with the widest scope not under the scope of negation in the premises, and then checks the result for consistency or satisfiability. All > statements with logically equivalent antecedents have to be considered together: either they can all be consistently turned into → statements or none of them can. Roughly a set of sentences Γ nonmonotonically implies a conclusion ϕ just in case, roughly, all acceptable ways of turning > statements in Γ into → statements, call these $\Gamma^{\rightarrow}$, are such that $\Gamma^{\rightarrow} \models \phi$. To give a very simple example in the propositional > logic, suppose that the premises Γ we have for our inference are $A > D, A > B, A$. To reason nonmonotonically with Γ, we have to add to Γ the pair $(A > D) \rightarrow (A \rightarrow D)$ and $(A > B) \rightarrow (A \rightarrow B)$ and check whether the result is consistent. It is, and since there are no other > to turn into →, we see that we can nonmonotonically infer B and D. A slightly more complex example is one where we suppose that the premises $\Gamma = \{A > D, A > B, B > C, A, \neg C\}$. There are three > statements all of which have a > with maximal scope. Now we could either consider whether adding $(B > C) \rightarrow (B \rightarrow C)$ to Γ preserves consistency or whether the pair $(A > D) \rightarrow (A \rightarrow D)$ and $(A > B) \rightarrow (A \rightarrow B)$. We can add either the A transformations or the B transformations and preserve consistency but not both. So from Γ nothing nonmonotically follows that doesn't follow from Γ in the underlying monotonic logic for our propositional > language.

The revised interpretation of the irrealis operator makes counterfactuals more modally robust than normality conditionals. For the counterfactual $irr(\phi) \Rightarrow would (\psi)$, to hold at σ, not only must the epistemic possibilities based on •$(\sigma, \|\phi\|)$ support ψ, but the epistemic possibilities ϵ in σ must be such that the epistemic possibilities based on •$(\epsilon, \|\phi\|)$ support ϕ. This makes a certain amount of sense, if we follow the idea that counterfactuals are lawlike and must hold not only in the actual world but in all other (in our case epistemically close) worlds as well. Normality conditionals describe modal facts about worlds, but they may be modally fragile—they may only hold at a given world.

The exact relation between normality conditionals and counterfactuals depends crucially on how the irrealis operator in the formula $irr\ \phi$ exploits the normal ϕ worlds in the epistemic possibilities $\mathcal{E}(\bullet(\sigma, \|\phi\|))$ it introduces. The first possibility is one where the epistemic possibilities are "cautious" or not completely omniscient with respect to the modal facts about normality.

- Lack of omniscience about normality: •$(\sigma, \|\phi\|)) \subseteq \bigcup(\mathcal{E}(\bullet(\sigma, \|\phi\|))) \subseteq \|\phi\|$.

This constraint leads to the following immediate fact.

Fact 3 *If Lack of omniscience holds, then* $\models (\mathrm{irr}(\phi) \Rightarrow would (\psi)) \rightarrow (\phi > \psi)$.

If $irr(\phi) \Rightarrow would(\psi)$, then every epistemic possibility based on $\bullet(\sigma, \|\phi\|)$ supports ψ. Given the Lack of Omniscience, this means $\bullet(\sigma, \|\phi\|)$ supports ψ.

Given Lack of Omniscience, we could simply use the nonmonotonic inference system for normality conditionals to deduce defeasible conclusions from counterfactuals. The defeasible inferences from counterfactuals would be exactly those that follow from the associated normality conditionals. In particular following Morreau (1995) or Asher and Morreau (1991), we predict the following sorts of reasoning patterns to be defeasibly valid for counterfactuals. To simplify notation we will use the $\square\rightarrow$ operator to symbolize the counterfactual

Fact 4 *Given Lack of Omniscience and letting $\approx$ stand for the defeasible, nonmonotonic inference relation based on normality conditionals,*

- *Defeasible Modus Ponens: $\phi\square\rightarrow\psi, \phi \approx \psi$, but not: $\phi\square\rightarrow\psi, \phi, \neg\psi, \approx \psi$*

- *Defeasible Transitivity: $\phi\square\rightarrow\psi, \psi\square\rightarrow\chi, \phi\approx\chi$, but not: $\phi\square\rightarrow\psi, \psi\square\rightarrow\chi, \phi, \neg\chi, \not\approx\chi$*

- *Defeasible Modus Tollens: $\phi\square\rightarrow\psi, \neg\psi \approx \neg\phi$, but not: $\phi\square\rightarrow\psi, \phi, \neg\psi, \approx \neg\phi$*

- *Penguin Principle: Provided $(p \subseteq q \wedge \bullet(\sigma, p) \cap \bullet(\sigma, q) = \emptyset) \rightarrow \bullet(\sigma, q) \subseteq p^c$ where p^c is the complement of the p elements,*

$$\frac{\vdash \phi \rightarrow \psi}{\phi\square\rightarrow\chi, \psi\square\rightarrow\neg\chi, \phi \approx \chi}$$

The Penguin Principle together with the frame condition on $\bullet$ that allows us to derive the Penguin Principle resolves a notorious problem with Lewis's semantics for counterfactuals.

(18) a.   If I were to drop the glass, it would break.

b.   If I were to drop the glass and the floor was thickly carpeted, it would not break.

Counterfactuals like normality conditionals allow for sets of consistent counterfacutals with conflicting consequents and antecedents that are linked by entailment. Now what inferences should we draw when we are in a situation where I drop the glass and the floor is thickly carpeted? Notoriously, Lewis's account of counterfactuals predicts an inconsistency in this case. The two normality conditionals entailed by the counterfactuals in (18), on the other hand, make the plainly sensible prediction that the glass does not break. Specificity is a constraint on $\bullet$ that ensures that $>$ conditionals with logically more specific antecedents always take priority in the case of conflicts between defaults. For instance, if I have a theory in which there are two defaults—e.g., that birds fly but that penguins don't and that it is a taxonomic fact or axiom of the theory that penguins are birds, then specificity together with the nonmonotonic inference proceedure defined for $>$ predicts that if I have something that is a penguin (and hence also a bird), I will

infer defeasibly that that bird doesn't fly.[8] Intuitively, counterfactuals pattern with the predictions of normality conditionals, as this analysis predicts.

The Lack of Omniscience constraint also places counterfactuals in a nice position along the scale of conditionals. Conditionals that are true at every world, necessary conditionals, clearly entail counterfactuals on my semantics, while counterfactuals, as we've just seen, entail normality conditionals. Neither counterfactuals nor normality conditionals entail nor are they entailed by the associated material conditional, though a normality conditional defeasibly entails the associated material conditional. And finally neither counterfactuals nor normality conditionals entail nor are they entailed by the indicative conditional $\Rightarrow$, though $\Rightarrow$ is a semantic constituent of a counterfactual.

With respect to the monotonic axioms and inference rules for counterfactuals, Lack of Omniscience on gives us at the following, sensible but minimal logic for counterfactuals:

Axioms for Counterfactuals:

- Idempotence: $\phi \boxright \phi$

- Necessitation:
$$\frac{\vdash \phi}{\vdash \psi \boxright \phi}$$

- Logical Closure on the right:
$$\frac{\vdash \psi \to \chi}{(\phi \boxright \psi) \vdash (\phi \boxright \chi)}$$

- Logical Equivalence on the Left:
$$\frac{\vdash \psi \leftrightarrow \chi}{(\psi \boxright \phi) \to (\chi \boxright \phi)}$$

- And: $((\phi \boxright \psi) \wedge (\phi \boxright \chi)) \to (\phi \boxright (\psi \wedge \chi))$

- $(\neg\phi \boxright \phi) \to (\psi \boxright \phi)$

Missing is Lewis's notorious axiom $(\phi \wedge \psi) \to (\phi \boxright \psi)$. We also don't have $((\phi \boxright \psi) \wedge \neg(\phi \boxright \neg\chi)) \to ((\phi \wedge \chi) \boxright \psi)$.

To get a stronger logic, a natural thing to do is to strengthen the Lack of Omniscience constraint so that the epistemic possibilities based on the modal facts about normality capture exactly those facts.

- Epistemic tracking of normality: $\bullet(\sigma, \|\phi\|) = \bigcup(\mathcal{E}(\bullet(\sigma, \|\phi\|))) \subseteq \|\phi\|$.

---

[8]The rule capturing this frame condition in Asher and Morreau 1991 is:
$$\frac{\vdash \phi \to \psi}{((\phi > \chi) \wedge (\psi > \neg\chi)) \to (\psi > \neg\phi)}$$
.

We can now impose a number of constraints on $\bullet$ to get different behaviors for the counterfactual. We could also add constraints to $\bullet$ as in Asher and Morreau (1991). If we add $\bullet(\sigma, p \cup q) \subseteq \bullet(\sigma, p) \cup \bullet(\sigma, q)$, we will verify as an axiom $((\phi \square\rightarrow \psi) \wedge (\chi \square\rightarrow \psi) \rightarrow ((\phi \vee \chi) \rightarrow \psi)$. If we add $(\bullet(\sigma, p) \subseteq q \wedge \bullet(\sigma, q) \subseteq p) \rightarrow \bullet(\sigma, p) = \bullet(\sigma, q)$, then we validate the axiom, $((\phi \square\rightarrow \psi) \wedge (\psi \square\rightarrow \phi)) \rightarrow ((\phi \square\rightarrow \chi) \leftrightarrow (\psi \square\rightarrow \chi))$. Note also that of course Facts 3 and 4 still hold if we replace the Lack of Omniscience Constraint with the Epistemic Tracking Constraint. And the Epistemic Tracking Constraint retains the set of relationships between the counterfactual, normality conditional, material conditional and the epistemic conditional that followed from the Lack of Omniscience Constraint.

On the other hand, we might assume that the epistemic possibilities are adventurous with respect to the normality facts. Namely, not only do the epistemic possibilities based on the normal ϕ worlds capture all of the modal facts about normality; they go further and introduce more information.

- Epistemic adventurousness about normality: $\bigcup(\mathcal{E}(\bullet(\sigma, \|\phi\|))) \subset \bullet(\sigma, \|\phi\|)$.

Although the axioms and rules for counterfactuals just given still hold, we can no longer infer the connection between counterfactuals and normality conditionals. This makes reasoning nonmonotonically with the counterfactual *sui generis*. And this calls for a much more involved investigation of nonmonotonic reasoning, as the link between antecedent and nonmodalized consequent of a counterfactual is a little more indirect than that between the antecedent and the consequent of a normality conditional. For one thing, the consequent of a counterfactual holds only in the epistemic possibilities provided for by the irrealis shift in the antecedent. In effect counterfactuals are sensitive to the epistemic possibilities furnished by the normal worlds where the antecedent is true; and thus, we might imagine that such counterfactuals are sensitive to a background theory or beliefs, whereas normality conditionals are less so. This means that to license MP even defeasibly for counterfactuals in the sense that $irr\phi \Rightarrow would\psi, \phi$ defeasibly yields ψ, we must first assume some sort of defeasible inference from ϕ to $irr\phi$, and secondly we must assume an *additional* defeasible inference from $would\phi$ to ϕ. The second of these looks quite plausible: it is sort of a reliabilist constraint on our epistemic possibilities—that they normally track the truth. Sometimes, however, our epistemic possibilities can go astray, in that what we take to be ways the actual world could be might all turn out to be unfaithful to the way the actual world is. This defeasible reliabilism is what we can capture in this framework by adding $>$ into our set of modal operators.[9]

More vexing is the defeasible inference from ϕ to $irr\phi$. This seems to go against the contrary to fact presupposition of the irrealis operator. On the other hand, if we take the contrary to fact information conveyed by the irrealis operator to be a defeasible

---

[9]The problem with nonmonotonic inference couched in this way and in most of the standard ways, is that we require a sound and complete proof theory for the underlying monotonic semantics for $>$ and other modal operators. Once we introduce several complex modal operators, it becomes practically difficult to get such a logic. For the defeasible inferences in this system, we will simply treat all the other modals besides $>$ as inert, and we will assume that the only sets of premises that we will be working with are at most Π_1. In this case consistency tests are possible—and an axiomatization of $>$ together with basic dynamic logic is in fact feasible. See Morreau's 1992 thesis, Asher and Morreau 1991.

implicature, then we can accommodate this information together with the defeasible inference ϕ to $irr\phi$ in certain circumstances—viz., those in which both the counterfactual and its antecedent are both true. In effect, we have here a particular circumstance of use in which a particular default is "blocked". This phenomenon is quite common in lexical semantics, but modelling it satisfactorily requires a rather sophisticated non-monotonic logic, especially when the language is rich enough to express defeasible conditionals (Asher and Lascarides 2001, Asher and Mao 1997). Using the semantics for > and the notion of defeasible inference appealed to above, all of the following are consistent.

- Defeasible Reliabilism: $would\phi > \phi$

- Contrary to fact defeasible implicature of irr : $\top > (irr\phi > \neg\phi)$

- Suspension of implicature in cases the antecedent of a counterfactual is true:
  $(((\text{irr}(\phi) \Rightarrow would(\psi)) \wedge \phi) > (\neg(irr\phi > \neg\phi) \wedge (\phi > irr\phi))$

Together these constraints will license the defeasible inference from $(\text{irr}(\phi) \Rightarrow would(\psi)) \wedge \phi$ to ψ, if we impose the specificity constraint on •, since the Suspension constraint is more specific that the Contrary to Fact implicature. We will now be able to infer from ϕ irr ϕ and by MP with $\Rightarrow$ $would$ ψ. By Defeasible Reliabilism, we'll infer defeasibly ψ. But a similar defeasible modus tollens inference with counterfactuals won't work, since the Suspension axiom can't be used to get the appropriate default irr $\phi > \phi$.

What about other forms of nonmonotonic inference with counterfactuals? To capture the intuitively obvious Penguin Principle, we would need additional axioms to block the Suspension axiom in the relevant cases—namely we want to block the Suspension axiom when it is applied to a counterfactual $\phi\square\rightarrow \psi$ and we have $(\phi \wedge \chi)\square\rightarrow \neg\psi$ and $\phi \wedge \chi$. We could write such axioms but all of this seems to be a leading to Rube Goldberg device that does something that a much simpler and more elegant assumption, namely either the Lack of Omniscience or the Epistemic Tracking constraints, accomplish.

5 Additional Wrinkles

5.1 Discourse Sensitivity of Counterfactuals

Kai von Fintel (2002) has developed a semantics of counterfactuals to deal with the following asymmetry noticed by Heim.[10] The Lewis examples:

(19) a. If the USA threw its weapons into the sea tomorrow, there would be war; but if all the nuclear powers threw their weapons into the sea tomorrow, there would be peace.

b. If all the nuclear powers threw their weapons into the sea tomorrow, there would be peace, but if the USA threw its weapons into the sea tomorrow, there would be war.

---

[10]As reported in von Fintel (2002).

For many speakers, the asymmetry is in any case quite delicate: replace *the USA* in (19b) by *only the USA* and the asymmetry disappears.To get these asymmetries out of the basic dynamic semantics, we need to complicate considerably our account of the conditional or the irrealis operator. Since similar effects can be observed with indicative conditionals, it would appear that it is the conditional operator should receive a more complex treatment.

(20) a. If the USA throws its weapons into the sea tomorrow, there will be war; but if all the nuclear powers throw their weapons into the sea tomorrow, there will be peace.

     b. If all the nuclear powers throw their weapons into the sea tomorrow, there will be peace, but if the USA throws its weapons into the sea tomorrow, there will be war.

Are the costs worth it? In this approach, we can take a different tack, which we mention briefly here. The option we have in mind here is this: it involves the claim that these examples are judged differently in terms of acceptability because of the way they attach at the level of discourse structure. In (19a) the contrast is easier to make between the (just) the US throwing its weapons into the sea and all the countries throwing their weapons into the sea. Lewis's pair of examples is simply a special case of contrasts that are known not to be commutative. As more compelling evidence for this approach consider the following examples.

(21) a. When the USA threw its weapons into the sea, there was war. But when all the nuclear powers threw their weapons into the sea, there was peace.

     b. ?When all the nuclear powers threw their weapons into the sea, there was peace. But when the USA threw its weapons into the sea, there was war.

     a. The US's throwing its weapons into the sea tomorrow is a cause for war; all the nuclear powers throwing their weapons into the sea tomorrow is a cause for peace.

     b. All the nuclear powers throwing their weapons into the sea tomorrow is a cause for peace., but the US's throwing its weapons into the sea tomorrow is a cause for war.

Many people find this minimal pair of examples to exhibit a similar asymmetry. But here we can't blame the asymmetry on the semantics of the conditional introduced by *if ... then*. Once again a solution based on discourse structure appears to be an attractive alternative. In all these cases, we are linking complex constituents by means of the discourse relation Contrast. As the first author argues in Asher (1993), when linking complex constituents Contrast requires a partial isomorphism between the discourse and predicational structures of the two constituents it links. To make a long story about these cases short, we pair up the antecedents and the consequents of the conditionals or the two subordination and the two main clauses of the main clauses and compare the constituents. The paired constituents must have some common or at least linked topical elements and some distinct bits that combine wtih the common elements and one pair

at least must involve an attribution of a property P and an attribution of some property that one would expect is incompatible with P. These examples show that distinct bits of information in the paired constituents are understood as distinct depending on the order in which they are presented: after one has introduced all the nuclear powers in one constituent, to say something about the US in the second element of the contrasting pair doesn't give us a distinct element (the US is just one of the nuclear powers), whereas mentioning the US first suffices, perhaps for Gricean reasons of manner (don't say more than you need to) to make it distinct from the mention of all the nuclear powers in the second pair of the contrasting constituents.

An appeal to contrasting discourse structures, even if the appeal needs to be developed in much more detail, is quite promising for the asymmetries noticed by Lewis for counterfactuals. But there are other pairs of counterfactuals that cause difficulty for our account. Here is one pair of examples that our account doesn't get right.[11]

(22) a.   If John were to come to the party, it would be fun.

       But of course.If John were to come to the party, he might have a heart attack and that wouldn't be fun.

If (22a) is true, then there aren't any epistemic possibilities in which John gets a heart attack and the party isn't fun. So this makes (22b) false. But intuitively both of these can be true. Such examples provide the founding intuitions for an accommodation account of *might*, according to which *mightϕ* expands the epistemic possibilities to make ϕ true. Our account of *might* is more restrictive than such an account (which threatens to make every *might* sentence that's logically consistent true!) One thing to note about such examples, however, is that they have a different flavor from the run of the mill specificity examples in (18); in (22b) we need the *but of course* to make the pair felicitous. Very speculatively, perhaps this means that we should understand (22b) as a modification of (22a) and a correction of the epistemic possibilities presented by John's coming to the party. More radically, we might propose that *would* doesn't canvass *all* of the epistemic possibilities given by the input state, but just almost all of them. Such a proposal would largely leave the logic of *might* and *would* intact, except that the semantic explanation of the incompatibility between *wouldϕ* and *might$\neg\phi$* would fail. If your intuitions can accomodate this, then this proposal might be the most reasonable to handle pairs of examples like (22).[12]

5.2   Modality and Time

A last wrinkle we'll mention here (actually it's much more than a wrinkle) is the interaction between modality and time. Modals can combine with tense as in (23) below.

(23)   John might have come to the party.

But how exactly does tense combines with modality? This is a big issue with lots written about it both in logical domains (cf. the literature on historical necessity) and also in linguistics (e.g., Abusch 1997). To address this issue, we need to extend our

---

[11]Thanks to Thony Gillies for these examples.
[12]thanks to Josh Dever for this point.

semantics to include a set of times. We propose to index epistemic possibilities to a time. Thus, we'll extend our structures to include a strict linear ordering of times and index each set of epistemic possibilities to a time.

From the surface structure of (23), you might suppose that the modals *would* and *might* scope over the tense feature and so its logical form would look like this:

(23') Might(Past(John come to the party))

Using a standard semantics for Past, we would predict that such a sentence yields a descendant for σ iff John's coming in the past is an epistemic possibility. But these are the wrong truth conditions. Suppose we know now that John didn't go to the party; it may still be true that he might have gone. What this means is that there was a past set of epistemic possibilities, in which Might(John goes to the party) is supported. Hence, it appears that contrary to what the surface order of the words suggests, the scopes of the relevant operators in (23) is:

(23'') Past(Might(John come to the party))

So the truth conditions for such sentences are that for some past time t, the set of epistemic possibilities at t supported John's coming to the party; more technically $\sigma(might(\text{Past}(\phi)))\sigma'$ iff σ''s epistemic possibilities at t, for some past time t, all support John's coming to the party, provided that there was in σ an epistemic possibility at t of John's coming to the party.

It appears possible to access both past and present epistemic possibilities:

(24) The grounds keeper might have committed the crime. He might strike again soon.

This suggests that indeed $\mathcal{E}$ be a function on times, perhaps from times and context elements to a set of sets of context elements—-a very classical view. We leave that for future work. (the problem is one of resetting).

6 Conclusions

There are lots of modals that we might include in this account. *must*, *ought ϕ*, *should ϕ* all suggest universal quantifications over *deontic* possibilities while *may*, as Kamp (1973) suggested, introduces deontic possibilities. It's a delicate matter to ground the deontic alternatives in the epistemic possibilities (Asher 1987). But it appears that the following approach, on which we extended a dynamic semantics with a dynamic account of modals, can accommodate these modals as well. The semantics we have developed handles both the modal subordination facts and Veltman-style update phenonmena and it provides a compositional account of counterfactuals that has at least some pleasing features, including a connection to normality conditionals and the non-monotonic notion of inference that has come to be associated with them.

7   References

Abusch, D. (1997) "Sequence of Tense and Temporal *De Re*", in *Linguistics and Philosophy*, 20.1, pp. 1-50.

Asher, N. (1993): *References to Abstract Objects in Discourse*, Kluwer Academit Publishers.

Asher, N. and D. Bonevac (2003): Free Choice Permission is Strong Permission, *Synthese*, 2003.

Asher, N. and A. Gillies (2003): Common Ground, Corrections et Coordination, *Argumentation* 17, 481-512.

Asher, N. and A. Lascarides (2001): Indirect Speech Acts, *Synthese* 128, 2001, pp. 183-228.

Asher, N. and A. Lascarides (2003): *Logics of Conversation*, Cambridge University Press.

Asher, N. and Y. Mao (2000): . Negated Defaults in Commonsense Entailment, Bulletin of the Section of Logic, 30, 2001, pp. 41-60. Also in *Proceedings of Logic Colloquium 2000*.

Asher, N. and E. McCready (2004): "Modality and Evidence: A study of modals and emotives in Japanese and Hindi", manuscript.

Asher, N. and M. Morreau (1991): "Commonsense Entailment: A Modal Theory of Non-Monotonic Reasoning", *IJCAI 1991*, pp. 387-392 Asher, N. and L. Wang (2003): "Ambiguity et Anaphora and Plurals in Discourse", *Proceedings of Semantics and Linguistic Theory (SALT) XIII*

Fernando, T. [1993] " Generalized quantifiers as second order programs—'dynamically' speaking, naturally", *Proceedings, Ninth Amsterdam Colloquium*, University of Amsterdam.

Frank, A. (1997) Context Dependence in Modal Constructions. Ph.D. thesis. University of Stuttgart, Germany.

Gillies, A. (2003): 'Modal Scorekeeping and 'Might' Counterfactuals, in P. Dekker (ed.) *Proceedings of Fourteenth Amsterdam Colloquium*, ILLC Publications, Amsterdam

Gillies, A. (2004): 'Epistemic Conditionals and Conditional Epistemics', *Nous*, forthcoming.

Groenendijk, J. and M. Stokhof (1991): " Dynamic Predicate Logic", *Linguistics and Philosophy* 14, pp. 39-100.

Groenendijk, J., M. Stokhof and F. Veltman (1994): "This Might Be It", manuscript.

Iatridou, S. (2000): The Grammatical Ingredients of Counterfactality. *Lingiustic Inquiry 31.2*

Kamp, H. and Reyle, U. (1993), *From Discourse to Logic: Introduction to Model Theoretic Semantics of Natural Language, Formal Logic and Discourse Representation Theory,* Dordrecht; Boston: Kluwer Academic.

Kratzer, A. (1977): "What 'must' and 'can' must and can mean", *Linguistics and Philosophy*, pp. 337-355.

Krifka, M., G. Carlson, and F.J. Pelletier (1995): "Introduction", in G. Carlson and J. Pelletier (eds.), The Generic Book (Chicago: University Press, 1995).

Morreau, M. (1995): "Allowed Arguments", *IJCAI 1995*, pp. 1466-1473.

Pelletier, F.J, and N. Asher (1998): "Generics and Defaults", in J. van Bentham and A. ter Meulen (eds.), *Handbook of Logic and Language*, Elsevier, pp. 1125-1175.

Roberts, C. (1987): *Modal Subordination, Anaphora, and Distributivity*, Ph.D. Dissertation, Department of Linguistics, University of Massachussetts at Amherst.

Veltman, F. (1985): *Logics for Conditionals*. Ph.D. Dissertation, Department of Philosophy, University of Amsterdam.

Veltman, F. (1996): "Defaults in Update Semantics", *Journal of Philosophical Logic* **25**, pp.225-261.

von Fintel, K. (2002): 'Counterfactuals in a Dynamic Context', in M. Kenstowicz (ed.), *Ken Hale: A Life in Language*, MIT Press.