Reference to Possible Worlds[1]

Matthew Stone
Center for Cognitive Science and Department of Computer Science
Rutgers University
152 Frelinghuysen Road, Piscataway NJ 08854-8020
mdstone@ruccs.rutgers.edu

Summary

In modal subordination, a modal sentence is interpreted relative to a hypothetical scenario introduced in an earlier sentence. In this paper, I argue that this phenomenon reflects the fact that the interpretation of modals is an ANAPHORIC process. Modal morphemes introduce sets of possible worlds, representing alternative hypothetical scenarios, as entities into the discourse model. Their interpretation depends on evoking sets of worlds recording described and reference scenarios, and relating such sets to one another using familiar notions of restricted, preferential quantification. This proposal relies on an extended model of environments in dynamic semantics to keep track of associations between possible worlds and ordinary individuals; it assumes that modal meanings and other lexical meanings encapsulate quantification over possible worlds. These two innovations are required in order for modals to refer to sets of possible worlds directly as static objects in place of the inherently dynamic objects—quite different from the referents of pronouns and tenses—used in previous accounts. The simpler proposal that results offers better empirical coverage and suggests a new parallel between modal and temporal interpretation.

1 Introduction

Modal statements in natural language admit quite precise construals in context. Discourse (1) illustrates this.

(1) a I want to hold a barbecue.
  b Some vegetarians may be coming.
  c What can I do for them?

[1]Thanks to Mark Steedman and Dan Hardt for extensive comments, and Filippo Beghelli, Maria Bittner, Anoop Sarkar, Roger Schwarzschild, Beverly Spejewski, and pragmatics and computational semantics seminars at Penn for helpful discussion. This work was supported by an NSF graduate fellowship, an IRCS graduate fellowship, and a postdoctoral fellowship from RUCCS, as well as NSF grant IRI95-04372, ARPA grant N6601-94-C6043, and ARO grant DAAH05-94-G0426. This paper is a revised submission to *Linguistics and Philosophy* of the original paper *The Anaphoric Parallel between Modality and Tense* of May, 1997. April 3, 1999.

To explain the meaning of (1c), general paraphrases such as that in (2) are unsatisfying.

(2) What actions for me to take towards vegetarians are logically possible?

Sentence (1c) means something much more specific. A better paraphrase is given in (3).

(3) Assuming vegetarians come to my barbecue, what actions will be available to me (in virtue of the properties of vegetarians and barbecues in the actual world) that will contribute towards making the event a success?

The context refines the interpretation of (1c) in at least three respects. First, context indicates that only a certain hypothetical prospect is under consideration: what to do if the vegetarians come to the barbecue. Second, context dictates that certain other general facts are to be taken as given in considering what is possible—facts about what vegetarians and barbecues are usually like. Third, context suggests that any possible action must be assessed for its relevance to my intention to have a successful barbecue.

In a seminal series of papers (notably [Kratzer, 1977], [Kratzer, 1981], and [Kratzer, 1991]), Kratzer characterizes and formalizes these contextual dependencies. Her tools are an ontology of alternative, total possible world histories, and two parameters that can be to supplied to modal quantifiers over such histories. One, the MODAL BASE, describes the set of possible worlds under general consideration. For (1c), the modal base picks out those worlds where vegetarians come to the barbecue, and that otherwise respect our knowledge and circumstantial constraints. The other parameter, the ORDERING SOURCE, describes the plausibility and present relevance of the different possibilities given by the modal base. For (1c), the ordering source ranks higher those possibilities that lead naturally to a successful barbecue. Kratzer goes on to show how these parameters can be given a precise, formal role in the semantics of modals. Each alternative history is interpreted model-theoretically as a possible world, and modals are interpreted as quantifiers that range over the ordering-source-best worlds in the modal base.

Although Kratzer's work outlines how modals can and do vary in interpretation in context, it remains largely open how language users arrive at the particular contextual interpretations they do. Answering this question is an important precondition to the construction of computational systems that participate in ordinary modal talk, and to the understanding of the play of context and coherence in discourse generally. This paper addresses the contextual dependency of modal meaning from this perspective. Of particular concern is the phenomenon of MODAL SUBORDINATION [Roberts, 1986], present in (1) but better illustrated by the famous exemplar in (4).[2]

[2](4) is based on [Roberts, 1986] (11), where it is attributed to Fred Landman. I will use Roberts's term modal subordination to refer to the phenomenon exhibited, not its analysis; as we shall see, the

(4) A wolf might walk into the house. It would eat you.

In context, the meaning of the second sentence can be paraphrased as in (5).

(5) If a wolf walked into the house, it would eat you.

The force of the assertion is restricted to a hypothetical scenario where a wolf walks in. Because of this restriction, the hypothetical wolf can serve as the referent for the pronoun *it*.

Here is an intuitive account of how (4) gets its meaning. The first sentence introduces consideration of a hypothetical possibility into the discourse. The next sentence refers to this possibility, and makes an additional claim about it. In other words, the interpretation of (4) arises from an ANAPHORIC connection between the two modal sentences: they refer to a common semantic object, in virtue of a formal link.

The recognition that modals can refer to evoked possibilities goes back at least to [Isard, 1974]. The contribution of this paper is to show how to formalize this explanation robustly and naturally, so as to describe meanings of discourses like (1) using principles familiar from other domains and a simple semantic ontology. I begin in section 2 by describing the obstacles facing an anaphoric extension of Kratzer's approach to modality, and the consequent awkwardness of previous accounts of modal subordination. Then in section 3, I present a novel account in which these difficulties are avoided by parameterizing both modal and lexical meaning by discourse referents for sets of possible worlds, and by using an extended model of environments in dynamic semantics to keep track of associations between these possible worlds and ordinary individuals. I conclude in section 4 by suggesting that this approach provides a promising avenue for future research in allowing a high-level formal parallel between modal and temporal meaning and interpretation.

2 Obstacles to an Anaphoric Account of Modal Subordination

An anaphoric explanation of a contextual dependency postulates a semantic parameter that specifies an entity represented in an evolving model of discourse. By restricting possible interpretations to use a small set of pre-established values that the listener is currently attending to, such an explanation helps constrain the otherwise open-ended computational problem of disambiguation of utterances in context. (For a similar characterization of anaphora for [computational] semantics, see [Webber, 1988].)

Kratzer's approach to modal semantics does not lend itself immediately to an anaphoric account of modal subordination because values for its parameters for modal meaning, the MODAL BASE and the ORDERING SOURCE, are determined in a complex way. In Kratzer's approach, modal words and conditionals alike get their meaning by contributing features both to the modal base and to the ordering source.

examples might better be said to illustrate modal anaphora.

Accordingly, neither the modal base nor the ordering source specifies a semantic parameter that can be contributed by prior discourse or recovered in modal subordination.

Consider first the modal base. The modal base provides the set of possible worlds that modals quantify over. In sentences like (6), this set of worlds is given in part by the *if*-clause and in part by the alternative modal verbs.

(6) If a wolf comes in, I { could / might } escape.

The *if*-clause restricts the modal base to those worlds where a wolf comes in. The circumstantial meaning of *could* further restricts the modal base to those worlds which are like the real world in those properties relevant to determining my ability to escape a wolf. Alternatively, the epistemic meaning of *might* further restricts the modal base to those worlds which share with the real world just our information as to whether I would in fact escape a wolf.

In modal subordination, a contribution to the modal base is recovered only for a prior *if*-clause; in other respects, the modal base retains the restriction suitable to the overt modal verb in the subordinated sentence. For instance, the modals *might* and *could* in (7) retain the epistemic and circumstantial modal bases that they exhibit in (6):

(7) If a wolf came in, I could escape. You might be eaten, though.

If the two sentences in (7) are related anaphorically, then, it is not because they share a common modal base.

Likewise, consider the ordering source. The ordering source ranks the possibilities in the modal base by compatibility with some ideal, like what the real world is like, what the law provides, or what the speaker wants. This allows the theory to describe the contextual variability found for example in deontic modals, as illustrated by (8).

(8) { Legally / Morally / Ethically } Sandy should leave.

The different possible adverbs in (8) disambiguate a range of deontic ordering sources with which *should* is compatible: ideals established by legality, morality or ethics. In [Kratzer, 1991], the ordering source also figures in the semantics of counterfactuals.

(9) If a wolf had come in, I would have escaped.

For example, (9) can be analyzed with an ordering source of similarity to the real world—and a modal base consisting of all possible worlds where a wolf came in—to obtain the semantics for counterfactuals proposed by Lewis [Lewis, 1973].

Once again, in modal subordination, a contribution to the ordering source is recovered only for a prior *if*-clause; in other respects the ordering source reflects the overt modal verb in the subordinated sentence. Take (10).

(10) If a wolf had come in, you could not legally have killed it. But you still would have.

Here *would* has the ordinary counterfactual meaning associated with the *if*-clause in the prior sentence, even though that sentence also involves ranking worlds against what our law provides. (Incidentally, it transpires that Kratzer's approach is somewhat murky concerning the combination of counterfactuality and other modal words.)

As I will show in section 3, a slightly different parameterization of modal meaning allows parameters to be supplied by a straightforward anaphoric discipline, while preserving Kratzer's flexible contextualization of modal meaning. The new parameterization works by distinguishing the action of modal verbs and *if*-clauses in EVOKING HYPOTHETICAL SCENARIOS that elaborate on or present an alternative to available information, from the action of modal verbs in particular to EVALUATE SCENARIOS against a range of ideals using a range of constraints. In particular, this approach retains constrained parameters for modal base and ordering source to capture the semantic variation in modal verbs; however, at the same time, it incorporates new parameters for reference and described scenarios, represented model-theoretically as sets of possible worlds, to record the common anaphoric dynamics of modal verbs and *if*-clauses in modal subordination. In fact, the most significant divergence involved in the new parameterization lies outside modals proper. We must regard sentences in general as characterizing sets of possible worlds (rather than worlds themselves); and we must percolate the changes down to the meanings of lexical items. Even here, it is straightforward to relate the new meanings to the more familiar ones by translations or meaning postulates.

But what if we stick more closely to Kratzer's semantics? Two avenues have been explored in the literature. Neither is satisfactory. The first, epitomized by [Roberts, 1986, Roberts, 1989], is essentially syntactic. Roberts offers a formal treatment in which the contents of DRSs may be COPIED from the matrix of a possibility operator to the implicit *if*-clause that provides the restrictor of a subsequent modal, in a process that Roberts likens to accommodation [Lewis, 1979]. These copied elements are then free to play the compound semantic role Kratzer's theory gives them. Although it accounts for a variety of examples of modal subordination, this analysis is ultimately unsatisfactory (as recognized by a variety of recent researchers). It explains modal subordination not as an ordinary, anaphoric process but in terms of a distinct mechanism—accommodation—by which Lewis treats cases which actually violate the ordinary rules of language use.

The second approach, which includes [Kibble, 1994, Portner, 1994, Geurts, 1995, Kibble, 1995, Frank and Kamp, 1997], adopts an anaphoric analysis of modals using dynamic semantics—a formalism in which the meaning of a sentence is described as a relation between input environments recording the initial state of the discourse and corresponding possible output environments. These theories are alike in taking modals to refer to complex, dynamic objects—objects that must be manipulated using an overt quantifier and which therefore must set up referents by a special mechanism. This complexity leaves the approaches open to critique for awkward theoretical asymmetries and empirical failure.

The theoretical asymmetries are illustrated by (a) and (b) examples of (11) and (12).

(11) a A wolf walked in. A victim was waiting. It ate him.
 b A wolf might walk in. A victim would be waiting. It would eat him.

(12) a John ate a cheesesteak. It {was, #is} very greasy.
 b John might be eating a cheesesteak. It {would be, #is} very greasy.

Naively, the possibilities for anaphoric connections in (a) and (b) sentences seem quite similar. In (11a), we introduce a wolf and a victim as part of an ongoing description of an actual past situation. As we continue the description of that same situation, we can continue to refer to those entities. However, as (12a) underscores, when we move from the description of one situation to that of another situation—the actual present—where the entities we have introduced will not exist, ascription of physical properties to those entities is no longer possible. Likewise, in (11b), we introduce a wolf and a victim as part of an ongoing description of a hypothetical future situation. As we continue the description of that situation, we can refer to those entities, but, as (12b) reveals, we can no longer ascribe physical properties to them when we move to the description of another situation where those entities will not exist. The semantics of section 3 allows this naive explanation to be formalized for the case of modals, much as it already is in standard accounts of temporal reference such as [Partee, 1984, Hinrichs, 1986, Webber, 1988]. Indeed, as sketched in section 4, this parallel is an instance of a far-reaching similarity between modal and temporal interpretation.

On an explicitly quantificational dynamic theory of modality, there is no such symmetry. While its approach to (11a) and (11b) implements the naive account, the possibility of reference in (11b) requires a different kind of explanation. The impossibility of reference in (12b), meanwhile, essentially needs no explanation at all.

The technical basis for the asymmetry is revealed by considering a simple, counterfactual sentence, such as (13).

(13) If a wolf came in, it would eat you.

On the Lewis-style semantics, many possible worlds are under consideration in (13): all the possible worlds which have a wolf that comes in and which are "closest" to some initial possibilities. These worlds differ, among other factors, in the identity of the wolf that comes in. Given this ontology and a usual lexical semantics in which the meanings contributed by sentences are evaluated for truth and falsehood at single worlds, (13) must be translated by a quantificational expression of dynamic logic. Roughly, this expression will construct a set of TEMPORARY environments that list the current world and a wolf that comes in there; and it will make sure that in each, the listed wolf eats you. The temporary environments will play no role outside this expression, and hence the possible worlds, and the wolves, will be added only temporarily. Modal anaphora will be impossible.

Let me be more precise, by first introducing a presentation of modal meaning that is formulated to work with dynamic objects, environments and updates, in addition to worlds and propositions. (Similar presentations are fleshed out with varying details in [Kibble, 1994, Portner, 1994, Geurts, 1995, Kibble, 1995, Frank and Kamp, 1997].)

To characterize generic *if*- and *when*-clauses, dynamic semantics uses an operator $\Rightarrow$ relating two updates. $p \Rightarrow q$ is a new update; intuitively, this update makes sure that all extensions of the current environment by p have extensions by q in turn. That is, $p \Rightarrow q$ outputs its input i, when for any j that you can get from i by updating by p, there is an environment k you can get to from j by updating by q. Formally, this is represented as the following translation or abbreviation:

(14) $$(p \Rightarrow q) \equiv \lambda i \lambda i'.i = i' \wedge \forall j(pij \supset \exists k qjk).$$

As a test, any environment that $p \Rightarrow q$ outputs is identical to its input (in this case, the generalization is true); for some inputs, however, there is no output (in this case, the generalization is false). Since the environment does not change, no new referents are introduced into a model of discourse using this operation. While the discourse model may evolve as i is transformed into j and then into k, these models are discarded after the test is complete.

For other conditionals, we must lift this propositional connective to a quantificational connective incorporating reference to possible worlds. Following [Chierchia, 1992], these too involve semantic tests. They take the form $Q\ (P,Q)$ where both P and Q are DYNAMIC properties, functions that take an entity to return an update. Using Q^* to represent the static quantifier corresponding to Q, the interpretation of Q as a dynamic quantifier is given in (15).

(15) $$Q\ (P,Q) \equiv \lambda i \lambda i'.i = i' \wedge Q^*(\lambda x \exists j Pxij, \lambda x \exists jk.Pxij \wedge Qxjk)$$

According to this definition, the restrictor of Q^* is the set of entities for which a P-update is possible, and the nuclear scope is the set of entities for which a P-update can be followed by a Q-update. For a sentence like (16a), where the entities are ordinary individuals, the definition gives an interpretation that can be paraphrased as (16b).

(16) a Every man who owns a donkey beats it.
 b Every man who owns a donkey beats some donkey he owns.

Assuming a Lewis-style semantics for counterfactuals, the entities are possible worlds and the static quantifier Q^* says the consequent is true in all closest worlds where the antecedent is true. We can apply this quantifier to Chierchia's presentation to analyze (13) in dynamic semantics. The resulting expression relates *i* and *i* at *w* by this condition:

(17) In all closest worlds to *w* where *i* can be updated to find a wolf that comes in, *i* can be updated to find a wolf that comes in and eats you.

This condition, of course, is also a test; no overall update occurs with (17).

In the literature, examples like (12b) are typically taken to show that the prediction of impossibility of modal reference is generally right, and that the anaphora in modal subordination depends on an exceptional mechanism. The parallel between modality and tense motivated by (11) and (12) suggests the opposite. Nobody would argue that the impossibility of reference in (12a) was best explained by assuming that times had a quantificational structure that blocked off further reference, or that times had a dynamic structure that supplied referents by a separate mechanism in extraordinary circumstances. The simple and uniform analysis of (11) and (12) is that entities are incrementally added to a one overall model of the discourse, but that modal presuppositions about referents, like temporal ones, modulate possibilities for further anaphora.

Having made precise the theoretical asymmetry involved in previous dynamic accounts of modal subordination, I can now illustrate the key empirical difficulty. It centers around discourses such as the following:

(18) A wolf might walk in. We would be safe because John has a gun. He would use it to shoot it.

Informally, (18) describes two situations: an actual present situation, in which John has a gun; and a possible future development of that situation, in which a wolf walks in. The last sentence of (18) illustrates that the speaker may refer both to the possible wolf and to John's gun in a description of that possible future.

In the semantics to be presented in section 3, I will account for (18) as follows. The discourse describes two scenarios, interpreted semantically as sets of possible worlds. These sets are static objects. Each sentence in (18) provides constraints on these scenarios; for example, the first sentence says that one set of worlds—where a wolf walks in—represents an epistemically possible refinement of the other set of worlds—corresponding to reality (as we know it). Now, the sentences that evoke entities (real or hypothetical) introduce these entities into a flat model of the DISCOURSE. Thus, for example, the sentence that describes John's real gun adds to the discourse a record of how to find John's gun in all the epistemically possible worlds. This automatically enables subsequent discourse to refer to the gun in the wolf-scenario, because the scenario is epistemically possible.

In previous dynamic approaches, scenarios are interpreted as sets of DYNAMIC objects, in which possible worlds are paired with assignments that indicate what entities are available for reference there. (Entities are introduced into a sequence of evolving SCENARIOS rather than into evolving representations of the DISCOURSE.) Because scenario referents explicitly inventory available referents, we can only refer to a gun in a scenario in which a gun has been explicitly added. This is incompatible with the pattern of reference in (18). First, the discourse describes a possible elaboration of what we know, where a wolf comes in (and we are safe). Then the discourse evokes a further elaboration of our information which includes a gun. Although this elaboration describes reality, it nevertheless leaves the original hypothetical scenario unchanged. There is therefore no gun to refer to in the wolf-scenario, cf. [Roberts, 1995].

Again, to be precise about the difficulty, it will help to describe more precisely the particular entities different theories propose to store in the discourse model, and to introduce some notation. We avail ourselves of discourse markers p, q, etc. whose values are sets of (world, environment) pairs: dynamic scenarios. (The apparent recursion on environments here requires some type-theoretic delicacy; it will be another virtue of the account of section 3 that it dispenses with this recursion and therefore involves straightforward types.) If P is a dynamic property of possible worlds, the condition in (19) characterizes environments where p is interpreted by the set of (world, environment) pairs obtained by updating the (world, environment) pairs in r by the dynamic meaning of P.

(19) $p = r + P$

Formally, using $v(q,i)$ to name the value of discourse marker q in environment i, (19) characterizes i such that $v(p,i) = \{(w,k) \mid \exists j.(w,j) \in v(r,i) \wedge Pwjk\}$. The notation is from [Geurts, 1995].

Dynamic scenario variables such as p can provide the restrictors and scopes for dynamic modal quantifiers in such a way as to give those quantifiers closely related interpretations to Chierchia's dynamic quantifiers $Q\ (P,Q)$. To illustrate, suppose we have a dynamic scenario r such that $v(r,i) = \{(w,i) \mid \text{for all worlds } w\}$. If $p = r + P$ then $(w,j) \in v(p,i)$ just in case $Pwij$. (Both are equivalent to $\exists i'.(w,i') \in v(r,i) \wedge Pwi'j$.) If $q = p + Q$, moreover, $(w,k) \in v(q,i)$ just in case $\exists j.Pwij \wedge Qwjk$. Thus, consider an update that firsts introduces p such that $p = r + P$ and introduces q such that $q = p + Q$, obtaining an environment i, and then tests $Q^*(\lambda w \exists j.(w,j) \in v(p,i), \lambda w \exists k.(w,k) \in v(q,i))$. The final test is equivalent to testing $Q\ (P,Q)$ in i; so the only new feature is that now we have also introduced p and q into the discourse, and within p and q appropriate discourse markers are defined.

With this ontology and notation, the anaphoric interpretation of *might D; would D′* is formalized roughly as in (20):

(20) $q = p + D; some(p,q); s = q + D'; all(q,s)$

(Of course, (20) suppresses how *some* and *all* may be parameterized to reflect Kratzer's semantics for *might* and *would*.) These ordinary, dynamic variables are used for modal anaphora in [Kibble, 1994, Geurts, 1995, Frank and Kamp, 1997].

Now consider how (18) is treated on this account. Schematizing the successive clauses of (18) by W, S, G, U, we must represent the first three clauses using the chain of updates in (21):

(21) $w = r + W; s = w + S; g = r + G$

We must take g to extend r because of its factual mood. If we took g to extend s, then we would know only that John would have a gun if a wolf came in (and we were safe). But John does in fact have a gun. However, by representing the discourse this way, we have ruled out the possibility of referring to the gun and the wolf at the same time. U can refer to the wolf only if it extends w or s, for these are the only dynamic scenarios in which a wolf appears. And U can refer to the gun only if it extends g, because this is the only dynamic scenario in which the gun appears. Thus, the anaphora in (18) cannot be accounted for.

An alternative, explored by [Portner, 1994, Kibble, 1995], is to use variables that essentially refer to the updates themselves, rather than a particular result of the update. This introduces conditions of the form $M = \lambda p.p + P$. Particular results of the update are obtained as part of processing modal conditions, by applying update variables to a representation of the current state, which we might notate again by a distinguished dynamic scenario marker r. The use of updating dynamic variables leads to the following kind of formalization for *might D; would D′*:

(22) $M_1 = \lambda p.p + D; some(r, M_1(r)); M_2 = \lambda q.q + D'; all(M_1(r), M_2(M_1(r)))$

This approach represents the updates involved in (18) as in (23).

(23) $M_1 = \lambda p.p + W; M_2 = \lambda p.p + S; M_3 = \lambda p.p + G$

When the final sentence is reached, three concrete updates have been instantiated previously in the discourse: $M_1(r)$, $M_2(M_1(r))$ or $M_3(r)$. These are the three dynamic scenarios stored on the previous account, and we have seen that none is appropriate. Because the discourse has now stored all the updates, however, these updates could potentially be reapplied to different arguments in different orders in interpreting U. For example, in the state represented by $M_3(M_2(M_1(r)))$, there is a wolf and a gun and a victim. This application, however, is not faithful to the history of the discourse because M_3 is applied to the scenario where a wolf comes in, not to reality. Similarly, wolf, gun and victim are available in the state $M_2(M_1(M_3(r)))$. Here each update applies to the same set of possible worlds as it did in the original discourse. This strategy is unfaithful to the history of the discourse because the updates must be processed in a new order. No previous approach has worked out in a satisfying way how such novel sequences of updates should be defined or constrained.

The quantificational analysis of sentences like (13) is not a theoretical necessity; neither are the increasingly abstruse representations of possibilities that we have just explored. These problems follow only from the decision to expose the logical ontology of possible worlds as the linguistic ontology too. As I will now show, a much simpler, more symmetrical and more robust account can by constructed by departing from this decision, and eliminating overt dynamic quantification over possible worlds.

3 A New Anaphoric Theory

In this section, I begin in section 3.1 by proposing an alternative ontology for dynamic modal semantics, which encapsulates generalization over possible worlds into the meanings of words, and uses structured environments to record associations between discourse referents and possible worlds. The new ontology enables Kratzer's semantics of modal verbs and conditionals to be reparameterized as described in in section 3.2. The new ontology and parameterization, in turn, makes possible a straightforward formal account of modal subordination, which I present in section 3.3.

3.1 Putting worlds in context

In order to represent anaphoric connections explicitly, the semantics of utterances can describe not only their truth conditions but also their change in context. In particular, I adopt a variant of dynamic semantics, which I formalize as an axiomatic theory in classical type theory, following [Muskens, 1995, Muskens, 1996].

The approach begins by adopting the primitive types of [Muskens, 1995], which include not only the ordinary types **e** for individuals and **t** for truth values but also a type **w** of possible worlds, a type ε of eventualities, and a type τ of times. I assume that possible worlds share a common domain, but that a predicate u **in** w is true of precisely those u that actually exist at a world w. These resources provide the familiar type-theoretic tools of Montague grammar.

We also have a further type of contextual ENVIRONMENTS, which abstract the evolving model of discourse that the listener constructs; designate this type **s**. (The term environment departs from Muskens's overloaded use of the term state, already a subtype of eventuality.) A key innovation of my proposal lies in how environments are related to the entities such as ordinary individuals that have been evoked by the discourse. For Muskens and most previous research, an environment directly specifies various entities by storing each in a cell metaphorically conceived as analogous to a location in a computer memory or file system. In contrast, I assume that each of these cells in an environment merely provides A WAY TO IDENTIFY AN APPROPRIATE ENTITY DEPENDING ON A POSSIBLE WORLD—in other words, an individual concept. The use of these structured environments means that the context required to store or to access the element in a cell must include not only the environment but also the possible world at which the element is to be found. Accordingly, as shown below, introductions of discourse markers and tests on discourse markers must be augmented by possible-world parameters.

Despite this difference, the overall dynamics of environments are unchanged. As the discourse evolves, new information is stored in new cells, perhaps nondeterministically, and perhaps the elements in old cells are reassigned; a semantic parameter interpreted using a cell derives its value from the current state (and a set of possible worlds supplied by the construction in which it occurs). For example, an indefinite NP will set up a cell that specifies the individual it refers to (across some set of possible worlds); subsequent pronouns may then refer to this individual by accessing that cell (at any of those worlds). The meaning of a sentence is then a relation between possible input environments, representing where the listener is before the sentence is processed, and corresponding possible output environments, representing where the listener is once the sentence has been assimilated.

Contextual items, or DISCOURSE MARKERS, are the semantic objects that retrieve the elements from cells. (This supersedes Muskens's idiosyncratic use of the term store.) As in [Muskens, 1995], for uniformity of types, I treat environments as primitives. I therefore treat discourse markers as functions on environments and possible worlds; a discourse marker must have type $\langle \mathbf{s} \times \mathbf{w}, \alpha \rangle$ when the entities picked out by the corresponding cell at each world would have type α. In the appendix, I present an axiomatic theory, closely following Muskens, that relates discourse markers and environments appropriately.

In outline, the theory begins by introducing a family of predicates $\mathbf{mk}_\alpha$ true of the meaningful discourse markers of type α. The types α include at least individuals, type $\mathbf{e}$, and sets of possible worlds, type $\langle \mathbf{w}, \mathbf{t} \rangle$. It continues by including axioms guaranteeing that any environment can be updated so that any variable v takes on any combination f of values across possible worlds.

(24) $\forall ivf\,(\mathbf{mk}_\alpha(v) \supset \exists j.i[v]j \wedge \forall w(v(j,w) = fw))$

Here $i[v]j$ appeals to Muskens's notion of environments that differ only in the value of v.

The next step adapts Muskens's strategy of specifying the meanings of sentences in terms of formulas inspired by the DISCOURSE REPRESENTATION STRUCTURES, or DRSs, of Discourse Representation Theory [Kamp, 1981, Heim, 1982, Kamp and Reyle, 1993]. (Again, details are provided in the appendix.) I begin by describing DRSs built according to the following BNF syntax definition in (25). The primitive structures differ from standard ones in referencing a discourse marker ω which contextualizes the statement to an appropriate set of possible worlds.

(25) $C ::= R\{\omega, t_1, \ldots, t_n\} \mid C, C$
$DRS ::= [\ \omega : v_1, \ldots, v_k \mid C\ ] \ \mid\ DRS; DRS$

(I shall add additional expansions to these definitions in section 3.2 to handle modals.)

The interpretations of these conditions and DRSs is made precise by regarding them as abbreviations of expressions in the raw language of the theory of types. First, for a term t, we define an expansion t°—a function whose value at environment i and world w is obtained from t by replacing each occurrence of a free marker

v_n in t by $v_n(i,w)$, its value at that point. Then we can define the abbreviations in (26):

(26) $(R\{\omega, t_1, \ldots, t_n\})$ for $\lambda i.\forall w_0 w(\omega(i,w_0)w \supset Rwt_1^{\circ}(i,w)\ldots t_n^{\circ}(i,w))$
(C_1, C_2) for $\lambda i.C_1 i \wedge C_2 i$
$[\,\omega : v_1, \ldots, v_k \mid C\,]$ for $\lambda i\lambda j.i[\omega : v_1, \ldots, v_k]j \wedge Cj$
$D_1; D_2$ for $\lambda i\lambda j\exists k.D_1 ik \wedge D_2 kj$

Regarding these definitions as abbreviations (instead of translations) is convenient for mixing DRS notation and type theoretic expressions.

Let us examine the definitions in (26) in detail. First, consider C, which represents the syntactic class of CONDITIONS imposed in a DRS. Semantically, conditions are predicates of environments. In environment i, the atomic condition $R\{\omega, t_1, \ldots, t_n\}$ considers any possible world w that might be picked out by ω at i (and some other world w_0). At w, the condition ensures that the relation R holds between the values at i and w of t_1 through t_n. In this way, all atomic conditions directly encapsulate quantification over possible worlds. This is what allows this approach to avoid the explicitly DYNAMIC quantification over possible worlds which complicates previous accounts of modal subordination.

Each DRS, meanwhile, is interpreted as a relation between environments. An atomic DRS $[\,\omega : v_1, \ldots, v_k \mid C\,]$ permits the output to differ from the input in allowing the values of markers v_1 through v_k to vary across the worlds specified by ω, so long as the output satisfies the condition C. The difference relative to ω is represented by a relation on environments $i[\omega : v_1, \ldots, v_k]\,j$ which reduces (as in Muskens) to the composition of relations $[\,\omega : v_a\,]$ allowing difference in a single discourse marker. This relation is relativized to ω in different ways depending on the type of v_a.

For a scenario marker ω', the relation $i[\omega : \omega']j$ ensures that i and j agree on all discourse markers but ω'. Further, at any world w where $\omega(j,w)$ picks out some nonempty set of worlds, $\omega'(j,w)$ may pick out any set of worlds; at other worlds $\omega'(j,w)$ must be empty. The restriction ensures that ω' is a possibility RELATIVE to ω. Formally, this is equivalent to the conjunction of the two conditions in (27).

(27) a $i[\omega']\,j$
b $\forall w_0(\neg\exists w_1(\omega(i,w_0)w_1) \supset \neg\exists w_2(\omega'(j,w_0)w_2)))$

For other markers δ, $i[\omega : \delta]j$ imposes the condition that the value of δ must exist in each world in ω; this gives it the formalization in (28).

(28) a $i[\delta]\,j$
b $\forall w_0 w(\omega(j,w_0)w \supset \delta(j,w)\ \mathbf{in}\ w)$.

Again, by encapsulating quantification over possible worlds within atomic statements, these formulations of the updating condition avoid the cumbersome technique of explicit dynamic quantification over worlds.

As in Muskens, the condition C_1, C_2 holds of exactly the environments that satisfy both C_1 and C_2. Likewise, a complex DRS $D_1; D_2$ relates an input with any output that may be obtained by processing D_1 to obtain an intermediate environment and using that intermediate as the input to D_2.

In the remainder of this section, I will show how these structured environments allow modal meaning to be formulated in a way that also encapsulates the set structure of worlds within atomic statements. But first I will use the technical apparatus thus far introduced to briefly situate this project of encapsulation.

In some sense, DRS representations already involve a certain degree of encapsulation in how expressions are evaluated. Specifically, the progression of environments involved in interpreting linguistic expressions is perfectly regular, and is therefore left implicit in DRS representations. Encapsulating worlds—another parameter of evaluation—is a strategy with a similar motivation: hiding a perfectly regular feature of the interpretation of linguistic expressions. With worlds as opposed to environments, the encapsulation just happens to involve quantification within linguistic meanings rather than sequencing between them.

With encapsulation, the analysis aims for a parallel anaphoric analysis of modality and pronouns, but it does so only at the cost of breaking an existing close parallel—the parallel between sets of worlds and other linguistic pluralities, which grounded the discussion of possible worlds and dynamic quantifiers in section 2. There are two reactions to take. One is that the apparent set structure required to represent partiality in possible-worlds semantics was always an artifact of the formalization. Many other approaches to information, such as data semantics [Landman, 1986, Veltman, 1981] and situation semantics [Barwise and Perry, 1983], have formalized partial information using atomic points. These formalisms involve no parallel between states of partial information and pluralities, and give us no reason to expect any such parallel in language. In fact, we can hope for adaptations of the abbreviatory conventions for DRSs presented here that would instead interpret the same DRS formulas in models derived from these alternative theories of modality.

The second reaction is to look for representations of plurality that have more in common with the new proposal. Van den Berg's work [van den Berg, 1993, van den Berg, 1996] can be regarded as providing just such a theory. To represent dependencies among plural individuals, van den Berg models environments using sets of assignment-functions. Collective predication is accomplished by accessing all values assigned to discourse referents in the current state, while distributive predication is accomplishing by subdividing sets of ASSIGNMENTS (rather than sets of individuals) in order to maintain appropriate dependencies. In fact, the alternative representations of discourse and context required in van den Berg's formalism for plurals help to highlight what makes reference to sets of possible worlds distinctive—and moreover particularly simple. Quantification over worlds is always distributive—this allows the structuring of assignments not in terms of sets of (world,value) pairs but in terms of FUNCTIONS from worlds to values (a restricted, independently interesting semantic class). Moreover, quantification over worlds can be represented always with narrow scope (lexically encapsulated), in such a way that overt operators of distribution do not need to be added to the language. Further investigation of such semantic disparities may permit a more informed answer to the linguistic relation between partiality and plurality.

3.2 Parameterizing modal meanings

I now turn to the parameterization of modal meaning. As emphasized earlier, this semantics must separate the dynamics of modal verbs and *if*-clauses, to introduce sets of possible worlds into the discourse model, from the evaluation of sets of possible worlds, which modal verbs semantically contribute. The dynamics is specified by a new DRS $\mathbf{if}(\omega, \omega', D)$ which introduces the scenario marker ω'. The modal semantics is specified by new conditions $\mathbf{poss}(\omega, \omega', M, \leq)$ and $\mathbf{necc}(\omega, \omega', M, \leq)$ that relate the values of ω and ω' in the current environment. I first outline in broad terms how these function in the theory, then provide the exact definitions.

If D is a DRS then $\mathbf{if}(\omega, \omega', D)$ is also a DRS; call it D'. D' updates its input environment so that (at each world) ω' holds the set of worlds closest to ω where D holds. The resulting environment reflects not only the introduction of ω' but also the update D performed in obtaining ω'. (In treating D as a DRS rather than a property of scenarios, I assume that ω' appears free in D.) DRSs like D' are used both to interpret conditionals and to introduce the hypothetical scenario whose possibility or necessity modal verbs describe.

The descriptions of possibility and necessity relate a reference scenario ω_r and a new scenario ω_e where the content of the sentence holds. They are formulated using conditions $\mathbf{poss}(\omega_r, \omega_e, M, \leq)$ and $\mathbf{necc}(\omega_r, \omega_e, M, \leq)$ that are true of an environment i when ω_e is a possibility (or necessity, respectively) relative to ω_r according to the restrictions represented by the contextual parameter M (supplying a modal base) and the ideals represented by the contextual parameter $\leq$ (supplying an ordering source).

Here is how the dual effect of a modal verb such as *may* is represented as a complete DRS. Consider (29) as an example.

(29) John may leave.

Under this account, the interpretation of (29) depends on introducing the hypothetical scenario where John leaves and describing it as a possibility relative to appropriate parameters M and $\leq$. In particular, some reference scenario ω_r is obtained from the context, and the scenario ω_e just like ω_r but where John leaves is constructed and introduced into the discourse:

(30) $\mathbf{if}(\omega_r, \omega_e, \textit{John leaves in } \omega_e)$

The force of *may* is to assert that ω_e is possible from ω_r:

(31) $[\,|\,\mathbf{poss}(\omega_r, \omega_e, M, \leq)\,]$

The meaning of (29) is then represented as a DRS sequencing (30) and (31) (by the ;-operator).

(29) may be taken as a statement of what is permitted for John, on appropriate values of the M and $\leq$. In this case, $\leq$ must rank the possible worlds in ω_r and ω_e according to the extent to which John meets the obligations that are imposed on him there. Alternatively, (29) may be taken as a description of our ignorance of John's intentions. In that case, the possibility assessed by M and $\leq$ indicates only that the scenario ω_e is simply compatible with the information in ω_r.

For a simple sentence like *John will leave*, the corresponding two-step translation is:

(32) $\mathbf{if}(\omega_r, \omega_e, \textit{John leaves in } \omega_e); [\,|\,\mathbf{necc}(\omega_r, \omega_e, M, \leq)\,]$

This can be simplified considerably if, as I assume, the necessity condition here forces ω_r and ω_e to be identical. Then *will*, *would* and *is* simply predicate the content of the clause directly of their reference scenario. For (32) then, we assert simply that John leaves in ω_r.

To flesh out the semantics for the modal connectives $\mathbf{if}(\omega, \omega', D)$, $\mathbf{poss}(\omega, \omega', M, \leq)$ and $\mathbf{necc}(\omega, \omega', M, \leq)$, I follow the ideas of Lewis on conditionals [Lewis, 1973] and Kratzer on parameterized modals [Kratzer, 1991] as closely as possible in this new framework.

Lewis explains conditionals using a ternary relation **close** on worlds; $\mathbf{close}(w, w', w'')$ holds when a world w' is as similar or more similar to some reference world w than w'' is. Given some fixed starting point w, the relation given by $\lambda w' \lambda w''.\mathbf{close}(w, w', w'')$ is then an ordering on possible worlds that can act as a Kratzer-style ordering source in modal quantification. That is, to determine the truth of *if p then q* at a world w, find the set of closest worlds to w that satisfy p, and make sure that q is true in all. This set of closest worlds is the scenario evoked by the conditional. (Note that the need to evoke a representative set of closest worlds to interpret modal subordination places some restrictions on the kinds of orderings we may use to interpret **close**.)

I now adapt this idea to describe an update $\mathbf{if}(\omega_1, \omega_2, D)$. First, observe that, given the earlier concrete definitions of assignment and updates, the relation

(33) $\lambda i \lambda j. \exists k (i[\omega_1 : \omega_2] k \wedge Dkj)$

updates the context so that in each world w_0 where $\omega_1(j, w_0)$ evokes a possibility, $\omega_2(j, w_0)$ evokes a possibility for which a D-update has succeeded. Of course, (33) does not require that ω_2 contain all and only the closest worlds with this property. We can impose this requirement, however, by comparing $\omega_2(j, w_0)$ against all other sets $\omega_2(h, w_0)$ obtained similarly. We look at each world w in $\omega_1(j, w_0)$. If no world in any such $\omega_2(h, w_0)$ is strictly closer to w than any world in $\omega_2(j, w_0)$, $\omega_2(j, w_0)$

must contain only closest worlds. And if every world in all $\omega_2(h, w_0)$ is at least as far as some world in $\omega_2(j, w_0)$, then $\omega_2(j, w_0)$ must contain all closest worlds. So we can realize a Lewis-style condition by the following definition for $\mathbf{if}(\omega_1, \omega_2, D)$:

(34) $$\begin{aligned}\lambda i \lambda j.(\ & \exists k(i[\omega_1 : \omega_2]k \wedge Dkj) \wedge \forall h(\exists k(i[\omega_1 : \omega_2]k \wedge Dkh) \supset \\ & \forall w_0 w(\ \omega_1(i, w_0)w \supset \\ & \quad \forall w_h w_j(\omega_2(h, w_0)w_h \wedge \omega_2(j, w_0)w_j \supset \\ & \quad\quad \mathbf{close}(w, w_h, w_j) \supset \mathbf{close}(w, w_j, w_h)) \wedge \\ & \quad \forall w_h(\omega_2(h, w_0)w_h \supset \\ & \quad\quad \exists w_j(\omega_2(j, w_0)w_j \wedge \mathbf{close}(w, w_j, w_h))))))\end{aligned}$$

As promised, this meaning for **if** includes an update of the overall context by its argument *D*. This explains why subsequent discourse ought to be able to refer to entities introduced by *D*, and accounts for the general possibility of anaphora in modal subordination.

When it comes to the existence of referents across worlds, however, this definition involves some subtlety. This meaning for **if** does not necessarily ensure that markers assigned existent values throughout worlds identified by ω_1 are also assigned existent values throughout the worlds identified by ω_2. This is not really a surprise: think of the meaning of counterfactuals that explicitly eliminate individuals, like *if Monica Lewinsky had never existed*. However, for indicative conditionals and modal verbs generally, we can justify the preservation of existence from ω_1 into ω_2.

Here is how. First, I assume that ω_1, whether it represents our knowledge of reality or of some other hypothetical situation, includes a set of possible worlds that capture a full range of consistent possible futures. This assumption makes ω_1 rather like a information state in data semantics or a point in a branching model of future time [Landman, 1986, Veltman, 1981, Thomason, 1970]. Moreover, I suppose that the ranking **close** has this relationship to ω_1: Any two worlds in ω_1 are closer to each other than either is to any world not in ω_1. Then if the antecedent is consistent with ω_1, every possible world picked out by ω_2 in *j* will already be picked out by ω_1 in *j*. So the values for discourse markers $v(j, w)$ for worlds in ω_2 will be determined by the values for worlds in ω_1. In particular if $v(j, w)$ exists across the worlds in ω_1, it must also exist across the worlds in ω_2. So everything that exists at the ω_1 worlds will exist at the ω_2 worlds.

Now, this account does not apply for counterfactual conditionals (and modals describing impossibilities). In introducing counterfactual scenarios we must be content to regard the preservation of entities merely as a default with exceptions both in semantics and in formal pragmatics. How such a default should be captured must depend on the details of counterfactual reasoning, and is beyond the scope of this paper. Note, however, that in fleshing out this default, it will be necessary to avoid reaching nearby worlds where environments happen to leave all entities completely unspecified, merely because those worlds have yet to be described directly in discourse. The natural solution is to generalize the relationship **close** and broaden the kinds of updates related by $[\omega:\omega']$, in such a way as to allow closely corresponding entities to be added to the environment in these cases at the closely corresponding worlds introduced in the conditional. In particular, one route would involve revising (34) to feature a relation $\textbf{close}(w,i,w_j,j,w_k,k)$ comparing not only the worlds w, w_j and w_k but also the cases identified by i in w by j in w_j and by k in w_k.

The conditions for $\textbf{necc}(\omega_1,\omega_2,M,\leq)$ and $\textbf{poss}(\omega_1,\omega_2,M,\leq)$ are more straightforward. The contextual parameters of these conditions are given in two parts M and $\leq$. At a world w and environment i, $M\,w$ is a modal base—a set of worlds that indicates the a restricted perspective that the modal takes on $\omega_1(i,w)$ and $\omega_2(i,w)$. (This encodes the difference between epistemic and circumstantial modals, for example.) Then there is a ternary ordering source $\leq$, where $w' \leq_w w''$ indicates that w' is as close or closer than w'' to some ideal for w. (This encodes the difference between for instance modals that evaluate possibilities based on conformity with the law and those based on conformity with desires.)

Using these two parameters, necessity can be formalized as in (35).

(35) $$\lambda i.\forall w_0(\exists w\omega_1(i,w_0)w \supset \forall w(\omega_1(i,w_0)w \wedge M\,w_0 w \supset \exists w'.(\omega_1(i,w_0)w' \vee \omega_2(i,w_0)w') \wedge M\,w_0 w' \wedge w' \leq_{w_0} w \wedge \forall z(\omega_1(i,w_0)z \wedge M\,w_0 z \wedge z \leq_{w_0} w \supset \omega_2(i,w_0)z)))$$

This mirrors Definition 6 of [Kratzer, 1991]; apart from notation it contains only one slight difference. (35) imposes the same condition at each world w_0 at which a hypothetical scenario is evoked by $\omega_1(i,w_0)$. For any world w in this scenario (as restricted by the modal base $M\,w_0$), we must be able to find a better world w' in the modal base—further restricted either by ω_1 or ω_2—so that any world at least this close in $\omega_1(i,w_0)$ is an element of $\omega_2(i,w_0)$. Semantically, the sole difference with Kratzer lies in the disjunction of the ω_1 and ω_2 terms; Kratzer's Definition 6 has the equivalent of the ω_1 term only. This reflects the difference that ω_2 on the new approach could in principle have been obtained by a counterfactual update, whereas Kratzer's equivalent to ω_2 is a proposition p which is effectively intersected with ω_1. Since on a counterfactual update ω_2 may contain a closer world than any ω_1 world, the broader restriction is appropriate.

On this approach, possibility, like necessity, imposes a constraint on all worlds w_0 at which a hypothetical scenario is evoked by $\omega_1(i,w_0)$. This constraint is dual to that of necessity, so the overall semantics for possibility is as in (36).

(36) $$\lambda i.\forall w_0(\exists w\omega_1(i,w_0)w \supset \exists w(\omega_1(i,w_0)w \wedge M\,w \wedge \forall w'.((\omega_1(i,w_0)w' \vee \omega_2(i,w_0)) \wedge M\,w' \wedge w' \leq_{w_0} w \supset \exists z(\omega_1(i,w_0)z \wedge \omega_2(i,w_0)z \wedge M\,w_0 z \wedge z \leq_{w_0} w'))))$$

It requires $\omega_1(i,w_0)$ to contain a world w such that any available world at least this close to w_0 allows the discovery of a further close world which is an element of $\omega_1(i,w_0)$ and $\omega_2(i,w_0)$ and the modal base.

We have thus opened up Kratzer's modal semantics to depend on a parameter ω_1 that can be supplied referentially or by an independent constituent such an an *if*-clause. It is worth observing that this depends on the general encapsulation of quantification over possible worlds in my proposal. Kratzer argues that modals and conditionals must be treated simultaneously using examples such as (37).

(37) a If a murder occurs, the jurors must convene. (= [Kratzer, 1991] (24))
b [A murder occurs] $\supset$ must [the jurors convene] (= [Kratzer, 1991] option 1)

As Kratzer observes, when we evaluate the truth of (37b) at a single, total possible world, we cannot supply the point of evaluation for *must* by a separate constituent like an *if*-clause. For example, it is wrong to translate (37a) as (37b), for then the sentence comes out vacuously true when no murder happens to occur. Yet (37b) seems quite close to the representation I propose for (37a):

(38) $\mathbf{if}(\omega, \omega', \textit{a murder occurs}); \mathbf{if}(\omega', \omega'', \textit{the jurors convene});$
$[\ |\ \mathbf{necc}(\omega', \omega'', M, \leq)\ ]$

The difference in (38) is that the discourse markers ω, ω' and ω'' represent SETS of possible worlds and are related by conditionalization. They don't supply a point of evaluation for *must*; instead they supply the whole domain of quantification. So, for example, ω' can describe a set of worlds where a murder takes place even when we happen to live in a world where in fact no murder ever occurs.

3.3 Interpreting modal subordination

The dynamic semantics of section 3.1 and the modal semantics of section 3.2 permit an account of modal subordination in which ordinary anaphoric mechanisms underlie reference to sets of possible worlds and restrictions on anaphora across modal contexts are a function of the range of worlds across which discourse markers take existent values.

To illustrate his account, we return to example (18), repeated below:

(18) A wolf might walk in. We would be safe because John has a gun. He would use it to shoot it.

(39) gives a formalization of the example ignoring temporality and causality:

(39) $^1\mathbf{if}(\omega, \omega', [\ \omega' : u, e \mid \textit{wolf}\{\omega', u\},\ \textit{walk-in}\{\omega', u, e\}]);$
$^2[\ |\ \mathbf{poss}(\omega, \omega', M, \leq)\ ];\ ^3[\ |\ \textit{safe}\{\omega', \textit{we}\}\ ];$
$^4[\ \omega : g \mid \textit{gun}\{\omega, g\},\ \textit{have}\{\omega, j, g\}\ ];\ ^5[\ \omega' : e' \mid \textit{shoot}\{\omega', j, u, g, e'\}]$

(The superscripts are an expository device to label the occurrences of sub DRSs.)

Given the semantics outlined thus far, we can regard (39) as an abbreviation for an type-theoretic expression describing changing discourse environments in which sets of possible worlds are stored. With some assumptions about the modal parameters ω, **close**, M and $\leq$ required by this expression, its correctness as a representation of (18) can be justified. Moreover, we can see how each part of the resulting expression plays a role in an intuitive discipline setting up and propagating referents for sets of possible worlds.

Assume that the initial environment i assigns a value to ω corresponding to the information implicitly available to the participants in the conversation, which we can represent by a proposition or set of worlds E. Formally, then, for each world $w \in E$, $\omega(i,w) = E$, and for any other world $\omega(i,w)$ is empty. Assume also that the ranking **close** which determines the possibility evoked by a conditional in this context considers each world in E to be an equally likely possibility, and, of course, to be more likely than any world outside of E. Formally, this means that for each pair of worlds w and w' in E and any other world w'', **close**(w,w',w'') holds just in case $w'' = w'$ or $w'' \notin E$.

Thus, according to the definition in (34), updating by DRS 1 sets up a state j in which referents ω', u and e have two key properties. First, at each possible world w that is in any $\omega'(j,w_0)$ (for $w_0 \in E$), $u(j,w)$ picks out a wolf that exists in w and that enters in a future event picked out by $e(j,w)$. Moreover, all and only the closest worlds to w_0 (for $w_0 \in E$) where a wolf walks in are represented in $\omega'(j,w_0)$.

To get a better idea of what set exactly any $\omega'(j,w_0)$ is, consider what it means for DRS 2 to output j successfully. DRS 2 incorporates a Kratzer-style requirement of contextually-relativized possibility between ω and ω'. Since the possibility involved is epistemic, for all worlds in E we can take the parameter $M\,w$ to supply an everywhere true modal base, and we can take $\leq_w$ to supply the identity relation. For these parameters, **poss**$(\omega,\omega',M,\leq)$ is true of j just in case $\omega(j,w) \cap \omega'(j,w)$ is nonempty for all $w \in E$. In context, we know that $\omega(j,w) = \omega(i,w) = E$, so if j passes this test the closest worlds to w where a wolf comes in are worlds in E. By our assumption about **close**, $\omega'(j,w)$ therefore contains exactly the set of worlds in E where a wolf comes in. Call this set E'.

Overall, then, with these contextual parameters the composition of DRS 1 and 2 relates an input i to an output j if and only if the following five conditions are met: 1) E contains a nonempty subset E' giving the worlds in E where a wolf comes in; 2) $\omega(j,w) = E'$ for all $w \in E$; 3) $u(j,w)$ is a wolf for all $w \in E'$; 4) $e(j,w)$ is the event of the wolf walking in for all $w \in E'$; and 5) j is otherwise identical to i. This makes the right introduction and assessment of the possibility of a wolf walking in.

Next, DRS 3. This provides another test of j: that we are safe in all the worlds in E'. Given the context that has set up E', this means that if a wolf walked in, we would be safe. Thus, this formalism gives an appropriate semantics to modal subordination generally.

Finally, DRS 4 and 5 reveal the streamlined dynamics of anaphora made possible by evoking sets of possible worlds as static objects. DRS 4 relates j to an additional environment k which differs from k only in the value of g. Because it is interpreted with respect to ω, DRS 4 imposes a condition on g and k across any $w \in E$—assuming j has been updated as an update of i satisfying our initial assumptions by DRS 1–3. At w, $g(k,w)$ must be a gun that John has. If John does not have a gun at all the worlds in E, then no output k is forthcoming; the sentence is a false characterization of the information available to speaker and hearer.

DRS 5, in turn, relates k to a final environment l, which differs in the value of e', and where moreover the following requirement is satisfied. In any world $w \in \omega'(l,w_0)$, $e'(l,w)$ must be an event in which John uses $g(l,w)$ to shoot $u(l,w)$. What does this mean, assuming k is related to an initial i by DRS 1–4? Then w ranges over all the worlds in E', and therefore $u(l,w)$ picks out the wolf that is recorded in that element of E' as walking in. Meanwhile, since $E' \subseteq E$ and $g(i,w)$ picks out John's gun in all the worlds in E, here too $g(i,w)$ picks out John's gun. Thus the final sentence means, correctly: if a wolf walked in, John would use his gun to shoot the wolf.

Similar assumptions allow an explanation of why anaphora is not felicitous on the ordinary interpretation of an example such as (12b), repeated as (40) below.

(40) John might be eating a cheesesteak. #It is very greasy.

Under the same assumptions as with (39), *John might be eating a cheesesteak* is translated as an update relating an input i to an output j under certain conditions, which include finding a nonempty subset E' of an initial set of worlds E and setting up a discourse referent u that picks out a cheesesteak John is eating in all the worlds in E'. Now the second sentence must describe the full set of worlds E, because of its verb *is*. On the intended reading, where *it* is interpreted as u, (40) should require for its felicity that u exist in all the worlds in E. But this presupposition is not met, since it is E' that specifies the set of worlds where the value for u is known to exist, and E' need not contain all the worlds in E. (Indeed, the implicature is that it does not.)

4 Towards a Parallel between Modality and Tense

In this paper, I have presented a formalism in which a simple idea is worked out: that modal verbs refer to and describe sets of possible worlds. This formalism requires rich structures for the state of the discourse and for the semantics of modal words; however, it represents the contextual dependencies and the semantic contributions of modal words more compactly and more robustly than previous approaches. A modal assertion depends for its interpretation on the modal force, the modal base, the ordering source as in Kratzer's semantics; and it also depends on a reference set of possible world that is supplied by context and a described set of possible worlds that is introduced into the context. By modeling these dependencies in terms of static objects (sets of possible worlds), I obtain a streamlined type-theory and a streamlined account of the introduction and retrieval of entities into the discourse model. Crucially, the model correctly predicts the ability to refer to real entities in real possibilities regardless of the order in which the entities and possibilities are introduced into the discourse.

The formalism encapsulates all quantification over worlds so that translations at the level of a DRS refer to discourse markers for sets of possible worlds as primitives. Because of this encapsulation, the temporal and modal relativization of DRS conditions can be presented symmetrically on this theory. For example, the account offers formally parallel explanations for the possibility of continued anaphora within temporal and modal descriptions such as (11).

(11) a A wolf walked in. A victim was waiting. It ate him.
b A wolf might walk in. A victim would be waiting. It would eat him.

The account offers parallel explanations for the impossibility in many cases of anaphora across multiple temporal and modal descriptions, as in (12).

(12) a John ate a cheesesteak. It {was, #is} very greasy.
b John might be eating a cheesesteak. It {would be, #is} very greasy.

Additional parallels between modal and temporal interpretation seem to be forthcoming on this approach. I mention three briefly.

The first concerns the pattern of referentiality in modality and tense. Since Reichenbach, temporal reference has been described in terms of three points: a speech point *S*, a reference time *R* and the event time *E* [Reichenbach, 1947]. I have already described parallels to *R* and *E* in the reference scenario and the described scenario in modal meaning. It is plausible, moreover, to regard the the difference between *will* and *would* (both describing future events) as a difference in whether the reference scenario represents a possibility compatible with our present information; with *will*, but not with *would*, the reference scenario really must be possible [Iatridou, 1997]. This suggests a role for a speech point in modal semantics as well, so that modality parallels tense in its triple referential context-dependency.

The second parallel concerns the modal and temporal interpretation of noun phrases. Enç has argued that the temporal interpretation of noun phrases is referential [Enç, 1986]. Beghelli [Beghelli, 1996] explores the parallel idea that the modal scope of noun phrases is also interpreted referentially. This gives a novel characterization of the well-known ambiguity between the *de re* and *de dicto* interpretations of embedded noun phrases, on which the content of the embedded noun phrase can describe real properties of a real entity, or subjective properties (perhaps of entities that exist only subjectively). On a referential account, this difference is explained not by the scope of the noun phrase, but rather by a set of possible worlds introduced by the context that is recovered anaphorically by the noun phrase itself.

The third parallel, the possibility for sloppy anaphora, is in fact a property shared by all anaphora. Anaphora is sloppy when the entity evoked by an anaphor differs from the referent of its antecedent because an anaphor contained in the antecedent is interpreted differently; both temporal and modal referents figure in sloppy anaphora:

(41) a **[VP [Tense]]** You Past *[think I [Past] was crazy]*. You probably still Pres do VPE.

b **[Modal [NP]]** John would use slides *if [[he] had to give the presentation]*. Bill would just use the chalkboard.

For example, in (41a), the anaphor is the VPE; its antecedent is the main VP of the previous sentence. The sloppy reading of the elided VP allows the reference point of the embedded verb *was crazy* to be reinterpreted as referring to the moment of speech. Dually, in (41b), the anaphor is the modal verb; on the sloppy reading, the pronoun *he* in the description of its scenario referent refers first to John, then to Bill. By treating sloppy identity as a general feature of the resolution of anaphoric devices, including tense and modals, [Stone and Hardt, 1997] give a simple, unified treatment of the sloppy identity patterns observed with tense and modals. ((41) repeats example (4) from that paper.)

The present paper has introduced a simple semantics that can can allow such parallels to be formalized in terms of ordinary static objects—sets of possible worlds. Scenario markers for these sets are primitive. They are introduced directly into the evolving discourse model by an update $\mathbf{if}(\omega, \omega', D)$ introduces ω' as what would happen in ω if D. They are evoked like other entities; they are described by linguistically-motivated, atomic conditions: $\mathbf{poss}(\omega, \omega', M, \leq)$ and $\mathbf{necc}(\omega, \omega', M, \leq)$. This abstract view crystallizes a formalism for reference to possible worlds in a concise form. It is hoped that the availability of these representations for modal discourse will allow future research to better address open problems in modal interpretation.

References

[Barwise and Perry, 1983] Barwise, J. and Perry, J. (1983). *Situations and Attitudes*. MIT Press, Cambridge.

[Beghelli, 1996] Beghelli, F. (1996). Mood and the interpretation of indefinites. Manuscript, University of Pennsylvania.

[Chierchia, 1992] Chierchia, G. (1992). Anaphora and dynamic binding. *Linguistics and Philosophy*, 15(2):111–183.

[Enç, 1986] Enç, M. (1986). Towards a referential analysis of temporal expressions. *Linguistics and Philosophy*, 9(3):405–426.

[Frank and Kamp, 1997] Frank, A. and Kamp, H. (1997). On context dependence in modal constructions. In *SALT-97*, Stanford, CA. Stanford.

[Geurts, 1995] Geurts, B. (1995). *Presupposing*. PhD thesis, University of Stuttgart.

[Heim, 1982] Heim, I. (1982). *The Semantics of Definite and Indefinite Noun Phrases*. PhD thesis, University of Massachusetts, Amherst. Published 1987 by Garland Press.

[Hinrichs, 1986] Hinrichs, E. (1986). Temporal anaphora in discourses of English. *Linguistics and Philosophy*, 9(1):63–82.

[Iatridou, 1997] Iatridou, S. (1997). The grammatical ingredients of counterfactuality. Manuscript, MIT.

[Isard, 1974] Isard, S. (1974). What would you have done if... *Theoretical Linguistics*, 1:233–255.

[Kamp, 1981] Kamp, H. (1981). A theory of truth and semantic representation. In Groenendijk, J., Janssen, T., and Stokhof, M., editors, *Truth, Interpretation and Information*. Foris, Dordrecht.

[Kamp and Reyle, 1993] Kamp, H. and Reyle, U. (1993). *From Discourse to Logic: Introduction to Modeltheoretic Semantics of Natural Language, Formal Logic and Discourse Representation Theory*. Kluwer, Boston.

[Kibble, 1994] Kibble, R. (1994). Dynamics of epistemic modality and anaphora. In Bunt, H., Muskens, R., and Rentier, G., editors, *International Workshop on Computational Semantics*, Tilburg, NL. ITK.

[Kibble, 1995] Kibble, R. (1995). Modal subordination, focus and complement anaphora. In *Proceedings of the Tbilisi Symposium on Language, Logic and Computation*.

[Kratzer, 1977] Kratzer, A. (1977). What 'must' and 'can' must and can mean. *Linguistics and Philosophy*, 1:337–355.

[Kratzer, 1981] Kratzer, A. (1981). The notional category of modality. In Eikmeyer, H. J. and Rieser, H., editors, *Words, Worlds and Contexts*, pages 38–74. de Gruyter, Berlin.

[Kratzer, 1991] Kratzer, A. (1991). Modality. In von Stechow, A. and Wunderlich, D., editors, *Semantics: An International Handbook of Contemporary Research*, pages 639–650. de Gruyter, Berlin.

[Landman, 1986] Landman, F. (1986). *Data Semantics: An Epistemic Theory of Partial Objects*, pages 1–96. Foris, Dordrecht.

[Lewis, 1979] Lewis, D. (1979). Scorekeeping in a language game. In *Semantics from Different Points of View*, pages 172–187. Springer Verlag, Berlin.

[Lewis, 1973] Lewis, D. K. (1973). *Counterfactuals*. Harvard University Press, Cambridge, Massachusetts.

[Muskens, 1995] Muskens, R. (1995). Tense and the logic of change. In Egli, U., Pause, P. E., Schwarze, C., von Stechow, A., and Wienold, G., editors, *Lexical Knowledge in the Organization of Language*, volume 114 of *Current Issues in Linguistic Theory*, pages 147–184. John Benjamins, Philadelphia.

[Muskens, 1996] Muskens, R. (1996). Combining Montague semantics and discourse representation. *Linguistics and Philosophy*, 19(2):143–186.

[Partee, 1984] Partee, B. H. (1984). Nominal and temporal anaphora. *Linguistics and Philosophy*, 7(3):243–286.

[Portner, 1994] Portner, P. (1994). Modal discourse referents and the semantics-pragmatics boundary. Manuscript.

[Reichenbach, 1947] Reichenbach, H. (1947). *Elements of Symbolic Logic*. Macmillan, New York.

[Roberts, 1986] Roberts, C. (1986). *Modal Subordination, Anaphora and Distributivity*. PhD thesis, University of Massachusetts, Amherst.

[Roberts, 1989] Roberts, C. (1989). Modal subordination and pronominal anaphora in discourse. *Linguistics and Philosophy*, 12(6):683–721.

[Roberts, 1995] Roberts, C. (1995). Domain restriction in dynamic interpretation. In Emmon Bach, Eloise Jelinek, A. K. and Partee, B., editors, *Quantification in Natural Languages*. Kluwer, Dordrecht.

[Stone and Hardt, 1997] Stone, M. and Hardt, D. (1997). Dynamic discourse referents for tense and modals. In *IWCS 2*, pages 287–299, Tilburg, NL.

[Thomason, 1970] Thomason, R. H. (1970). Indeterminist time and truth-value gaps. *Theoria*, 36:265–281.

[van den Berg, 1993] van den Berg, M. H. (1993). Full dynamic plural logic. In Bimbó, K. and Máté, A., editors, *Proceedings of the Fourth Symposium on Logic and Language*, Budapest.

[van den Berg, 1996] van den Berg, M. H. (1996). Generalized dynamic quantifiers. In van der Does, J. and van Eijk, J., editors, *Quantifiers, Logic and Language*. CSLI.

[Veltman, 1981] Veltman, F. (1981). Data semantics. In Groenendijk, J., Janssen, T., and Stokhof, M., editors, *Truth, Interpretation and Information*. Foris, Dordrecht.

[Webber, 1988] Webber, B. L. (1988). Tense as discourse anaphor. *Computational Linguistics*, 14(2):61–73.

5 Appendix: Complete Definitions

The complete formalism used in this paper can be specified, following [Muskens, 1995, Muskens, 1996], as an axiomatic theory in type theory, and a specification of DRS expressions as abbreviations of type-theoretic expressions.

5.1 Type theory

We start with a type of environments **s** and a set Θ of types that discourse markers can take on, namely $\{\mathbf{e}, \tau, \varepsilon, \langle \mathbf{w}, \mathbf{t} \rangle\}$. We introduce for each $\alpha \in \Theta$ a predicate $\mathbf{mk}_\alpha$ identifies the elements of type $\langle \mathbf{s} \times \mathbf{w}, \alpha \rangle$ that name discourse markers of type α. (42) makes use of this predicate to describe environments that differ only in v:

(42) If v is a term of type $\langle \mathbf{s} \times \mathbf{w}, \alpha \rangle$ for any $\alpha \in \Theta$, then $i[v]j$ abbreviates the conjunction of

- $\forall u((\mathbf{mk}_\alpha u \wedge u \neq v) \supset \forall w(u(i,w) = u(j,w)))$
- and the conjunction of all $\forall u(\mathbf{mk}_\beta u \supset \forall w(u(i,w) = u(j,w)))$ for all $\beta \in \Theta - \{\alpha\}$

This adapts [Muskens, 1995, def 1]. We also add abbreviations that relativize updates to sets of possible worlds.

(43) a If ω and ω' are terms of type $\langle \mathbf{s} \times \mathbf{w}, \langle \mathbf{w}, \mathbf{t} \rangle \rangle$, then $i[\omega : \omega']j$ abbreviates the conjunction of $i[\omega']j$ and
$\forall w_0(\neg \exists w_1(\omega(i, w_0)w_1) \supset \neg \exists w_2(\omega'(j, w_0)w_2)))$

b If ω is a term of type $\langle \mathbf{s} \times \mathbf{w}, \langle \mathbf{w}, \mathbf{t} \rangle \rangle$, and δ is a term of type $\langle \mathbf{s} \times \mathbf{w}, \alpha \rangle$, for any $\alpha \in \Theta - \{\langle \mathbf{w}, \mathbf{t} \rangle\}$, then $i[\omega : \delta]j$ abbreviates the conjunction of $i[\delta]j$ and $\forall w_0 w(\omega(j, w_0)w \supset \delta(j, w)$ **in** $w)$.

These compose together:

(44) If ω is a term of type $\langle \mathbf{s} \times \mathbf{w}, \langle \mathbf{w}, \mathbf{t} \rangle \rangle$, δ is a term of type $\langle \mathbf{s} \times \mathbf{w}, \alpha \rangle$ for any $\alpha \in \Theta$ and σ is a nonempty sequence of terms of type $\langle \mathbf{s} \times \mathbf{w}, \alpha \rangle$ for any $\alpha \in \Theta$, then $i[\omega : \delta, \sigma]j$ abbreviates $\exists k. i[\omega : \delta]k \wedge k[\omega : \sigma]j$.

Three axioms then describe the possibilities for updates.

(45) a $\forall ivf(\mathbf{mk}_\alpha(v) \supset \exists j. i[v]j \wedge \forall w(v(j, w) = fw))$

b $\mathbf{mk}_\alpha v$ for each store name v of type $\langle \mathbf{s} \times \mathbf{w}, \alpha \rangle$, for each $\alpha \in \Theta$

c $u \neq v$ for any two different store names u and v of any type $\langle \mathbf{s} \times \mathbf{w}, \alpha \rangle$

These adapt [Muskens, 1995, AX1,2,3].

5.2 DRS Abbreviations

First we define an expansion t° which is always term of some type $\langle \mathbf{s} \times \mathbf{w}, \alpha \rangle$.

(46) If t is a constant symbol c, $t^\circ = \lambda p.c$. If t is a discourse marker name v, $t^\circ = \lambda p.v(p)$. If t is an application $f(a)$, $t^\circ = \lambda p.f^\circ(p)(a^\circ(p))$. If t is a tuple (a,b), $t^\circ = \lambda p.(a^\circ(p), b^\circ(p))$.

Then we have the following abbreviations:

(47) a

$R\{\omega, t_1, \ldots, t_n\}$ for $\lambda i.\forall w_0 w(\omega(i,w_0)w \supset Rwt_1^\circ(i,w)\ldots t_n^\circ(i,w))$

$\mathbf{necc}(\omega_1, \omega_2, M, \leq)$ for $\lambda i.\forall w_0(\exists w \omega_1(i,w_0)w \supset \forall w(\omega_1(i,w_0)w \wedge M\, w_0 w \supset \exists w'.(\omega_1(i,w_0)w' \vee \omega_2(i,w_0)w') \wedge M\, w_0 w' \wedge w' \leq_{w_0} w \wedge \forall z(\omega_1(i,w_0)z \wedge M\, w_0 z \wedge z \leq_{w_0} w \supset \omega_2(i,w_0)z)))$

$\mathbf{poss}(\omega_1, \omega_2, M, \leq)$ for $\lambda i.\forall w_0(\exists w \omega_1(i,w_0)w \supset \exists w(\omega_1(i,w_0)w \wedge M\, w \wedge \forall w'.((\omega_1(i,w_0)w' \vee \omega_2(i,w_0)) \wedge M\, w' \wedge w' \leq_{w_0} w \supset \exists z(\omega_1(i,w_0)z \wedge \omega_2(i,w_0)z \wedge M\, w_0 z \wedge z \leq_{w_0} w'))))$

C_1, C_2 for $\lambda i.C_1 i \wedge C_2 i$

$[\,\omega : v_1, \ldots, v_k \mid C\,]$ for $\lambda i\lambda j.i[\omega : v_1, \ldots, v_k]j \wedge Cj$

$\mathbf{if}(\omega, \omega', D)$ for $\lambda i\lambda j.(\exists k(i[\omega_1 : \omega_2]k \wedge Dkj) \wedge \forall h(\exists k(i[\omega_1 : \omega_2]k \wedge Dkh) \supset \forall w_0 w(\omega_1(i,w_0)w \supset \forall w_h w_j(\omega_2(h,w_0)w_h \wedge \omega_2(j,w_0)w_j \supset \mathbf{close}(w, w_h, w_j) \supset \mathbf{close}(w, w_j, w_h)) \wedge \forall w_h(\omega_2(h,w_0)w_h \supset \exists w_j(\omega_2(j,w_0)w_j \wedge \mathbf{close}(w, w_j, w_h))))))$

$D_1; D_2$ for $\lambda i\lambda j \exists k.D_1 ik \wedge D_2 kj$