Knowledge, Causality and Temporal Representation

Alex Lascarides*
Centre for Cognitive Science,
2 Buccleuch Place,
Edinburgh EH8 9LW
Tel: +44 31 668 4432
e-mail: `alex@uk.ac.ed.cogsci`

October 12, 1993

Abstract

In this paper, a formal semantic framework is developed in order to account for the temporal semantics of text. The theory is able to represent and reason about both semantic issues, which are independent of world knowledge (WK), and pragmatic issues, which are not, within a single logical framework. The theory will allow a text's semantic entailments to differ from its pragmatic ones, even though they are all derived within the same logic. I demonstrate that this feature of the theory gives rise to solutions to several puzzles concerning the temporal structure of text.

1 The Problem

The purpose of this paper is to provide a formal account of the temporal semantics of text. The chief goal is to explain when a text is *temporally coherent*: it should not mislead the reader as to the order of the events reported. If John hits Max, causing Max to turn round (to face John), then text (1) reflects this while (2) distorts it:

(1) John hit Max. Max turned round.

(2) Max turned round. John hit him.

So the order in which such 'atomic' clauses[1] appear is crucial. It gives rise to an Ordering Question: given a particular order in which events are described, what are the constraints in interpretation on their temporal and causal relations in the world?

---

*Many thanks are due to Jon Oberlander, Marc Moens and other members of the Temporal Mechanics workshop in HCRC for their useful comments on this work. All the mistakes are my own.

[1]For the purposes of this discussion, an atomic clause is one that consists of a simple tense as opposed to the progressive or perfect, and contains no temporal adverbials.

Texts (1) and (2) suggest that the ordering constraints are satisfied if clause order reflects temporal order. But this constraint, although sufficient, is not necessary. Text (3a,b) is temporally coherent, even though the event described by (3b) does *not* follow the event described by (3a).

(3) a. The council built the bridge.
b. The architect drew up the plans.

The situation is further complicated by the fact that, contrary to what (2) indicates, we *can* describe events in the reverse order to their *causal* sequences while still using atomic clauses in certain cases. Texts (4) and (5) are examples of this.

(4) Max fell. John pushed him.

(5) Max died. John poisoned him.

(3), (4) and (5) show that the coherent clause order and temporal order need not match; whereas the mismatch in (2) leads to incoherence. These examples attempt to capture the same temporal relations between events (i.e. the second event described occurs before the first), but only (3), (4) and (5) are coherent; thus, something more than temporal relations between events must be involved in the ordering constraints on text.

The above data is puzzling, for it is commonly held that if (2), (4) and (5) have similar logical forms, then they should describe similar temporal relations between events. The puzzle is that this is not so. Our aim is to explain this: why can't the event of hitting and turning round be described in the reverse to the causal order using atomic clauses, but falling and pushing and dying and poisoning can?

The above texts show that the current theories are insufficient. The relevant theories expressed in Discourse Representation Theory (DRT) (Kamp and Rohrer (1983), Partee (1984) Hinrichs (1986)) and Dowty's (1986) interval-based analysis are all unable to account for the fact that in certain cases the linguistic order of events mismatches their temporal order, even though no explicit synactic markers are evident in the clauses to indicate this. Webber (1987) allows interpretations where descriptive order mismatches temporal order, but her theory is unable to predict that mismatching descriptive order and temporal order is allowed in some cases (e.g. (3), (4) and (5)) but not in others (e.g. (2)). Clearly, the current theories must be modified, as they fail to account for a significant portion of the linguistic data.

To exacerbate the situation, we may add other marks in syntax so that the event sequence of hitting and turning round is described in reverse to causal order, thus 'cancelling' the constraint on (2). Examples (6) and (7) demonstrate this.

(6) Max turned round. John *had* hit him.

(7) Max turned round *just after* John hit him.

The puzzle is: how can the constraints on the interpretation of temporal order be affected by the semantics of the perfect and temporal connectives?

We might also choose to describe some things in terms of states, rather than events; compare text (8a,b,c) (which is acceptable) with the alternative (8a,b,d) (which is awkward), and (8a,b,e) (which is much better):

(8) a. Max crossed the finish line in first place.
b. He punched the air with joy.
c. He arrived home with the cup.
d. He was home with the cup.
e. A moment later, he was home with the cup.

How can the different semantics of events and states lead to an explanation of our distinct interpretations of (8a,b,d), (8a,b,d) and (8,a,b,e)?

To adequately characterise temporal coherence and answer the Ordering Question, we would have to determine constraints between the linguistic order of events and their temporal order, and on the use of aspect and other markers of non-atomic syntax.

One might argue that all the above texts that I have claimed are awkward are simply *pragmatically* misleading as opposed to semantically contradictory. This may be so; but it is undesirable to simply dismiss the phenomena to pragmatics and say nothing more on the subject. I aim in this paper to provide a unified logical account of both the semantics of the above texts (i.e. their truth conditions), *and* their pragmatic implications (which for our purposes will be what one believes about the texts given certain world knowledge). In so doing, I hope to demonstrate why some of the above texts are misleading, and thus provide an answer to the Ordering Question.

2 The Need for Default Reasoning

I will first concentrate on the problem of representing texts (1), (2), (4) and (5); the theory will subsequently be extended to deal with the rest of the above examples.

As we have already mentioned, (2), (4) and (5) ought to have similar logical forms, and yet they describe different temporal relations between the events; our task is to explain why this is so.

(2) Max turned round. John hit him.

(4) Max fell. John pushed him.

(5) Max died. John poisoned him.

Consider the above on an intuitive level. How can we account for their different interpretations? One possibility is that the difference is derived from the difference in the relations that typically hold between the events being described. More specifically, if there is a causal relation (but we don't know in which direction) between a pushing and a falling, or between a poisoning and a dying, then it is much more likely that the pushing (poisoning) is the cause than the falling (dying). The same generality does not hold for hitting and turning round; if there is a causal relation (but we don't know in which direction) between hitting and turning round, then we are unable to conclude, even tentatively, the direction of the causal relation. It is this difference between (2), (4) and (5) that may yield an account of their different interpretations.

Let us investigate this further. One might assume the following rule R holds, which in turn could be colloquially characterised using a belief operator B as in (9), where the operator B has a Hintikka-style semantics.[2]

R I believe that in the absence of contrary information, John pushed Max and Max fell entails the pushing caused the falling.

(9) $B(((\text{Max fell}) \wedge (\text{John pushed Max}) \wedge \neg B \neg(\text{John's pushing Max caused Max's fall})) \Rightarrow (\text{John's pushing Max caused Max's fall}))$

Note that if I believed that the pushing and falling events were too far apart, either temporally or spatially, for there to be any causal relation between them at all, then I would believe in particular that the pushing did not cause the falling, and so I would not believe the antecedent in (9), and therefore wouldn't necessarily believe the consequent. Hence the case where the two events are too far apart for there to be any causal relation between them is not a counterexample to the truth of (9).

There is a similar rule R′ (characterised in turn as (10)) for the poison/dying case:

R′ I believe that in the absence of contrary information, John poisoned Max and Max died entails the poisoning caused the dying.

(10) $B(((\text{Max died}) \wedge (\text{John poisoned Max}) \wedge \neg B \neg(\text{John's poisoning Max caused Max's death})) \Rightarrow (\text{John's poisoning Max caused Max's death}))$

However, a similar rule would not hold for John hitting Max and Max turning round. If we believe that John hit Max and that Max turned round and that these events occurred in close spatio-temporal proximity (but we don't know in which order) then we cannot derive any conclusions, even tentative ones, as to their order.

The above causal laws may yield a strategy for explaining the differences between (2) and (4) and (5). If we were to describe the semantics of text in a framework where rules (9) and (10) could be represented and used in deduction, then perhaps we could use them to deduce via the logic of belief that from believing (4) and (5) the reader believes the situations are ones where the second events described cause the first. Informally, the deduction would go as follows: the semantics of text (4) (text (5)) entails that a falling and pushing (dying and poisoning) occurred. Rule (9) (rule (10)) entails that if one believes a falling and pushing (dying and poisoning) occured, as one does if one believes text (4) (text (5)), then contrary to other information, one believes the pushing (poisoning) caused the falling (dying), and so one believes the second event described causes the first and therefore precedes it. On the

[2] As B would have a Hintikka-style semantics, (9) entails
$(B(\text{Max fell}) \wedge B(\text{John pushed Max}) \wedge \neg B \neg(\text{John's pushing Max caused Max's fall})) \Rightarrow B(\text{John's pushing Max caused Max's fall})$.

other hand, from believing (2) the reader does *not* believe that the second event described causes the first, as there is no appropriate rule similar to (9) or (10) that could be used as part of the deduction.

Rules (9) and (10) are classified in the AI literature as *default rules.* Levesque (1985) characterises the standard default reasoning example in a similar way to the above: if one believes that Tweety is a bird, then one believes that Tweety flies, unless one already believes something to the contrary (e.g. Tweety is a penguin and so doesn't fly). So the following formula is true.

(11) $B((bird(tweety) \wedge \neg B \neg fly(tweety)) \Rightarrow fly(tweety))$

(9), (10) and (11) are called default rules because they are about making tentative conclusions from an incomplete source of knowledge. For example, I may not know all there is to know about John's pushing Max, or about the kind of bird Tweety is, but I can nevertheless jump to certain conclusions, e.g. that the pushing caused the falling or that Tweety flies.

So, in claiming that rules (9) and (10) play a crucial role, we are essentially claiming that the difference in coherence between (2) on the one hand and (4) and (5) on the other derives from the fact that we have in our world knowledge (WK), independently of our linguistic knowledge, *default rules* concerning the causal relations between falling and pushing and between poisoning and dying, but we have no such rule for hitting and turning round.

The above default rules (9) and (10) potentially explain why the descriptive order of events mismatch temporal order in (4) and (5). But we have yet to show why the interpretation of (2) is one where descriptive order *matches* temporal order. I claim that in addition to the rules (9) and (10), which constitute aspects of the reader's defeasible *world* knowledge, the reader has defeasible *linguistic* knowledge; in particular, the default rule which is expressed at an intuitive level as follows:

- *Narration*: If a text describes an eventuality e_1 before it describes e_2, then unless the reader has reason to believe something to the contrary (e.g. from WK or from linguistic features like temporal adverbials), he believes that there are times t_1 and t_2 such that e_1 holds at t_1, e_2 holds at t_2, and $t_1 \leq t_2$.

This default rule states that in the absence of contrary information, the reader assumes that the text is purely narrative, where linguistic order reflects temporal order. In other words, we assume (2) is narrative because it is *possible* for the events described to occur in narrative order; our WK tells us nothing to the contrary.

(2) Max turned round. John hit him.

Text (5), on the other hand, is an example of where *Narration*'s antecedent is *not* satisfied due to contrary *indefeasible* information (you can't poison someone who has died).

(5) Max died. John poisoned him.

It is because the narrative structure of text is cancellable by WK, as (5) demonstrates, that the above rule is represented as a default rule, and not as a semantic principle independent of WK.

In the case of text (4), we now have two default rules that have 'conflicting' conclusions, both of whose antecedents are potentially satisfied by (4). The formalisation will have to ensure that in cases where there are two conflicting default rules, we gain a weaker conclusion; either the pushing caused the falling (due to (9)), or the descriptive order and temporal order match (due to *Narration*).[3]

The above rule of*Narration* is similar to Dowty's Temporal Discourse Interpretation Principle (TDIP), but there are two important differences. First, *Narration* is expressed in purely in terms of event times rather than *reference times*. In constrast to Dowty's, our analysis of tense will not be anaphoric. But we will see that the reader has beliefs about the relative temporal order of events, and so tense will behave *as if* it were anaphoric at the pragmatic level. The second important difference is that the applicability of *Narration* is dependent on the reader's WK; WK allows the application of *Narration* to (2) but not to (5). Dowty's TDIP is expressed as a principle that is independent of WK, and he therefore cannot account for the difference in the natural interpretations of (2) and (5).

Similarly, the current theories in DRT encode in the logical form of 'atomic' text that the descriptive order and temopral order of events match. Because it is encoded in logical form, it is a principle of semantic interpretation that is not cancellable. But this is contrary to intuitions, cf. text (5). To assign the correct interpretation to (5), DRT would have to assign it a different logical structure from other atomic texts that have the same surface syntax, or 'weaken' the semantic principle that descriptive order and temporal order match, so that it is cancellable, depending on what events are being talked about. The former option is undesirable, for texts with similar syntax would have distinct logical forms. But the latter option is more plausible. This is one reason why we make *Narration* a *defeasible* rule.

But what is the motivation for our assumption that narrative is the default structure for text? There are two motivations: the first concerns the order of perception of events, and the second concerns one of Grice's pragmatic maxims. As regards to the first, the narrator typically describes events in the order in which the protagonist views them. For events, which are punctual, the protagonist views them in exactly the order they occur, as is predicted by *Narration*. For states, which occur over extended periods, the protagonist may well not be aware of the state until after it has started to occur. In these situations, the narrator typically indicates the protagonist's inability to view the state at the point when it started by introducing it in the text only at the point where the protagonist perceives it. Text (12) is an example of this:

(12) Max entered the room. It was pitch dark.

By describing the state after the event, even though the room started to be dark before Max entered, the narrator is indicating that Max did not perceive darkness until *after* he had entered the room. Because the narrator typically describes eventualities in the order

[3]We could attempt to form a hierarchy of default rules so that (9) would take priority over *Narration* (or vice versa), but forming a hierarchy of default rules is a considerable problem in its own right, and is beyond the scope of this paper.

in which the protagonist views them, the temporal structure of text typically satisfies the above rule of *Narration*. Note that *Narration* does not exclude the possibility that a state described after an event (e.g. (12)) actually started to occur before the event did; it would merely describe the state as still holding sometime after the event occurred. But it must be emphasised that there are exceptions to narrative order; (5) is one of them.

The second motivation for *Narration* arises from the fact that it corresponds to Grice's (1975) conversational maxim of Manner, where it is suggested that discourse should be *orderly*: one should describe events in the right order. So in essence what we aim to do here is incorporate into the formal system one of Grice's pragmatic principles in the form of a default rule. The proposal that Grice's pragmatic principles should be thought of as default rules has already been suggested by Joshi & Webber (1984).

We have argued in favour of certain causal laws forming part of the reader's defeasible WK and the default rule of *Narration* forming part of the reader's defeasible linguistic knowledge (LK). We have claimed that these play a central role in explaining the natural interpretations of texts (2), (4) and (5). So to follow the proposed strategy, the formal framework must have at least two properties that correspond to the criteria *Epistemic Expressiveness* and *Causal Realism*:

- *Epistemic Expressiveness*: The framework must be able to represent WK and LK in the form of default rules (e.g. *ceteris paribus* John's poisoning Max causes Max's death), and the accompanying logic must be able to use these within the natural deduction system.

- *Causal Realism*: The framework must invoke a notion of causality (as some of the above default rules are about causality), and therefore it must treat events as entities that can stand in causal relations.

Although there are current theories in linguistics that incorporate events in a declarative framework, they have not so far exploited the notion of causality in characterising the interpretation of text, and so do not satisfy the above criterion of *Causal Realism*. Furthermore, no current logical treatment of the semantics of natural language (NL) fragments satisfies the above criterion of *Epistemic Expressiveness*. They all treat the interpretation of text as independent of the reader's WK. We aim to construct a framework that satisfies both criteria and so provides the first satisfactory explanation of the above texts.

In order to construct an event-based semantic representation scheme, I will capture some of Parsons' (1989) ideas in the theory presented here. Parsons' framework constitutes a thematic treatment of a natural language fragment, and he includes in his ontology events that are treated on a par with individuals, and referred to by terms in the language. Although he does not represent causality, it can easily be added to his scheme as a two-place predicate *cause* between terms that denote events (cf. Davidson (1967a)). We will argue below that in adding the causal predicate we satisfy the above criterion of *Causal Realism*.

In order to incorporate default reasoning techniques into the framework, I will present the theory within an extension of Levesque's (1985) default logic. The extension will arise from incorporating events in the ontology and causality in the language.

Contrary to the traditional non-monotonic approach to default reasoning (Reiter (1980), McDermott & Doyle (1980)), Levesque's logic is completely monotonic, in that adding clauses to the antecedent of a true conditional never requires consequences to be removed in order to preserve truth. This monotoncity is a property of traditional logics which already have well defined models and proof theoretic results. Levesque's strategy is to represent default rules by using a standard modal belief operator that has well understood logical properties (cf. the auto-epistemic approach to default reasoning pioneered by Moore (1985) and Konolige (1987)). This gives Levesque's logic a distinct advantage over non-monotonic logics, which are not so well established from a logical point of view.

By representing the semantics of text within an extension of Levesque's theory we will achieve the following important goals: first, we will describe the semantics of text within a logic that can handle default rules and so we will achieve the requirement of *Epistemic Expressiveness* described above. Secondly, one can regard the agent of Levesque's belief operator as the reader of the text. So we are able to represent within a single unified theory both the truth conditions of a text, and the reader's beliefs about what is described by a text given his other beliefs (such as the above default rules).

The truth conditions of the text define the text's semantic properties. The reader's beliefs that follow via the logic of belief from his knowledge base (KB) and his belief of the text intuitively define the text's *pragmatic* properties. We will show that the entailments from the text's truth conditions (i.e. semantic entailments) won't necessarily match the reader's inferences from the text (i.e. pragmatic entailments).

The details of the analyses of (2), (4) and (5) are set out below. These three texts will share common semantic entailments, but (2) will not share the pragmatic implications derived from (4) and (5). The theory will be able to derive via *natural deduction* that in the case of (2) the reader will believe the turning around precedes the hitting (this is due to *Narration*); from believing (4) the reader will believe that either the falling precedes the pushing or the pushing causes the falling (this is due to *Narration* and (9)); from believing (5) the reader will believe that the poisoning causes the death (this is due to (10) plus the *indefeasible* knowledge that one cannot poison someone who is already dead and so one cannot apply *Narration*). This will hold in spite of the fact that other temporal relations would have been warranted by semantics. So in essence by representing both the semantic properties (i.e. truth conditions) *and* the pragmatic properties (i.e. the reader's beliefs) of text within a single modal logic with well understood logical properties, we will distinguish the text's semantics from its pragmatics, and this will lead to an account of the distinct natural interpretations of (2), (4) and (5) in spite of the fact that they are assigned similar logical forms: this is just as required.

2.1 The Scope of the Theory

It is important to stress the limited scope of the problem area that this paper will address. We will only explore the problem of characterising the temporal relation between events described within the *same* discourse segment; we will not represent those texts that feature embedded discourse structures, e.g. text (13).

(13) a. John went to Mary's house.
b. She told him about her trip to Iceland.
c. She saw many geesers, and was very impressed with the landscape.
d. She got dreadful food poisoning, and came home early to recover.
e. John was very sorry to hear about how ill she had been.

(13) has an embedded discourse segment in the sense that sentences (13c,d) are about a different event (Mary's trip to Iceland) from the main segment (which is about John's visit to Mary's house). Webber (1987) argues that in order to represent texts like (13), the notion of a *temporal focus* must be incorporated into the theory. In Webber's terms, the temporal focus is the event which is most salient and to which the eventuality being described is most closely associated. Webber's temporal focus is the analogue in the temporal domain to Sidner's (1982) *discourse focus* in the nominal domain. The temporal focus would indicate the embedding in text (13) by predicting that (13c,d) has a different temporal focus from the main segment.

But the problem of keeping track of the temporal focus as text gets processed is a deep one, especially if it is to be explained in a non-procedural framework like the one we will develop here. So for reasons of exposition, we will not investigate the problem of representing embeddeings. We will address only the related but separate problem of how to characterise the temporal relation between events that are described within the *same* discourse segment. In other words, we will only consider those texts that have a 'flat' discourse structure, i.e. no embeddings, like texts (14) and (15).

(14) Max arrived at the house.
Mary ran into the driveway.
She greeted him with a huge smile
They walked into the house arm in arm.

(15) Max went to the florists' shop.
He picked out three red roses.
He took them to Mary's house.

In limiting the scope of this paper to those texts with flat discourse structure, we will be able to consider the relation between *pairs* of sentences as in (2), (4) and (5), and the subsequent theory will straighforwardly extend to those texts with more than two clauses that do not contain embeddings.

3 The Treatment of Events and Causality

How are we to represent a sentence like (16)?

(16) John hit Max.

We want to represent (16) using a term that refers to the event of John hitting Max. We will therefore include in our language, as Parsons (1989) does, terms that refer to *eventualities*. Following Parsons, the set of eventualities includes the subclasses of events (e.g. *John hit Max*) and states (e.g. *Max was upset*). The subclass of processes (e.g. *Max cried*) will not concern us in this paper; the theory will eventually have to be extended to deal with these. The symbols $e_1, e_2, \ldots$ represent the eventuality constants, and $e'_1, e'_2, \ldots$ are the variables.

Parsons persuasively argues in favour of a thematic representation of natural language clauses such as (16). We will also follow this strategy. The logical form of (16) will be (16a).

(16a) $(\exists t)(hitting(e) \wedge subject(john, e) \wedge object(max, e) \wedge hold(e, t) \wedge t < now)$

In words, (16a) describes an eventuality e that is a hitting, the subject of e is John, the object of e is Max, and e holds at a point of time t that is earlier than now (i.e. the time of speech). *now* is a constant referring to a point of time. We use an event constant in (16a) to indicate that the sentence is about a *particular* event identified by the context of utterance. Contrary to Partee's (1973) thesis, however, the tense in (16a) is not anaphoric with respect to time; at least, not with respect to the past time. Nevertheless, we will subsequently show that the reader (i.e. the agent of our belief operator) will behave *as if* tenses are anaphoric, in that he will infer beliefs about the identity of the time t where the event holds given the existing context.

Parsons' representation of (16) differs from ours in at least one important respect; he uses the 2-place predicate *cul* (standing for *culminates*) in place of our predicate *hold*. He does this because he distinguishes the culmination of an event from its development portion (i.e. that part of the event that leads to the culmination) by including in his language the logically-unrelated formulae $cul(e, t)$ (meaning e culminates at t), and $hold(e, t)$ (meaning e is in development at t).

In Lascarides (1988), we demonstrated that irresolvable problems arise from Parsons' strategy of naming the development portion and culmination by the same event term e; it becomes impossible to give an adequate characterisation of point adverbials such as *at 3pm*. We showed that many problems could be solved if the event term e refers to the *culmination alone*, making e a punctual entity holding at an isolated point of time. The development portion of the event e would have to be denoted by some term other than e itself; for example we could introduce a function PR from event terms to process terms where $PR(e)$ denotes the development portion of e. This option is explored in Lascarides (1988). Our assumption that e refers only to the (punctual) culmination is in contrast to Parsons' treatment of event terms e, where e holds at all points of time within a connected extended interval identifying the development portion as well as the culmination. According to our interpretation, $hold(e, t)$ means the event e's *culmination* holds at time t.

However, by far the most significant difference between Parsons' theory and the one presented here arises from the fact that Parsons' interpretation of (16a) is first order and ours must be modal. As we have mentioned, we intend to satisfy the above critierion of *Epistemic Expressiveness* by representing default rules in terms of a belief operator with a modal interpretation (cf. Hintikka (1962), Levesque (1990)). So (16a) will have to be assigned a modal interpretation, where its truth value must be defined relative to possible worlds.

Only then can (16a) be an object of belief. We are therefore obliged to clarify how events will behave with respect to the possible worlds semantics.

3.1 Events and Possible Worlds

Incorporating terms that denote eventualities into a modal framework that evaluates semantics relative to possible worlds has not been explored elsewhere. However, because eventualities are treated exactly on a par with individuals, the problems of accounting for the eventualities' behaviour with respect to possible worlds semantics are no worse than those for individuals.

We will treat all constants and variables as *rigid designators*. In so doing, we avoid a lot of the traditional puzzles of identifying objects across possible worlds (cf. Kripke (1972)).[4] However, both individual objects and eventualities may be qualified by different predicates in different possible worlds.

Save its name, an eventuality need not have any common properties across possible worlds. In particular, $hold(e,t)$ may be false for every time t with respect to w. So we have eventualities that need not occur at any time in some worlds. The fact that eventualities may lack occurrence in some worlds may appear mysterious at first glance, but perhaps such eventualities can best be thought of as those that are referred to in counterfactual statements like (17).

(17) The premature resignation of Reagan would have been a good thing.

We have chosen to interpret well-formed expressions of the language modally with respect to possible worlds, but *not* modally with respect to times as well. This is in contrast to the standard Montagovian treatment of modal interpretation, where expressions are interpreted with respect to a world-time index. We have decided for expository reasons to maintain object language expressions denoting times and use these to express when events occur. Following the object language route will enable the semantics of temporal expressions to be clear from their syntactic configeration. Also, by preventing a modal interpretation on a world-time index we stay faithful to Levesque's belief logic, and so we are not obliged to add axioms saying how beliefs can change with respect to time. That problem deserves considerable attention in its own right, and since it is not our concern here, it is best avoided.

The truth value of $hold(e,t)$ can vary across possible worlds, even though the denotations of e and t do not. We will nevertheless apply certain restrictions on the set of times where an eventuality holds in any given world w. One restriction is that this set of points of time forms a single connected set. In the case of events, the formula (18) is true relative to any given possible world w and for every event e'.

(18) $(\forall t'_1)(\forall t'_2)((hold(e',t'_1) \wedge hold(e',t'_2)) \Rightarrow t'_1 = t'_2)$

So events hold on at most a unique point of time in any given world w. Because of this, the terms referring to events (and, in fact, all eventualities) are best thought of as referring

[4]Because terms are rigid designators, our theory will be unable to account for referential opacity within belief contexts. But we are not concerned with this problem here.

to event tokens rather than event-types, even though we can describe the properties of an eventuality e differently in different worlds. Events are tokens, for if they occur in a world at all, then they do so only once.

Our formalism will reflect the intuition that, in contrast to the punctual structure of culminations, *states* occur over an extended period of time with no definite endpoints (cf. Moens (1987)). We will assume that the set of points where a state s holds relative to a possible world w either forms the empty set (in which case s doesn't hold) or an open interval.

In contrast to Parsons, our theory will include terms that refer to intervals of points of time. $i_1, i_2, \ldots$ will be the constants and $i'_1, i'_2, \ldots$ the variables. The intervals in the model will constitute the connected sets of points of time, where the set of points of time is dense so as to gain a distinction between open and closed intervals.

A 2-place predicate *occ* will identify relative to a world w the *interval* (as opposed to point of time) over which an eventuality e occurs in w, from e's start to its finish. The predicates *occ* and *hold* are related in the following way:

(19) $$(\forall e')(\forall i')(\forall t')(occ(e', i') \Leftrightarrow (t' \in i' \Leftrightarrow hold(e', t')))$$

$occ(e, i)$ will be true of an event e and interval i only if i is a singleton, and it will be true of a state e and interval i only if i is open. Furthermore, for every eventuality e, $occ(e, i)$ will be true relative to a world w for at most a unique interval i.

For notational convenience, we also introduce an ordering predicate on eventualities which will be defined in terms of the ordering on intervals. More specifically, $e_1 < e_2$ will intuitively mean that (the whole of) e_1 occurs before (the whole of) e_2. The ordering on eventualities and the ordering on intervals are related as follows:[5]

(20) $$(\forall e_1)(\forall e_2)(e_1 < e_2 \Leftrightarrow (\exists i_1)(\exists i_2)(occ(e_1, i_1) \wedge occ(e_2, i_2) \wedge i_1 < i_2))$$

3.2 Causality

We will in line with Davidson (1967a) incorporate into the formal language a two-place predicate *cause* which takes eventualities as arguments. The formula $cause(e_1, e_2)$ is interpreted as the eventuality e_1 causes the eventuality e_2. As we have mentioned, e_1 and e_2 are best thought of as singular eventualities rather than eventuality-types, so the predicate *cause* represents singular causal statements, even though $cause(e_1, e_2)$ may have different truth values in different worlds.

Our representation of causality will not illuminate any of the philosophical puzzles associated with the notion. We merely wish to utilize causality in order to account for the temporal structure of text. We will reveal almost nothing about what causality itself actually *means*, for its semantics will be treated as primitive.

In order to reflect in our theory the intuition that a cause must precede its effect, we assume that (21) holds, i.e. if e_1 causes e_2 and e_1 is an event, then e_1 occurs before e_2.

[5] Where there is no confusion, I will occasionally abandon placing primes ($'$) on all the variables.

(21) **Law of Forward Causality**
$(\forall e_1)(\forall e_2)((cause(e_1, e_2) \wedge event(e_1)) \Rightarrow e_1 < e_2)$

(21) will be true for every possible world; it has the status of a non-logical axiom. Note that if e_1 is a *state* and $cause(e_1, e_2)$ is true, then it doesn't necessarily follow that $e_1 < e_2$. This agrees with intuitions, cf. the road is slippery causing John to slip.[6]

3.3 The Model

Now that the important aspects of the object language are in place, we will define the model that will be used in evaluating its semantics. The model will be a 9-tuple $\langle W, D, T, I, E, <_T, <_I, \Re, f\rangle$ where:

- W is a non-empty set of possible worlds.
- D is a non-empty set of individuals.
- T is a non-empty dense set of points of time.
- I is a non-empty set of intervals of points of time; it constitutes all the connected sets in the power set of T apart from $\emptyset$.
- E is the set of eventualities, which includes the set of events and set of states.
- $<_T$ is a temporal ordering relation on T.
- $<_I$ is a temporal ordering relation on I, such that $i_1 <_I i_2$ if and only if for all t_1 in i_1 and t_2 in i_2, $t_1 <_T t_2$.[7]
- $\Re$ is a function from W to the power set of W (this will eventually be used to define the semantics of the belief operator, and will characterise the doxastic accessibility relation).
- f is the interpretation function that assigns intensional values to the non-logical constants.

All well-formed expressions in the language are interpreted with respect to a pair $\langle M, g\rangle$ where M is a model for the language and g is a variable assignment function.

The semantic assignments of the expressions of the language are evaluated relative to possible worlds in the usual way, subject to the constraint that constants and variables are rigid designators. The complete truth definition is given in the appendix. We will use $[\![\phi]\!]^{\langle M,g\rangle}(w)$ to denote the semantic value of the well-formed expression ϕ relative to the possible world w with respect to the model M and variable assignment function g. We may abbreviate this occasionally to $[\![\phi]\!](w)$ where there is no confusion.

---

[6]The predicates *event* and *occ* are related as follows:
$(\forall e)(event(e) \Rightarrow (\forall i)(occ(e, i) \Rightarrow singleton(i)))$

[7]Where there is no confusion we may use $<$ instead of $<_T$ and $<_I$.

4 The Representation of Text

As we have mentioned, the logical form of sentence (16) is (16a).

(16) John hit Max.

(16a) $(\exists t)(hitting(e) \wedge object(max, e) \wedge subject(john, e) \wedge hold(e, t) \wedge t < now)$

In words, (16a) is true with respect w just in case the eventuality e is a hitting whose subject is John and object is Max in w, e holds at a point of time t in w, and t is earlier than now (i.e. the time of speech).

The logical form of *text* will be an ordered sequence of the logical form of its sentences: If $c_1 \ldots c_n$ is a text of n sentences, then its logical form is $\langle c'_1, \ldots, c'_n \rangle$, where $c'_1, \ldots, c'_n$ are the logical forms of $c_1, \ldots, c_n$ respectively.

$\langle c'_1, \ldots, c'_n \rangle$ can be thought of as an ordered syntactic array of object language expressions; the ordering itself is not part of the object language and so does not require a semantics. We represent text as an ordered sequence of the logical forms of its sentences in order to keep track of the *descriptive order* of events. This will be important in our attempt to answer the Ordering Question, which concerns the relation between descriptive order and temporal order.

(2), (4) and (5) are all texts and their logical forms as such are (2a), (4a) and (5a).

(2) Max turned round. John hit him.

(2a) $\langle (\exists t_1)(turning - round(e_1) \wedge subject(max, e_1) \wedge hold(e_1, t_1) \wedge t_1 < now),$
$(\exists t_2)(hitting(e_2) \wedge subject(john, e_2) \wedge object(max, e_2) \wedge hold(e_2, t_2) \wedge t_2 < now) \rangle$

(4) Max fell. John pushed him.

(4a) $\langle (\exists t_1)(falling(e_1) \wedge subject(max, e_1) \wedge hold(e_1, t_1) \wedge t_1 < now),$
$(\exists t_2)(pushing(e_2) \wedge subject(john, e_2) \wedge object(max, e_2) \wedge hold(e_2, t_2) \wedge t_2 < now) \rangle$

(5) Max died. John poisoned him.

(5a) $\langle (\exists t_1)(dying(e_1) \wedge subject(max, e_1) \wedge hold(e_1, t_1) \wedge t_1 < now),$
$(\exists t_2)(poisoning(e_2) \wedge subject(john, e_2) \wedge object(max, e_2) \wedge hold(e_2, t_2) \wedge t_2 < now) \rangle$

(2a) is true with respect to w just in case e_1 is a turning round by Max that holds at a time t_1 earlier than now in w; and e_2 is a hitting by John of Max that holds at a time t_2 earlier than now in w. Furthermore (2a) indicates that the text described e_1 before it described e_2; this is reflected in (2a)'s ordered sequence. (4a) and (5a) are exactly analogous to this, save that turning round/hitting are replaced by falling/pushing and dying/poisoning. Note that (2a), (4a) and (5a) do not place *any* conditions whatsoever on the order in which e_1 and e_2 occur. In this sense the semantics of (2), (4) and (5) are weaker than one might expect.

Furthermore, the semantics of (2), (4) and (5) are all similar, contrary to what one might expect given their distinct natural interpretations. We will subsequently show that (2),

(4) and (5) have very different pragmatic properties however, arising from the reader's distinct defeasible knowledge concerning (2), (4) and (5). The formalisation of this defeasible knowledge will now be developed.

5 Levesque's Default Logic

5.1 The Belief Operators

We will represent the reader's defeasible WK and LK using Levesque's belief logic. Levesque claims that traditional logics of belief cannot adequately represent default reasoning tasks. For (a) and (b) are not sufficient to derive (c).

(a) I believe Tweety is a bird.
(b) If a bird can be consistently believed to fly, then I believe it flies.
(c) I believe Tweety flies.

Rather, one needs (a), (b) and (a′) to derive (c); (a′) will tell one that (b) applies to Tweety.

(a′) (a) is *all* I believe about Tweety.

But (a′) cannot be expressed in traditional logics of belief.

Levesque augments a logic of belief so that propositions similar to (a′) can be expressed directly within the language. There are two modal operators, B and O, where $B\phi$ is read as "ϕ is believed" and $O\phi$ is read as "ϕ is *all* that is believed". (a′) is expressed in terms of O, and so O will be used to represent part of the default reasoning process.

The semantics of O looks remarkably like that of B; we will examine B first. Levesque uses the traditional procedure for providing the truth conditions of B. That is, he defines $B\phi$ in terms of an accessibility relation over possible worlds:

- $[\![B\phi]\!]^{\langle M,g\rangle}(w) = 1$ if and only if $w' \in \Re(w)$ entails $[\![\phi]\!]^{\langle M,g\rangle}(w') = 1$

This is exactly analogous to the strategy used by Hintikka (1962) in defining the semantics of belief.[8] Beliefs are those sentences that are true in all accessible worlds. The more that is believed, the more conditions are satisfied by the accessible worlds, and so the smaller the set of accessible worlds must be. So if ϕ is *all* that is believed, i.e. $O\phi$ is true, then the set of accessible worlds is *as large as possible* subject to being consistent with believing ϕ, since the larger the set, the less knowledge is represented. This leads to the following truth conditions for $O\phi$:

- $[\![O\phi]\!]^{\langle M,g\rangle}(w) = 1$ if and only if $w' \in \Re(w)$ if and only if $[\![\phi]\!]^{\langle M,g\rangle}(w') = 1$

Levesque assumes that $\Re$ corresponds to the relation defined by weak S5 (also known as K45); i.e. it is serial and transitive, but unlike S5 proper, it is not reflexive. Furthermore, $\Re$

[8] Levesque explicitly refers to an accessible set of worlds W rather than to the function $\Re$ in his definition of B. We have included $\Re$ in the above definition for expository reasons, to show that Levesque's semantics for B is the classical one.

characterises a Euclidean relation rather than symmetry: if $w', w'' \in \Re(w)$, then $w'' \in \Re(w')$ and $w' \in \Re(w'')$. But it is not necessarily the case that if $w' \in \Re(w)$, then $w \in \Re(w')$.

Levesque shows how O can be defined in terms of two belief operators B and N where N is defined as follows:

- $[\![N\phi]\!]^{\langle M,g\rangle}(w) = 1$ if and only if for every w', $[\![\phi]\!]^{\langle M,g\rangle}(w') = 0$ entails $w' \in \Re(w)$

Levesque suggests that $N\phi$ can be read as "at most ϕ's belief status is false". But given the truth conditions of $N\phi$, a more accurate description is "if ϕ is *falsifiable*, then either it is not believed, or believed to be false (hence at most its belief status is false); otherwise, if ϕ is valid, then $N\phi$ is true trivally." So $B\phi \wedge N\neg\phi$ should be read as "at least ϕ is believed to be true and at most ϕ is believed to be true"; that is, "exactly ϕ is believed to be true", which is just what $O\phi$ is supposed to represent. And sure enough, one can prove that $O\phi = B\phi \wedge N\neg\phi$.

N behaves exactly like a belief operator, but with respect to the compliment of the $\Re$ relation, for the above definition is equivalent to the following:

- $[\![N\phi]\!]^{\langle M,g\rangle}(w) = 1$ if and only if for every w', $w' \in \overline{\Re}(w)$ entails $[\![\phi]\!]^{\langle M,g\rangle}(w') = 1$

This allows Levesque to produce a proof theory that is very similar to what would be done for two separate believers. But there are two important differences. First, the two 'agents' are fully introspective and know about each other's beliefs and non-beliefs; to handle this axioms such as $N\phi \Rightarrow BN\phi$ are added, as well as the usual $B\phi \Rightarrow BB\phi$. Secondly, for any falsifiable ϕ, ϕ's belief status cannot be at most false *and* at least true; this yields axiom (2) below. Overall, then, the proof theory is formed by adjoining the standard objective axioms with the following:

(1) The axioms for weak S5, for both B and N, which are the following:
Where L is either B or N:
  (a) $L\phi$, where ϕ is any valid objective sentence,
  (b) $L(\phi \Rightarrow \psi) \Rightarrow (L\phi \Rightarrow L\psi)$
  (c) $(\forall x)L\phi \Rightarrow L(\forall x)\phi$
  (d) $\sigma \Rightarrow L\sigma$, where σ is subjective
  (σ is subjective if all non-logical symbols are within the scope of a belief operator.
  One can in fact show that $BB\phi \Rightarrow B\phi$ is also a theorem),

(2) $N\phi \Rightarrow \neg B\phi$, where ϕ is any objective sentence that is falsifiable.

(3) $O\phi \Leftrightarrow (B\phi \wedge N\neg\phi)$, for any ϕ.

According to axiom (1), both B and N are logically omniscient; they essentially correspond, logically speaking, to the strongest possible belief operators one can have, without assuming that belief implies truth. So neither B nor N characterise a psychologically plausible notion of belief. But then our aim is not one of psychological modelling.

Levesque shows that $O\phi$ is true with respect to $\langle M, g\rangle$ at the index w only if ϕ is a *complete* statement of the agent's beliefs at w, in the sense that all the agent's beliefs are derivable from ϕ using the above axioms (1) and (2) alone. So the formal semantics of $O\phi$ matches the formula's intended intuitive meaning, that ϕ is *all* that is believed.

Levesque defines the notion of theoremhood, satisfaction and validity in the obvious way: ϕ is satisfiable if there is an interpretation $\langle M, g\rangle$ and a world w such that $[\![\phi]\!]^{\langle M,g\rangle}(w) = 1$. ϕ is valid if its negation is not satisfiable. ϕ is a theorem if it is derivable from the above axioms. The results below follow:

- *Soundness*: Every theorem is valid.
- *Propositional Completeness*: If ϕ is in the propositional subset, then it is a theorem if and only if it is valid.

5.2 The Representation of Default Rules

Levesque represents default rules like (b) as (22) (cf. (9) and (10)):

(b) One believes for any object that if it is a bird, then it flies unless one already has reason to believe something to the contrary.

(22) $B(\forall x)((bird(x) \wedge \neg B \neg fly(x)) \Rightarrow fly(x))$

So the object of belief in (22) is part of the reader's knowledge base (KB).

Now having the above default rule together with Tweety is a bird in the KB does not entail the belief that Tweety flies. But Levesque shows that if these two things are *all* that is believed, then the belief that Tweety flies does follow:

Theorem 1

Let $\beta = (\forall x)((bird(x) \wedge \neg B \neg fly(x)) \Rightarrow fly(x))$.
Then
$O(bird(tweety) \wedge \beta) \Rightarrow Bfly(tweety)$
is a theorem.

Moreover, in the case when one believes that Chilly is a bird that can't fly, Levesque shows that one can derive that Tweety can still fly as long as Tweety is not Chilly.

We have now seen how Levesque accounts for so-called non-monotonic reasoning.[9] We will incorporate his belief operators into our framework, so that we too can handle default reasoning within the natural deduction system in our theory, and so fulfil the critierion of *Epistemic Expressiveness* as defined above. Our framework can be viewed as an extension of Levesque's; the extension being that we have four sorts of constants and variables — individuals, eventualities, points of time and intervals of time — as opposed to Levesque's one sort (which represents individuals).

[9]The non-monotonicity is not in the logic, but instead is pushed into the scope of the operator O; $O\alpha$ can entail $B\beta$ even if $O(\alpha \wedge \gamma)$ does not.

6 Causal Defaults and the Narrative Default

We have already argued for the following default rules:[10]

(9) $B(((\text{Max fell}) \wedge (\text{John pushed Max}) \wedge \neg B\neg(\text{John's pushing Max caused Max's fall})) \Rightarrow (\text{John's pushing Max caused Max's fall}))$

(10) $B(((\text{Max died}) \wedge (\text{John poisoned Max}) \wedge \neg B\neg(\text{John's poisoning Max caused Max's death})) \Rightarrow (\text{John's poisoning Max caused Max's death}))$

- *Narration*: If a text describes an eventuality e_1 before it describes e_2, then unless the reader has reason to believe something to the contrary, he believes that there are times t_1 and t_2 such that e_1 holds at t_1, e_2 holds at t_2, and $t_1 < t_2$.

Our task now is to represent these using Levesque's operators B, N and O, and use these rules to account for the differences in our natural interpretations of texts (2), (4) and (5).

(2) Max turned round. John hit him.

(4) Max fell. John pushed him.

(5) Max died. John poisoned him.

The rules (9), (10) and *Narration* can be represented in an exactly analogous way to the default rule β in the Tweety example. That is, the objects of belief corresponding to (9), (10) and *Narration* are β_1, β_2 and β_3 respectively.

$$\begin{aligned}\beta_1 \quad (\forall e_1' e_2' xy) \quad & ((falling'(e_1') \wedge subject(x, e_1') \wedge (\exists i_1')occ(e_1', i_1') \wedge \\ & pushing'(e_2') \wedge subject(y, e_2') \wedge object(x, e_2') \wedge (\exists i_2')occ(e_2', i_2') \wedge \\ & \neg B \neg cause(e_2', e_1')) \Rightarrow cause(e_2', e_1'))\end{aligned}$$

$$\begin{aligned}\beta_2 \quad (\forall e_1' e_2' xy) \quad & ((dying'(e_1') \wedge subject(x, e_1') \wedge (\exists i_1')occ(e_1', i_1') \wedge \\ & poisoning'(e_2') \wedge subject(y, e_2') \wedge object(x, e_2') \wedge (\exists i_2')occ(e_2', i_2') \wedge \\ & \neg B \neg cause(e_2', e_1')) \Rightarrow cause(e_2', e_1'))\end{aligned}$$

$$\begin{aligned}\beta_3 \quad (\forall e_1' e_2') \quad & (\langle \phi_1(e_1'), \phi_2(e_2') \rangle \wedge \neg B \neg (\exists t_1' t_2')(hold(e_1', t_1') \wedge hold(e_2', t_2') \wedge t_1' < t_2')) \Rightarrow \\ & (\exists t_1' t_2')(hold(e_1', t_1') \wedge hold(e_2', t_2') \wedge t_1' < t_2')\end{aligned}$$

where ϕ_1 and ϕ_2 are λ-expressions of the form $\lambda e \psi$.

So β_1, β_2 and β_3 will be part of the reader's KB, and so can be viewed as conjuncts that feature in the formula ϕ such that $O\phi$ is true.

β_1 in words states that for all eventualities e_1' and e_2' and individuals x and y, if e_1' is a falling by x that occurs and e_2' is a pushing by y of x that occurs, and if it is not the case that e_2' doesn't cause e_1' is believed, then e_2' causes e_1'. β_2 is the same save that we substitute dying/poisoning in place of falling/pushing. β_3 states that for any e_1' and e_2', if e_1' is described before e_2' in a text (as shown by the ordered sequence $\langle \phi_1(e_1'), \phi_2(e_2') \rangle$), and it is not the case that e_2' occurs *after* e_1' is believed, then e_1' holds at a time t_1' and e_2' holds at t_2', where t_1' is earlier than t_2'.

[10] One can think of "Max" and "John" as variables in (9) and (10).

The motivation for having the rules β_1 and β_2 as objects of belief is *independent* of linguistic knowledge. β_1 and β_2 merely state generalisations about our perception of how events in the world happen. They constitute defeasible WK, whereas β_3 is defeasible LK.

In addition to the above default rules, the reader's KB must include rules that are indefeasible; after all, not all conclusions that we draw from the world are tentative. An example of such knowledge is that one cannot poison someone who has already died. This fact is represented in the rule α_1 below, and so we will assume that α_1 is a conjunct in the KB:

$$\alpha_1 \quad (\forall e'_1 x t'_1 e'_2 t'_2) \quad \begin{array}{l} (dying'(e'_1) \wedge subject(x, e'_1) \wedge hold(e'_1, t'_1) \wedge \\ poisoning'(e'_2) \wedge object(x, e'_2) \wedge hold(e'_2, t'_2) \Rightarrow t'_2 < t'_1) \end{array}$$

Rule α_1 will shortly be used to show that the reader will believe that the linguistic order of (5) mismatches temporal order.

More generally, in our theory we are able to represent: the semantics of text, the reader's defeasible knowledge (some of which is world knowledge and some of which is linguistic knowledge), and the reader's indefeasible knowledge. We will show how these interact given our logic of belief to yield the natural interpretations of the above texts.

7 The Application of the Default Rules

We have already given the truth conditions of texts (2), (4) and (5). Now that the logic of belief and the KB statements are in place, we can also state what beliefs follow from believing (2), (4) or (5) together with the KB.

We start with (2). The logical form of (2) is (2a) below.

(2) Max turned round. John hit him.

(2a) $\langle turning-round'(e_1) \wedge subject(max', e_1) \wedge hold(e_1, t_1) \wedge t_1 < now,$
$hitting'(e_2) \wedge subject(john', e_2) \wedge object(max', e_2) \wedge hold(e_2, t_2) \wedge t_2 < now\rangle$

(2a) places no conditions on the temporal relations between the turning-round and hitting events. However, the reader's beliefs will not be so weak. We will show that according to the logic of belief and KB stated above, the reader will believe that the turning-round *precedes* the hitting. This agrees with intuitions.

Theorem 2

Let KB_2 =(2a)$\wedge\beta_1 \wedge \beta_2 \wedge \beta_3 \wedge \alpha_1$

Then
$O(\text{KB}_2) \Rightarrow B(e_1 < e_2)$
is a theorem.

That is, if the reader believes (2a) and the α and β rules alone, then he believes that the turning-round (e_1) precedes the hitting (e_2).

Proof 2

In what follows, LO stands for logical omniscience and PL for predicate logic.

a.	$O(\text{KB}_2)$	assumption
b.	$B(\text{KB}_2)$	a; axiom (3)
c.	$B((2\text{a}) \wedge \beta_3)$	b; LO
d.	$B(\neg B \neg (\exists t'_1 t'_2)(hold(e_1, t'_1) \wedge hold(e_2, t'_2) \wedge t'_1 < t'_2) \Rightarrow$ $(\exists t'_1 t'_2)(hold(e_1, t'_1) \wedge hold(e_2, t'_2) \wedge t'_1 < t'_2))$	c; LO
e.	$B \neg B \neg e_1 < e_2 \Rightarrow B(e_1 < e_2)$	d; LO, punctual events
f.	$\neg B \neg B \neg e_1 < e_2 \vee B(e_1 < e_2)$	e; PL
g.	$B \neg e_1 < e_2 \vee B e_1 < e_2$	f; axiom (1d)
h.	$N \neg (\text{KB}_2)$	a; axiom (3)
i.	$N(((2\text{a}) \wedge \beta_1 \wedge \beta_2 \wedge \alpha_1) \Rightarrow \neg \beta_3)$	h; LO

Let $\beta_3 =_{def} (\forall e'_1 e'_2)(\beta_3^{ante} \Rightarrow \beta_3^{conseq})$

j.	$N(((2\text{a}) \wedge \beta_1 \wedge \beta_2 \wedge \alpha_1) \Rightarrow (\exists e'_1 e'_2) \neg (\beta_3^{ante} \Rightarrow \beta_3^{conseq}))$	i; LO
k.	$N(((2\text{a}) \wedge \beta_1 \wedge \beta_2 \wedge \alpha_1) \Rightarrow (\exists e'_1 e'_2)(\beta_3^{ante} \wedge \neg \beta_3^{conseq}))$	j; LO
l.	$N(((2\text{a}) \wedge \beta_1 \wedge \beta_2 \wedge \alpha_1) \Rightarrow$ $(\exists e'_1 e'_2)(\beta_3^{ante} \wedge e'_1 \neq e_2 \wedge e'_2 \neq e_1 \wedge \neg \beta_3^{conseq}))$	k; LO[11]
m.	$N(((2\text{a}) \wedge \beta_1 \wedge \beta_2 \wedge \alpha_1) \Rightarrow (\exists e'_1 e'_2)(e'_1 \neq e_2 \wedge e'_2 \neq e_1 \wedge \neg \beta_3^{conseq}))$	l; LO
n.	$\neg B(((2\text{a}) \wedge \beta_1 \wedge \beta_2 \wedge \alpha_1) \Rightarrow (\exists e'_1 e'_2)(e'_1 \neq e_2 \wedge e'_2 \neq e_1 \wedge \neg \beta_3^{conseq}))$	m; axiom(2)
o.	$\neg (B((2\text{a}) \wedge \beta_1 \wedge \beta_2 \wedge \alpha_1) \Rightarrow B(\exists e'_1 e'_2)(e'_1 \neq e_2 \wedge e'_2 \neq e_1 \wedge \neg \beta_3^{conseq}))$	n; LO
p.	$B((2\text{a}) \wedge \beta_1 \wedge \beta_2 \wedge \alpha_1) \wedge \neg B(\exists e'_1 e'_2)(e'_1 \neq e_2 \wedge e'_2 \neq e_1 \wedge \neg \beta_3^{conseq})$	o; PL
q.	$\neg B(\exists e'_1 e'_2)(e'_1 \neq e_2 \wedge e'_2 \neq e_1 \wedge \neg \beta_3^{conseq})$	p; PL
r.	$(\forall e'_1 e'_2) \neg B(e'_1 \neq e_2 \wedge e'_2 \neq e_1 \wedge \neg \beta_3^{conseq})$	q; axiom (1c), PL
s.	$\neg B \neg (\exists t'_1 t'_2)(hold(e_1, t'_1) \wedge hold(e_2, t'_2) \wedge t'_1 < t'_2)$	r; PL, substitute for β_3^{conseq}
t.	$\neg B \neg e_1 < e_2$	s; punctual events
u.	$B(e_1 < e_2)$	g,t; PL
v.	$O(\text{KB}_2) \Rightarrow B(e_1 < e_2)$	a,u; PL
		Q.E.D.

Because belief is logically omniscient in this theory, one cannot believe contradictory formulae. Therefore, since the law of Forward Causality is a non-logical axiom of the system,

(21) The Law of Forward Causality
$(\forall e_1)(\forall e_2)(cause(e_1, e_2) \Rightarrow e_1 < e_2)$

the belief that the hitting does *not* cause the turning round logically follows from believing KB_2 alone. Thus our theory predicts that text (2) misleads the reader if the situation being described is one where John's hitting Max caused Max to turn round (to face John). This is as we required. A similar proof to the above shows that by believing text (1), the reader believes the hitting precedes the turning round.

[11] l follows from k via LO because given that (2a) is believed, we have
$(\exists e'_1 e'_2)(\beta_3^{ante}) \Rightarrow (\exists e'_1 e'_2)(\beta_3^{ante} \wedge e'_1 \neq e_2 \wedge e'_2 \neq e_1)$.
This is because the substitution for e'_1 as e_2 and e'_2 as e_1 will not satisfy β_3^{ante} since it stipulates that the description of e'_1 *precedes* the description of e'_2. We need to add the conjunct $e'_1 \neq e_2 \wedge e'_2 \neq e_1$ in the above proof so that the formula in the scope of the operator N in line m is falsifiable.

(1) John hit Max. Max turned round.

This is compatible with the hitting *causing* the turning round. And so in contrast to (2), (1) does *not* mislead the reader if the situation being described is one where the hitting caused the turning round.

Note that the only default rule that played a significant part in the above proof was β_3. This is as one would expect, since WK about falling, pushing, poisoning and dying should not affect one's beliefs about (2). The above results are in sharp contrast to those gained from believing texts (4) and (5), as I will now show. We start with text (4), whose logical form is (4a).

(4) Max fell. John pushed him.

(4a) $\langle falling'(e_1) \wedge subject(max', e_1) \wedge hold(e_1, t_1) \wedge t_1 < now,$
$pushing'(e_2) \wedge subject(john', e_2) \wedge object(max', e_2) \wedge hold(e_2, t_2) \wedge t_2 < now\rangle$

Theorem 4

Let $\text{KB}_4 =$(4a) $\wedge \beta_1 \wedge \beta_2 \wedge \beta_3 \wedge \alpha_1$

Then
$O(\text{KB}_4) \Rightarrow B(cause(e_2, e_1) \vee e_1 < e_2)$
is a theorem.

That is, if the reader believes (4a) and the α and β rules alone, then he believes that either the pushing (e_2) caused the falling (e_1) or the falling preceded the pushing.

Proof 4

a.	$O(\text{KB}_4)$	assumption
b.	$B(\text{KB}_4)$	a; axiom (3)
c.	$B((4a) \wedge \beta_1 \wedge \beta_3)$	b; LO
d.	$B((\neg B\neg e_1 < e_2 \Rightarrow e_1 < e_2) \wedge$ $(\neg B\neg cause(e_2, e_1) \Rightarrow cause(e_2, e_1)))$	c; LO, punctuality of events
e.	$(\neg B\neg e_1 < e_2 \Rightarrow Be_1 < e_2)) \wedge$ $(\neg B\neg cause(e_2, e_1) \Rightarrow Bcause(e_2, e_1))$	d; LO
f.	$(B\neg e_1 < e_2 \vee Be_1 < e_2) \wedge (B\neg cause(e_2, e_1) \vee Bcause(e_2, e_1))$	e; PL
g.	$(B\neg e_1 < e_2 \wedge B\neg cause(e_2, e_1)) \vee (B\neg e_1 < e_2 \wedge Bcause(e_2, e_1)) \vee$ $(Be_1 < e_2 \wedge B\neg cause(e_2, e_1)) \vee (Be_1 < e_2 \wedge Bcause(e_2, e_1))$	f; PL
h.	$B\neg(e_1 < e_2 \vee cause(e_2, e_1)) \vee$ $(B\neg e_1 < e_2 \wedge Bcause(e_2, e_1)) \vee (Be_1 < e_2 \wedge B\neg cause(e_2, e_1))$	g; LO, PL, (21)

i.	$B\neg(e_1 < e_2 \vee cause(e_2, e_1)) \vee Bcause(e_2, e_1) \vee Be_1 < e_2$	h; LO, (21)
j.	$B\neg(e_1 < e_2 \vee cause(e_2, e_1)) \vee B(e_1 < e_2 \vee cause(e_2, e_1))$	h; LO
k.	$N\neg(\text{KB}_2)$	a; axiom (3)
l.	$N(((4a)\wedge\beta_2 \wedge \alpha_1) \Rightarrow \neg(\beta_1 \wedge \beta_3))$	k; LO
m.	$N(((4a)\wedge\beta_2 \wedge \alpha_1) \Rightarrow$ $((\exists e_1' e_2' xy)\neg(\beta_1^{ante} \Rightarrow \beta_1^{conseq}) \vee (\exists e_1' e_2')\neg(\beta_3^{ante} \Rightarrow \beta_3^{conseq})))$	l; LO
n.	$N(((4a)\wedge\beta_2 \wedge \alpha_1) \Rightarrow$ $((\exists e_1' e_2')\neg cause(e_2', e_1') \vee (\exists e_1' e_2')(e_1' \neq e_2 \wedge e_2' \neq e_1 \wedge \neg\beta_3^{conseq})))$	m; as before
o.	$\neg B(((4a)\wedge\beta_2 \wedge \alpha_1) \Rightarrow$ $((\exists e_1' e_2')\neg cause(e_2', e_1') \vee (\exists e_1' e_2')(e_1' \neq e_2 \wedge e_2' \neq e_1 \wedge \neg\beta_3^{conseq})))$	n; axiom(2)
p.	$\neg(B((4a)\wedge\beta_2 \wedge \alpha_1) \Rightarrow B((\exists e_1' e_2')\neg cause(e_2', e_1') \vee (\exists e_1' e_2')(e_1' \neq e_2 \wedge$ $e_2' \neq e_1 \wedge \neg\beta_3^{conseq})))$	o; LO o; LO
q.	$\neg B((\exists e_1' e_2')\neg cause(e_2', e_1') \vee (\exists e_1' e_2')(e_1' \neq e_2 \wedge e_2' \neq e_1 \wedge \neg\beta_3^{conseq}))$	p; PL
r.	$\neg B(\exists e_1' e_2')\neg cause(e_2', e_1')) \wedge \neg B(\exists e_1' e_2')(e_2' \neq e_1 \wedge e_1' \neq e_2 \wedge$ $\neg\beta_3^{conseq})$	q; LO, PL
s.	$\neg(\exists e_1' e_2')B\neg cause(e_2', e_1') \wedge \neg(\exists e_1' e_2')B(e_1' \neq e_2 \wedge e_2' \neq e_1 \wedge \neg\beta_3^{conseq})$	r; axiom (1c)
t.	$\neg B\neg cause(e_2, e_1) \wedge \neg B\neg(\exists t_1' t_2')(hold(e_1, t_1') \wedge hold(e_2, t_2') \wedge t_1' < t_2')$	s; PL, substitution for β_3^{conseq}
u.	$\neg B\neg cause(e_2, e_1) \wedge \neg B\neg e_1 < e_2$	t; LO, punctual events
v.	$\neg B\neg(cause(e_2, e_1) \vee e_1 < e_2)$	u; PL, LO
w.	$B(cause(e_2, e_1) \vee e_1 < e_2)$	j,v; PL
x.	$O(\text{KB}_4) \Rightarrow B(cause(e_2, e_1) \vee e_1 < e_2)$	a,w; PL
		Q.E.D.

What follows from the reader's belief in KB_4 alone is weaker than one might expect. Why can't the reader make a choice between the two exclusive disjuncts $cause(e_2, e_1)$ and $e_1 < e_2$? Note that neither $O\text{KB}_4 \Rightarrow Bcause(e_2, e_1)$ nor $O\text{KB}_4 \Rightarrow Be_1 < e_2$ are theorems (if one were a theorem then the other would be by a similar proof, but if they were both theorems, then $B(cause(e_2, e_1) \wedge e_1 < e_2)$ would follow, but this is impossible). The inability to choose between $cause(e_2, e_1)$ and $e_1 < e_2$ is due to the fact that the reader knows that the antecedents of two conflicting default rules are satisfied, and he has no way of assessing which default case is the most likely. This is as it should be, for (4) is ambiguous, just as (23) is.

(23) Max fell. John kicked him.

There are two natural interpretations for (23); either the falling (23a) precedes the kicking (23b), so that John kicks Max when he is still on the ground, or the kicking caused Max to fall. For both (4) and (23), the reader needs more information before he can decide which of the two default rules β_1 and β_3 applies.

But one wouldn't want the interpretation of (5) to be ambiguous in the same way, at least not if one is informed of the fact that death precludes poisoning (which is what α_1 is supposed to capture).

(5) Max died. John poisoned him.

(5a) $\langle dying'(e_1) \wedge subject(max', e_1) \wedge hold(e_1, t_1) \wedge t_1 < now,$
$poisoning'(e_2) \wedge subject(john', e_2) \wedge object(max', e_2) \wedge hold(e_2, t_2) \wedge t_2 < now\rangle$

This is predicted by our analysis:

Theorem 5

Let $\text{KB}_5 = (5\text{a}) \wedge \beta_2 \wedge \beta_3 \wedge \alpha_1$

Then
$O(\text{KB}_5) \Rightarrow Bcause(e_2, e_1)$
is a theorem.

That is, if the reader believes text (5) together with the knowledge described in the α and β rules, then he believes the poisoning (e_2) caused the death (e_1).

Proof 5

We first show that $O(\text{KB}_5) \Rightarrow B\neg e_1 < e_2$ is a theorem:

a.	$O(\text{KB}_5)$	assumption
b.	$B(\text{KB}_5)$	a; axiom (3)
c.	$B((5\text{a}) \wedge \alpha_1)$	b; LO
d.	$B(hold(e_1, t_1) \wedge hold(e_2, t_2) \wedge t_2 < t_1)$	c; LO
e.	$B(e_2 < e_1)$	d; LO and the punctuality of events
f.	$B\neg e_1 < e_2$	e; LO
		Q.E.D.

Now we will use this result to prove Theorem 5:

a.	$O(\text{KB}_5)$	assumption
b.	$B(cause(e_2, e_1) \vee e_1 < e_2)$	a; as proof 4
c.	$Bcause(e_2, e_1)$	b; LO, above lemma
d.	$O(\text{KB}_5) \Rightarrow Bcause(e_2, e_1)$	a,c; PL
		Q.E.D.

We have accounted for the natural interpretations of texts (1), (2), (4) and (5), not by a judicious statement of their truth conditions alone, but by providing a framework in which we can explore how their semantics interact with the reader's WK (both defeasible and indefeasible) and the pragmatic principles he uses for understanding language.

We will now extend the theory to deal with the other data we initially discussed, as this data has further implications for the relation betwen descriptive order and temporal order.

8 Extending the Theory

8.1 Preparatory Phases

Just as falling/pushing and dying/poisoning can be described in reverse order to their occurrences, so can building a bridge and drawing the plans:

(3) The council built the bridge. The architect drew up the plans.

And yet, there is no causal relation between drawing the plans and building the bridge, and so we cannot use the same strategy for explaining (3) as we used for texts (4) and (5), namely exploiting default rules about *causal* relations.

However, something very similar to causality can be invoked. The idea that certain events, such as building a bridge, have *preparatory phases* is now well established in the literature. The intuition is that events like building a bridge are preceded by a phase where events leading to the culmination occur; for example, drawing plans, mixing cement, laying foundations and so on (cf. Parsons' development portions). While these events are going on, one could truthfully say "The council is building the bridge". These events, that contribute to the culmination but do not strictly cause it, are what constitute the preparatory phase. Moens (1987) suggests that many temporal phenomena occurring in language can be more favourably explained than in current theories if the preparatory phase of an event is viewed as an irreducible concept in the ontology.

Intuitively, then, the explanation of the temporal structure of (3) is similar to that of (4) and (5), except for the fact that we use the notion of preparatory phases in place of causality. According to real WK, there is a default rule concerning the relation that typically holds between building a bridge and drawing the plans, and that is that the latter typically is in the preparatory phase of the former. More specifically, if one believes a bridge is built and plans were drawn (but one does not necessarily know for sure in which order these events occurred), then unless there is reason to believe something to the contrary, one believes the plans were part of the preparatory phase of the building of the bridge (and therefore preceded it).

We will represent this default rule in a similar manner to the above defeasible causal laws. This is done by using a two-place predicate *prep* which will take an eventuality and an event as its arguments, where $prep(e_1, e_2)$ means the eventuality e_1 is in the preparatory phase of the event e_2 and so contributes to the occurrence of e_2's culmination. We will treat the semantics of preparatory phases as primitive in this paper, as we are not concerned here with teasing out what criteria eventualities in preparatory phases satisfy. These criteria are explored in Lascarides (1988) in detail.[12]

A well known puzzle concerning preparatory phases is that some preparatory events occur even if the culmination to which they contribute does not. For example, someone can be crossing the street, but never achieve it because he gets hit by a truck. This means that $prep(e_1, e_2)$ should be satisfiable in those cases when e_2 does not occur.

The default rule concerning building a bridge and drawing the plans is expressed in β_4 below (where *iobject* stands for indirect object).

$$\beta_4 \quad (\forall e'_1 e'_2 x) \quad \begin{array}{l} ((building'(e_1) \wedge object(x, e'_1) \wedge (\exists i'_1) occ(e'_1, i'_1) \wedge \\ drawing'(e_2) \wedge object(plans', e'_2) \wedge iobject(x, e_2) \wedge (\exists i'_2) occ(e'_2, i'_2) \wedge \\ \neg B \neg prep(e'_2, e'_1)) \Rightarrow prep(e'_2, e'_1)) \end{array}$$

In words, β_4 states that if any building (e'_1) of an object x occurs and a drawing of plans (e'_2) of x also occurs, then e'_2 is in the preparatory phase of e'_1 unless one has reason to believe

[12] One could think of events in preparatory phases as those that *ceteris paribus* are insufficient but necessary for the culmination to occur. So one could potentially define preparatory phases in terms of default rules.

something to the contrary. We must add this WK about building bridges and drawing plans to the KB; it becomes an object of belief.

The logical form of (3) is (3a).

(3) The council built the bridge. The architect drew up the plans.

(3a) $\langle building'(e_1) \wedge subject(council', e_1) \wedge object(bridge', e_1) \wedge hold(e_1, t_1) \wedge t_1 < now, drawing'(e_2) \wedge subject(architect', e_2) \wedge object(plans', e_2) \wedge iobject(bridge', e_2) \wedge hold(e_2, t_2) \wedge t_2 < now\rangle$

Again, (3a) does not stipulate any temporal relations between the building of the bridge and the drawing of the plans. But nevertheless, one can prove the following theorem which shows that the reader has stronger beliefs than the semantics of (3) alone would warrant:

Theorem 3

Let $\text{KB}_3 = (3a) \wedge \alpha_1 \wedge \beta_1 \ldots \wedge \beta_4$

Then
$O(\text{KB}_3) \Rightarrow B(e_1 < e_2 \vee prep(e_2, e_1))$
is a theorem.

That is, if the reader's knowledge base is KB_3, then he believes that either the drawing of the plans (e_1) was part the preparatory phase of the building of the bridge (e_2) and so preceded it, or the building of the bridge preceded the drawing of the plans.

Proof 3

The proof is exactly analogous to Proof 4.

Again, the beliefs of the reader appear weaker than one might expect. Why can't the reader make a choice between the exclusive disjuncts $prep(e_2, e_1)$ and $e_1 < e_2$? The reasons are similar to those given for text (4); there are two 'conflicting' default rules β_3 and β_4 whose premises are satisfied by (3a), and so the reader needs further information if he is to choose which default case applies.

8.2 States

Describing a situation using an event expression or a state expression can affect the interpretation of text: the temporal structure of (24) is one of temporal precedence, and in (25) it is one of inclusion.

(24) Max arrived at the house. Mary stepped outside to greet him.

(25) Max arrived at the house. Mary was outside to greet him.

The problem is: how can the distinct semantics of events and states yield their distinct contributions to the interpretation of the temporal structure described in text?

The problem is more complicated than one might think, because the temporal relation between a state and other events as described in text depends on what those other events are: Why is the natural interpretation of text (26) one where the state of darkness includes the time when Max enters the room, but in (27) the state of darkness does *not* include the time when Max switches off the light?

(26) Max entered the room. It was pitch dark.

(27) Max switched off the light. It was pitch dark.

In text (27), we assume that the event causes the state, and so from this causal structure one can tease out via the law of Forward Causality the information that the event and state don't overlap (i.e. hold at a common time). But in text (26), we assume a different causal structure; entering a room does not typically cause darkness. And in interpreting (26), we assume the event and state overlap. So text (26) indicates that perhaps the following holds:

- *The Stative Principle*
  Contrary to other information, one assumes that the state overlaps previously described events.

Informally, according to the above Stative Principle, the event and state hold at a common time in (26) because there is no information to the contrary. In (27) on the other hand, if one believes the switching off the light caused the darkness, so by the law of Forward Causality one believes the event preceded the state, then one believes there is information contrary to event and state overlapping, and so the consequent of the Stative Principle does not apply for (27).

We have already represented defeasible linguistic knowledge (LK) in the framework in form of the rule of Narration. We now argue in favour of further defeasible LK: the Stative Principle.

We have not so far offered motivation for the Stative Principle that is independent of the fact that it potentially explains the above data. But there is independent motivation: we argue that the Stative Principle is a realisation of Grice's Maxim of Relevance. The argument goes as follows: In order to construct a full picture of the situation described by NL text, one must understand the relative occurrences of all the eventualities, including where they start and stop. In the case of events, which are punctual, one can fully determine the start and stop once one knows the point of time where it holds. These points of time are given directly in the semantic representation of NL text. But because *states* extend in time, knowing a *point* of time where the state holds is not sufficient for determining where it starts relative to other eventualities described in the text.

There are several linguistic mechanisms one can use to indicate where a state starts relative to the other eventualities. One could explicitly refer in the text to what caused the state, and so from the law of Forward Causality one knows the relative place where the state starts. This is what's going on in (27). Alternatively, one can use temporal adverbials to indicate where a state starts. But these two mechanisms are not enough: what about determining where states start in texts that do *not* feature adverbials or causes? The Stative Principle is another vital mechanism for determining where a state starts relative to other eventualities described in text. For the Stative Principle essentially states that if there is no 'explicit' indication of where the state starts, in the form of mentioning causes or temporal adverbials, then the start of the state is assumed to be irrelevant; i.e. the state started to hold *before* the situation that the text is concerned with occurs. Thus the Stative Principle can be viewed as a manifestation of Grice's Maxim of Relevance, for it asserts that unless there is indication in the text to the contrary, the point where a state starts is assumed to be irrelevant.

We represented Grice's Maxim of Orderliness as defeasible LK. Similarly, the Stative Principle, which we claim is closely related to Grice's Maxim of Relevance, is defeasible LK. It is represented as the following default rule:

Representation of the Stative Principle

$$\beta_5 \quad (\forall e_1' e_2') \quad ((\langle \phi_1(e_1'), \phi_2(e_2') \rangle \wedge occ(e_2', i_2) \wedge initial - bound(i_2) = t \wedge \neg hold(e_2', t) \wedge \neg B \neg overlap(e_1', e_2')) \Rightarrow overlap(e_1', e_2'))$$

The predicates *overlap* and *hold* are related in the following way, so that e_1 overlaps e_2 means that e_1 and e_2 hold at a common time.

(28) $overlap(e_1, e_2) \Rightarrow (\exists t)(hold(e_1, t) \wedge hold(e_2, t))$

β_5 is added to the reader's KB. It asserts that if e_1' is described before e_2', and the initial bound of the interval i_2 where e_2' occurs is not a point where e_2' holds (so e_2' must be a state), then unless there is reason to believe something to the contrary, e_1' and e_2' overlap. Note that we interpret the start of an eventuality as the initial bound of the interval where the (whole) eventuality occurs.

If e_2' is an event then it will not satisfy the antecedent of β_5 because e_2' would hold at the initial bound of the interval i_2 where (the whole of) e_2' occurs. So β_5 will play no part in determining where events start. Thus this new defeasible LK does not affect the semantics of texts (1), (2), (3), (4) and (5), as they describe events alone.

However, β_5 plays a central role in the reader's assessments of the temporal structures described by texts (25), (26) and (27).

(24) Max arrived at the house. Mary stepped outside to greet him.

(25) Max arrived at the house. Mary was outside to greet him.

(26) Max entered the room. It was pitch dark.

(27) Max switched off the light. It was pitch dark.

Let's look at (24) and (25) first, and our explanation of why the reader's beliefs about these texts differ. The logical forms of (24) and (25) are (24a) and (25a) respectively.

(24) Max arrived at the house. Mary stepped outside to greet him.

(24a) $\langle arriving'(e_1) \wedge subject(max', e_1) \wedge at(house', e_1) \wedge hold(e_1, t_1) \wedge t_1 < now,$
$stepping - outside'(e_2) \wedge subject(mary', e_2) \wedge hold(e_2, t_2) \wedge t_2 < now\rangle$

(25) Max arrived at the house. Mary was outside to greet him.

(25a) $\langle arriving'(e_1) \wedge subject(max', e_1) \wedge at(house', e_1) \wedge hold(e_1, t_1) \wedge t_1 < now,$
$be - outside'(e_2) \wedge subject(mary', e_2) \wedge hold(e_2, t_2) \wedge t_2 < now\rangle$

In words, (24a) is true at the index w if e_1 is an arriving at the house of Max in w that holds at a time t_1 in w that is earlier than now, and e_2 is a stepping outside (to greet Max) by Mary in w that holds at the time t_2 in w that is earlier than now. The truth conditions of (25a) are the same, save that e_2 becomes the *state* of being outside as opposed to the *event* of stepping outside.

Despite the fact that we have distinguished the semantics of events and states in terms of their different behaviours with the predicate *occ*, the above semantic representations of (24) and (25) on their own do not account for these texts' distinct natural interpretations. This is because no conditions are imposed on the temporal relation between e_1 and e_2.

We have claimed, however, that the reader's KB contains certain defeasible LK, namely the Principle of Narration (β_3) and the Stative Principle (β_5). This defeasible LK will yield different pragmatic entailments for the above two texts. That is, the reader's beliefs about the temporal relations for the above two texts will differ, in spite of the fact that they do not differ at the semantic level. Theorems 26 and 27 show that our theory predicts the natural interpretations of the two texts.

Theorem 26

Let KB_{26} =(24a)$\wedge\beta_3 \wedge \beta_5$.
Then
$O(\text{KB}_{26}) \Rightarrow B(e_1 < e_2)$
is a theorem.

That is, the reader believes that the arriving at the house *precedes* the stepping outside.

Proof 26

The proof is exactly like Proof 2.

Theorem 27

Let KB_{27} =(25a)$\wedge\beta_3 \wedge \beta_5$.
Then
$O(\text{KB}_{26}) \Rightarrow Boverlap(e_1, e_2)$
is a theorem.

That is, the reader believes that the arriving at the house and the being outside hold at a *common time.*

Proof 27

a.	$O(\mathrm{KB}_{27})$	assumption
b.	$B(\mathrm{KB}_{27})$	a; axiom (3)
c.	$B((25\mathrm{a}) \wedge \beta_5)$	b; LO
d.	$B(\neg B \neg overlap(e_1, e_2) \Rightarrow overlap(e_1, e_2))$	c; LO
e	$B \neg B \neg overlap(e_1, e_2) \Rightarrow Boverlap(e_1, e_2)$	d; axiom (1b)
f	$\neg B \neg B \neg overlap(e_1, e_2) \vee Boverlap(e_1, e_2)$	e; PL
g.	$B \neg overlap(e_1, e_2) \vee Boverlap(e_1, e_2)$	f; axiom (1d)
h.	$N \neg(\mathrm{KB}_{27})$	a; axiom (3)
i.	$N(((25\mathrm{a}) \wedge \beta_3) \Rightarrow \neg \beta_5)$	h; LO
j.	$N(((25\mathrm{a}) \wedge \beta_3) \Rightarrow (\exists e'_1 e'_2)(\beta_5^{ante} \wedge \neg \beta_5^{conseq}))$	i; LO
k.	$N(((25\mathrm{a}) \wedge \beta_3) \Rightarrow (\exists e'_1 e'_2) \neg overlap(e'_1, e'_2))$	j; LO, β_5^{conseq}
l.	$\neg B(((25\mathrm{a}) \wedge \beta_3) \Rightarrow (\exists e'_1 e'_2) \neg overlap(e'_1, e'_2))$	k; axiom(2)
m.	$\neg(B((25\mathrm{a}) \wedge \beta_3) \Rightarrow B(\exists e'_1 e'_2) \neg overlap(e'_1, e'_2))$	l; LO
n.	$B((25\mathrm{a}) \wedge \beta_3) \wedge \neg B(\exists e'_1 e'_2) \neg overlap(e'_1, e'_2)$	m; PL
o.	$\neg B(\exists e'_1 e'_2) \neg overlap(e'_1, e'_2)$	n; PL
p.	$(\forall e'_1 e'_2) \neg B \neg overlap(e'_1, e'_2)$	o; axiom (1c), PL
q.	$\neg B \neg overlap(e_1, e_2)$	p; PL
r.	$Boverlap(e_1, e_2)$	g,q; PL
s.	$O(\mathrm{KB}_{27}) \Rightarrow Boverlap(e_1, e_2)$	a,r; PL
		Q.E.D.

So, (24) and (25) are assigned similar semantics, but the reader has distinct beliefs about the temporal relations described in these texts because of his defeasible LK; the beliefs he has about the temporal relations are more specific than the semantics of these texts alone would warrant.

Now let's look at texts (26) and (27), and the puzzle that the temporal relation between and event and state described in text appears to depend on what the event is. The logical forms are (26a) and (27a) respectively.

(26) Max entered the room. It was pitch dark.

(26a) $\langle entering'(e'_1) \wedge subject(max', e_1) \wedge into'(room', e_1) \wedge hold(e_1, t_1) \wedge t_1 < now,$
$darkness'(e'_2) \wedge subject(room', e_2) \wedge hold(e_2, t_2) \wedge t_2 < now\rangle$

(27) Max switched off the light. It was pitch dark.

(27a) $\langle switching - off'(e'_1) \wedge subject(max', e_1) \wedge object'(light', e_1) \wedge hold(e_1, t_1) \wedge t_1 <$
$now,$
$darkness'(e'_2) \wedge subject(room', e_2) \wedge hold(e_2, t_2) \wedge t_2 < now\rangle$

(26) and (27) have similar logical forms, and therefore have similar semantic entailments. However, if the reader believes that the switching off of the light caused the darkness on the one hand, but does not believe that entering the room causes darkness on the other, then the following theorems apply:

Theorem 28

Let KB$_{28}$ =(26a)$\wedge\beta_3 \wedge \beta_5$.
Then
$O(\text{KB}_{28}) \Rightarrow Boverlap(e_1, e_2)$
is a theorem.

Proof 28

The proof is similar to Proof 27.

Theorem 29

Let KB$_{29}$ =(27a)$\wedge cause(e_1, e_2) \wedge \beta_3 \wedge \beta_5$.
Then
$O(\text{KB}_{29}) \Rightarrow B(e_1 < e_2)$
is a theorem.

Proof 29

$B(e_1 < e_2)$ follows directly from $Bcause(e_1, e_2)$ and the law of Forward Causality.

The above theorems show that the reader's beliefs concerning causal structure yield an explanation of why the temporal relation between the event and state described in text depends on what the event is; more specifically it depends on whether the reader thinks the event caused the state or not.

Now let's turn to the question of the awkwardness of text (8a,d); compared with acceptable (8a,c) and (8a,e).

(8)	a	Max crossed the finish line in first place.
	c	He arrived home with the cup.
	d	??He was home with the cup.
	e	A moment later, he was home with the cup.

The logical form of (8a,d) is (29).

(29) $$\langle crossing'(e_1') \wedge subject(max', e_1) \wedge object'(line', e_1) \wedge hold(e_1, t_1) \wedge t_1 < now,$$
$$being - home'(e_2) \wedge subject(max', e_2) \wedge hold(e_2, t_2) \wedge t_2 < now\rangle$$

(29) is satisfiable, so we are not explaining the awkwardness of (8a,d) by assigning it a representation that cannot be true. Instead, we shall explain the awkwardness of (8a,d) in terms of lack of information; (29) together with the reader's KB do not give rise to any beliefs about the temporal relations between e_1 and e_2.

The explanation is as follows: we add to the reader's KB Grice's Maxim of Quantity. It would be an unusal situation for the finish line to be at the winner's house. So according to the Maxim of Quantity, if this is what the narrator intends, then he must indicate this in some way. The narrator may indicate it by adding temporal adverbials (cf. (8a,e)), or he may explicitly refer to what caused Max to be at home (cf. (8a,f,d)):

(8) a Max crossed the finish line in first place.
f He opened his front door in great excitment.
d He was home with the cup.

It is unclear how a general version of the Maxim of Quantity would be represented in our theory. But we can represent the manifestation of the Maxim of Quantity for the particular situation of Max winning the race that we are discussing here. The Maxim of Quantity with respect to this situation yields the following:

- Contrary to other information, the reader believes that the finish line is in a different place from the winner's home.

This is represented in our theory as defeasible linguistic knowledge, given in β_6:

$$\begin{array}{lll} \beta_6 & (\forall e'_1 e'_2 x) & (crossing(e'_1) \wedge subject(x, e'_1) \wedge object(finishline(e'_1) \wedge \\ & & beathome(e'_2) \wedge subject(x, e'_2) \wedge \\ & & \neg Boverlap(e'_1, e'_2) \Rightarrow \neg overlap(e'_1, e'_2) \end{array}$$

β_6 stipulates that if there is a crossing of the finish line by an individual x and a being at home state of the individual x then if nothing to the contrary is believed, the event and state do not overlap.

Suppose we add β_6 to the KB, so the reader believes the following:

KB$_6$ (29) $\wedge\beta_3 \wedge \beta_5 \wedge \beta_6$

Then one can show that not only is (29) satisfiable, but it is consistent to believe KB$_6$. However, the awkwardness of (8a,d) is explained by the fact that KB$_6$ is uninformative, in that the reader cannot glean anything about the relation between the eventualities described, in spite of the fact that he has all sorts of default rules that are there to help him tease out the temporal structure of text. The reason for this is that the reader has two conflicting default rules that apply to (29)—β_5 (*ceteris paribus* the event and state overlap) and β_6 (*ceteris paribus* the event and state do not overlap)— and no way of deciding which default case is the most likely. So the reader infers from KB$_6$ that $B(overlap(e_1, e_2) \vee \neg overlap(e_1, e_2))$; the proof is similar to proof 4. But this is true by axiom (1a) of the logic of belief anyway. And there are no other more 'specific' beliefs about the temporal relation between e_1 and e_1 that follow from KB$_6$. So our explanation of the awkwardness of (8a,d) is that it does not contain enough information for any temporal structure to be revealed even if all WK and LK are brought to bear. So to summarise, (8a,d) is not awkward due to inconsistency. Instead, it is awkward due to uninformativeness.

We showed earlier that the reader believed a disjunction of temporal relations for text (4) that were mutually exclusive.

(4) Max fell. John pushed him.

Namely, (30) followed from the reader's KB.

(30) $B(cause(e_2, e_1 \vee e_1 < e_2)$

But the reader cannot decide between the two exclusive disjuncts. This may appear to be similar to the case for (29), because in (29), once again, the reader cannot decide between two exclusive disjuncts. But there is a big difference; the object of belief in (30) is not valid, while the object of belief following from (29) is valid. The difference in informativeness between texts (4) and (8a,d) is as big as the difference in informativeness between (31) and (32).

(31) The block is blue or it is green.

(32) The block is blue or it is not blue.

8.3 Temporal Adverbials

We have offered an explanation of the awkwardness of (8a,d) by claiming that it does not offer enough information so that the reader can construct the temporal structure of the eventualities involved. This yields a puzzle: how does the temporal adverbial in (8e) invoke enough information so that there is an improvement in interpretability between (8a,d) and (8a,e)?

(8) a Max crossed the finish line in first place.
d ??He was home with the cup.
e A moment later, he was home with the cup.

In order to explain why (8a,e) is better than (8a,d), we would have to tell a story about how the Maxim of Quantity is not violated in cases when adverbials are present. In other words, we must allow the presence of certain temporal adverbials to entail that the defeasible LK β_6 does not apply to the text.

To tell such a story would involve adding more linguistic knowledge to the knowledge base. We would have to add rules that empahsise the role of adverbials such as *two hours later*, *yesterday* and *when Max entered the room* as not only helping one to identify when the eventuality being described occurred, but also indicating that there is a 'lapse in time' between this eventuality and the ones described previously in the text and the author does not wish to describe events that occur in this lapse of time. For example, (8a,e) asserts that there is a lapse in time of a moment between the event of crossing the finish line and the state of being at home, and one knows from WK that Max must have gone home in that moment, even though the author doesn't state this. So parts of the situation are left undescribed.

We would have to add rules to the reader's existing LK in our theory that essentially amount to the following:

- Adverbials allow the author to leave certain events unmentioned, even if their existence is crucial to the understanding of the text.

If the above were added, then we would have an account of how the adverbial in (8a,e) indicates that the event and state don't overlap. It is unclear how the above could be expressed in the formal theory though. The question of how one might formally represent temporal adverbials deserves considerable attention in its own right, and is beyond the scope of this paper.

9 Conclusion and New Problems

In this paper we have addressed the Ordering Question: Given a particular order in which events are described in text, what are the constraints in interpretation on their temporal and causal relations in the world?

We have developed a semantic account of text where we have shown that many of the problems associated with the Ordering Question can be solved if the logical framework is one where we not only have an account of the text's truth conditions, but also have an account of how the reader's knowledge about the world and use of language contributes to his interpretation of text. In other words, we also have the means within the logic to derive the reader's beliefs about the text given his existing KB.

This theory supplies a logical account of how WK, LK and semantic representation interact to yield the temporal structure of text. The strategy of bringing both the logical form of NL expressions *and* WK and LK to bear in interpreting of NL fragments could have wider applicability. Perhaps the similarities between temporal and nominal reference (cf. Partee (1984)) would enable the same technique to be used in the nominal case.

However, the constraints on the relation between descriptive order and temporal order that we have provided are incomplete. First, we have failed to address the problem of how non-atomic syntax such as the past perfect and progressive affect the constraints. Second, our theory as it currently stands predicts that contrary to intuitions, *both* the texts (33) and (34) are acceptable to the reader, assuming that he knows that death precludes poisoning, and darkness precludes switching off the light.

(33) It was pitch dark. Max switched off the light.

(34) Max was dead. John poisoned him.

We predict they have similar natural interpretations because both (33) and (34) have similar logical forms *and* the reader has similar WK about the events described (i.e. the first mentioned event precludes the second). But in fact, whereas (34) is acceptable, (33) is awkward. Thus (33) can be considered to be a counterexample to our answer to the Ordering Question.

How could one rectify this? The difference between the acceptability of (33) and (34) intuitively arises from a difference between their *rhetorical structures* (cf. Mann and Thompson (1987), Hovy (1987)). In (34) the narrator is creating drama; he is teasing the reader by leaving the crucial event of the story until last. Indeed, in a crime story, there may be several pages between the first event described in (34) and the second. It is more difficult to find a context where Max's switching off the light is as significant to the story being told as John's poisoning Max. So the same rhetorical ploy isn't used in (33). Consequently, it seems that not only is knowledge about the typical temporal and causal relations between the events described crucial to the constraints on their descriptive order, but also rhetoric plays a non-trivial role. It is the combination, then, of WK, LK and rhetorical knowledge (RK) that will yield a full answer to the Ordering Question.

Note that RK alone is not sufficient for answering the Ordering Question, for even though the rhetorical ploy used in (35) is the same as that used in (34), (35) is not acceptable if the situation being described is intended to be one where the second event described causes the first. Intuitively, this is because the (lack of) WK about the typical relations between the events prevents its accpetability as it did in (2).

(2) Max turned round. John hit him.

(35) Max turned round. (Just imagine!) John hit him.

In this paper, we have dealt with only two of the three types of reader's knowledge that yield a full answer to the Ordering Question; namely WK and LK. In a subsequent paper, the problem of how rhetorical structure affects the interpretation of temporal reference will be addressed.

Appendix

The Truth Definition

The well-formed expressions of the language are defined with respect to a pair $\langle M, g\rangle$ where M is a model for the language (i.e. a 9-tuple $\langle W, D, T, I, <_T, <_I, \Re, f\rangle$ as we defined in the paper), and g is a variable assignment function.

(a) Where β is a variable, $[\![\beta]\!]^{\langle M,g\rangle}(w) = g(\beta)$.

(b) Where β is either a name constant or a predicate constant, $[\![\beta]\!]^{\langle M,g\rangle}(w) = f(\beta)(w)$.

(c) Where β is an atomic wff $p^n(d_1, \ldots, p_n)$,
$[\![\beta]\!]^{\langle M,g\rangle}(w) = 1$ iff $\langle [\![d_1]\!]^{\langle M,g\rangle}(w), \ldots, [\![d_n]\!]^{\langle M,g\rangle}(w)\rangle$ belongs to $[\![p^n]\!]^{\langle M,g\rangle}(w)$.

(d) Where β is a wff $t_1 < t_2$,
$[\![\beta]\!]^{\langle M,g\rangle}(w) = 1$ iff $[\![t_1]\!]^{\langle M,g\rangle}(w) <_T [\![t_1]\!]^{\langle M,g\rangle}(w)$.
Where β is a wff $i_1 < i_2$,
$[\![\beta]\!]^{\langle M,g\rangle}(w) = 1$ iff $[\![i_1]\!]^{\langle M,g\rangle}(w) <_I [\![i_1]\!]^{\langle M,g\rangle}(w)$.

(e) Where β is a wff $\neg A$,
$[\![\beta]\!]^{\langle M,g\rangle}(w) = 1$ iff $[\![A]\!]^{\langle M,g\rangle}(w) = 0$.

(f) Where β is a wff $A \wedge B$,
$[\![\beta]\!]^{\langle M,g\rangle}(w) = 1$ iff $[\![A]\!]^{\langle M,g\rangle}(w) = 1$ and $[\![B]\!]^{\langle M,g\rangle}(w) = 1$.

(g) Where β is a wff $A \vee B$,
$[\![\beta]\!]^{\langle M,g\rangle}(w) = 1$ iff $[\![A]\!]^{\langle M,g\rangle}(w) = 1$ or $[\![B]\!]^{\langle M,g\rangle}(w) = 1$.

(h) Where β is a wff $A \Rightarrow B$,
$[\![\beta]\!]^{\langle M,g\rangle}(w) = 1$ iff $[\![A]\!]^{\langle M,g\rangle}(w) = 0$ or $[\![B]\!]^{\langle M,g\rangle}(w) = 1$.

(i) Where β is a wff $\exists xA$ with the individual variable x,
$[\![\beta]\!]^{\langle M,g\rangle}(w) = 1$ iff $[\![A]\!]^{\langle M,g(x,e)\rangle}(w) = 1$ for some e belonging to D.
The truth definition for $\exists xA$ for other sorts of variables x is similar.

(j) Where β is a wff $\forall xA$ with the individual variable x,
$[\![\beta]\!]^{\langle M,g\rangle}(w) = 1$ iff $[\![A]\!]^{\langle M,g(x,e)\rangle}(w) = 1$ for all e belonging to D.
The truth definition for $\forall xA$ for other sorts of variables x is similar.

(k) Where β is a wff $B\phi$,
$[\![B\phi]\!]^{\langle M,g\rangle}(w) = 1$ iff for every w' such that $w' \approx w$ and $w' \in \Re(w)$, $[\![\phi]\!]^{\langle M,g\rangle}(w) = 1$.

(l) Where β is a wff $N\phi$,
$[\![N\phi]\!]^{\langle M,g\rangle}(w) = 1$ iff for every w' such that $w' \approx w$ and $w' \in \overline{\Re}(w)$, $[\![\phi]\!]^{\langle M,g\rangle}(w) = 1$.

(m) Where β is a wff $O\phi$,
$[\![O\phi]\!]^{\langle M,g\rangle}(w) = 1$ iff for every w', $w' \approx w$ and $w' \in \Re(w)$ if and only if $[\![\phi]\!]^{\langle M,g\rangle}(w) = 1$.

Axioms

(1) The axioms for weak S5, for both B and N, which are the following:
- (a) $L\phi$, where ϕ is any valid objective sentence,
- (b) $L(\phi \Rightarrow \psi) \Rightarrow (L\phi \Rightarrow L\psi)$
- (c) $(\forall x)L\phi \Rightarrow L(\forall x)\phi$
- (d) $\sigma \Rightarrow L\sigma$, where σ is subjective
  (one can in fact show that $BB\phi \Rightarrow B\phi$ is also a theorem),

Where L is either B or N.

(2) $N\phi \Rightarrow \neg B\phi$, where ϕ is any objective sentence that is falsifiable.

(3) $O\phi \Leftrightarrow (B\phi \wedge N\neg\phi)$, for any ϕ.

(21) Law of Forward Causality
$(\forall e_1 e_2)(cause(e_1, e_2) \Rightarrow e_1 < e_2)$

Defeasible World Knowledge

Typically, pushing causes falling:

$$\beta_1 \quad (\forall e'_1 e'_2 xy) \quad ((falling'(e'_1) \wedge subject(x, e'_1) \wedge (\exists i'_1)occ(e'_1, i'_1) \wedge$$
$$pushing'(e'_2) \wedge subject(y, e'_2) \wedge object(x, e'_2) \wedge (\exists i'_2)occ(e'_2, i'_2) \wedge$$
$$\neg B \neg cause(e'_2, e'_1)) \Rightarrow cause(e'_2, e'_1))$$

Typically, poisoning causes death:

$$\beta_2 \quad (\forall e'_1 e'_2 xy) \quad ((dying'(e'_1) \wedge subject(x, e'_1) \wedge (\exists i'_1)occ(e'_1, i'_1) \wedge$$
$$poisoning'(e'_2) \wedge subject(y, e'_2) \wedge object(x, e'_2) \wedge (\exists i'_2)occ(e'_2, i'_2) \wedge$$
$$\neg B \neg cause(e'_2, e'_1)) \Rightarrow cause(e'_2, e'_1))$$

Typically, drawing the plans is in the preparatory phase of building the bridge:

$$\beta_4 \quad (\forall e'_1 e'_2 x) \quad ((building'(e_1) \wedge object(x, e'_1) \wedge (\exists i'_1)occ(e'_1, i'_1) \wedge$$
$$drawing'(e_2) \wedge object(plans', e'_2) \wedge iobject(x, e_2) \wedge (\exists i'_2)occ(e'_2, i'_2) \wedge$$
$$\neg B \neg prep(e'_2, e'_1)) \Rightarrow prep(e'_2, e'_1))$$

Defeasible Linguistic Knowledge

The Narrative Principle:

$$\beta_3 \quad (\forall e'_1 e'_2) \quad (\langle \phi_1(e'_1), \phi_2(e'_2) \rangle \wedge \neg B \neg (\exists t'_1 t'_2)(hold(e'_1, t'_1) \wedge hold(e'_2, t'_2) \wedge bt'_1 < t'_2)) \Rightarrow$$
$$(\exists t'_1 t'_2)(hold(e'_1, t'_1) \wedge hold(e'_2, t'_2) \wedge t'_1 < t'_2)$$

where ϕ_1 and ϕ_2 are λ-expressions of the form $\lambda e \psi$.

The Stative Principle:

$$\beta_5 \quad (\forall e'_1 e'_2 t'_1 t'_2) \quad (((\langle \phi_1(e'_1), \phi_2(e'_2) \rangle \wedge occ(e'_2, i_2) \wedge initial - bound(i_2) = t \wedge \neg hold(e'_2, t) \wedge$$
$$\neg B \neg overlap(e'_1, e'_2)) \Rightarrow overlap(e'_1, e'_2))$$

The 'Maxim of Quantity':

$$\beta_6 \quad (\forall e'_1 e'_2 x) \quad (crossing(e'_1) \wedge subject(x, e'_1) \wedge object(finishline(e'_1) \wedge$$
$$beathome(e'_2) \wedge subject(x, e'_2) \wedge$$
$$\neg Boverlap(e'_1, e'_2) \Rightarrow \neg overlap(e'_1, e'_2)$$

Indefeasible Knowledge

Death precludes Poisoning:

$$\alpha_1 \quad (\forall e'_1 x t'_1 e'_2 t'_2) \quad \begin{array}{l} (dying'(e'_1) \wedge subject(x, e'_1) \wedge hold(e'_1, t'_1) \wedge \\ poisoning'(e'_2) \wedge object(x, e'_2) \wedge hold(e'_2, t'_2) \Rightarrow t'_2 < t'_1) \end{array}$$

References

Davidson, D. [1967a] Causal Relations. *Journal of Philosophy*, **64**. Reprinted in Davidson (1980) pp149-162.

Davidson, D. [1967b] The logical form of action sentences. In Rescher, N. (ed.) *The Logic of Decision and Action.* Pittsburgh: University of Pittsburgh Press.

Dowty, D. [1986] The effects of aspectual class on the temporal structure of discourse: semantics or pragmatics? *Linguistics and Philosophy*, **9**, 37-61.

Grice, H. P. [1975] Logic and Conversation. In Cole, P. and Morgan, J. L. (eds.) *Syntax and Semantics*, Volume 3: *Speech Acts*, pp41-58. New York: Academic Press.

Hinrichs, E. [1986] Temporal anaphora in discourses of English. *Linguistics and Philosophy*, **9**, 63-82.

Hintikka, J. [1962] *Knowledge and Belief.* Ithaca, N.Y.: Cornell University Press.

Hovy, E. H. [1987] Some pragmatic decision criteria in generation. Chapter 1 in Kempen, G. (ed.) *Natural Language Generation*, pp3-18. Dordrecht: Martinus Nijhoff.

Joshi, A. and Webber, B. [1984] Default reasoning in interaction. In *Proceedings of the Non-Monotonic Reasoning Workshop*, AAAI, New York, October, 1984, pp144-150.

Kamp, H. and Rohrer, C. [1983] Tense in texts. In Bauerle, R., Schwarze, C. and von Stechow, A. (eds.) *Meaning, Use and Interpretation of Language*, pp250-269. Berlin: de Gruyter.

Konolige, K. [1987] On the Relation Between Default and Autoepistemic Logic. Report No. CSLI-87-105, Center for the Study of Language and Information, September, 1987.

Kripke, S. [1972] Naming and Necessity. In Davidson, D. and Harman, G. (eds.) *Semantics of Natural language*, pp253-355. Dordrecht: D. Reidel.

Lascarides, A. [1988] A Formal Semantic Analysis of the Progressive. PhD Thesis, Centre for Cognitive Science, University of Edinburgh.

Levesque, H. J. [1985] All I know: An Abridged Report. In *Readings in Knowledge Representation.* Los Altos, Ca.: Morgan Kaufmann Publishers, Inc..

Mann, W. C. and Thompson, S. A. [1987] Rhetorical Structure Theory: A Framework for the Analysis of Texts. Technical Report No. RS-87-185, USC Information Sciences Institute, Marina Del Rey, Ca., April, 1987.

McDermott, D. and Doyle, J. [1980] Nonmonotonic logic I. *Artificial Intelligence*, **13**, 41-72.

Moens, M. [1987] Tense, Aspect and Temporal Reference. PhD Thesis, Centre for Cognitive Science, University of Edinburgh.

Moore, R. C. [1985] A Formal Theory of Knowledge and Action. Report No. CSLI-85-31, Center for the Study of Language and Information, September 1985, 1985.

Parsons, T. [1989] The Progressive in English: Events, States and Processes. *Linguistics and Philosophy*, **12**, 213 - 241.

Partee, B. H. [1984] Nominal and temporal anaphora. *Linguistics and Philosophy*, **7**, 243-286.

Reiter, R. [1980] A Logic for Default Reasoning. *Artificial Intelligence*, **13**, 81-132.

Sidner, C. L. [1983] Focusing in the comprehension of definite anaphora. In Berwick, R. C. and Brady, M. (eds.) *Computational Models of Discourse*, pp267-330. Cambridge, Mass.: MIT Press.

Webber, B. L. [1987] The Interpretation of Tense in Discourse. In *Proceedings of the 25th Annual Meeting of the Association for Computational Linguistics*, Stanford University, Stanford, Ca., July, 1987, pp147-154.