Technical Report

UCAM-CL-TR-325
ISSN 1476-2986

Number 325

UNIVERSITY OF CAMBRIDGE

Computer Laboratory

The temporal properties of English conditionals and modals

Richard Crouch

January 1994

15 JJ Thomson Avenue
Cambridge CB3 0FD
United Kingdom
phone +44 1223 763500

http://www.cl.cam.ac.uk/

This technical report is based on a dissertation submitted April 1993 by the author for the degree of Doctor of Philosophy to the University of Cambridge.

Technical reports published by the University of Cambridge Computer Laboratory are freely available via the Internet:

http://www.cl.cam.ac.uk/techreports/

ISSN 1476-2986

Abstract

The Temporal Properties of English Conditionals and Modals

Richard Crouch

This thesis deals with the patterns of temporal reference exhibited by conditional and modal sentences in English, and specifically with the way that past and present tenses can undergo deictic shift in these contexts. This shifting behaviour has consequences both for the semantics of tense and for the semantics of conditionals and modality.

Asymmetries in the behaviour of the past and present tenses under deictic shift are explained by positing a primary and secondary deictic centre for tenses. The two deictic centres, the assertion time and the verification time, are given independent motivation through an information based view of tense. This holds that the tense system not only serves to describe the way that the world changes over time, but also the way that information about the world changes. Information change takes place in two stages. First, it is asserted that some fact holds. And then, either at the same time or later, it is verified that this assertion is correct.

Typically, assertion and verification occur simultaneously, and most sentences convey verified information. Modals and conditionals allow delayed assertion and verification. *If A, then B* means roughly: suppose you were now to assert A; if and when A is verified, you will be in a position to assert B, and in due course this assertion will also be verified. Since A and B will both be tensed clauses, the shifting of the primary and secondary deictic centres leads to shifted interpretations of the two clauses.

The thesis presents a range of temporal properties of indicative and subjunctive conditionals that have not previously been discussed, and shows how they can be explained. A logic is presented for indicative conditionals, based around an extension of intuitionistic logic to allow for both verified and unverified assertions. This logic naturally gives rise to three forms of epistemic modality, corresponding to *must*, *may* and *will*.

Declaration

This report differs in a small number of details from the PhD thesis lodged with the Cambridge University Library. Besides formatting differences, it incorporates some useful observations made by my examiners, Prof. Frank Veltman and Sir John Lyons, for which I thank them.

This dissertation is the result of my own work and includes nothing which is the outcome of work done in collaboration. I declare that the thesis is not substantially the same as any that I have submitted for a degree of diploma or other qualification at any other university. I further state that no part of my thesis has already been or is being concurrently submitted for any such degree, diploma or other qualification.

Acknowledgements

The acknowledgements are invariably the most interesting part of any thesis, especially if you are mentioned in them. There are a number of people whose readership I would like to attract by thanking them.

Without Steve Pulman this thesis would never have been started, though I am sure there were times when we both wondered if it would ever be finished. I would like to thank him for his patience and tolerance and for his good advice as a supervisor. Jan van Eijck also had a hand in supervising me. I would like to thank him for some enjoyable evenings discussing logic along with Ben Macias, Victor Poznanski and David Milward, and more recently for putting me up on a number of trips to Amsterdam.

Thanks must also go to the SRI Cambridge Computer Science Research Centre, whose laser printer and photocopier I have greatly abused. If the gratitude is ambivalent, it is only because since I started working there I have been given too many other interesting things to do. Without them, I might have finished this thesis sooner, but I would have been the poorer. Special mention should go to Hiyan Alshawi, who always managed to come up with better things to think about than tense and conditionals. Ian Lewin made many helpful comments on a draft of this thesis. Further thanks go to David Carter for making a number of corpora available on-line, to Arnold Smith, Manny Rayner and once again to Steve Pulman.

Thanks also to the people who worked with me on the Alvey natural language planner interface project, especially Nigel Pirie, Dave Moffat, Barry Crabtree and Graeme Ritchie. The project suggested a topic of research, but at the same time gave me precious little opportunity to do much about it.

It is invidious to single out people, past and present, from the natural language group at the Computer Laboratory, so here it goes. Ann Copestake has always been a source of sound advice, about formal semantics and home maintenance equally. Ben Macias has been a fount of heterodox and stimulating ideas, and also made helpful comments on a draft of this thesis. Derek Bridge and Victor Poznanski were a comedy double act that made tea-time worthwhile. Mark Hepple, the Alexei Sayle of computational linguistics, was always on for a drink, as were Antonio Sanfillipo and John Carroll (before fatherhood beckoned). Ben Johnson-Laird helped undo some of the effects of this in forlorn attempts at exercise along the banks of the River Cam. Ted Briscoe's unfailing sense of cynicism was occasionally a source of discomfort, but more often a source of consolation. Other people from Cambridge who would have a right to feel aggrieved if they were not mentioned include: Mayumi Masuko, Thomas Forster, Arturo Trujillo, Graham Russell, David Elworthy, Karen Sparck Jones and Julia Galliers, though the list is not exhaustive.

From further afield, thanks to Per Hasle and Peter Øhrstrøm for two invitations to

Denmark, and to Alex Lascarides and Jon Oberlander for encouragement and letting me steal one of their examples.

By far my greatest debt of gratitude belongs to Valeria de Paiva. Without her constant emotional and intellectual support, I don't know where I would have been. I cannot thank you enough.

Contents

Chapter 1

Introduction

This thesis deals with the patterns of temporal reference exhibited by conditional and modal sentences in English. More specifically, it is concerned with the way that the past and present tenses behave in these sentences. This behaviour turns out to have consequences both for the semantics of tense, and for the semantics of conditionals and modals.

As an example of what is at issue, consider the following sentence, which has two quite distinct temporal interpretations:

(1) If the bimetallic strip bent, then the temperature rose.

Recall that a bimetallic strip in a thermostat bends one way when the temperature falls, completes an electrical connection and turns on a heater. It bends the other way when the temperature rises, breaks the connection and turns the heater off.

The first reading of (1) has the antecedent event of the strip bending preceding / causing the consequent event of the temperature rising; the strip bent and turned the heater on. The second reading reverses this order; the fact that the bimetallic strip bent is used as evidence for the fact that the temperature rose beforehand, causing the strip to bend. The first reading is predictive, reasoning forward from causes to effects. The second is explanatory, reasoning backwards from effects to causes.

If the past tenses in (1) are changed to present tenses, as in (2)

(2) If the bimetallic strip bends, then the temperature rises.

the second, explanatory reading is no longer available. Only a predictive, antecedent precedes/causes consequent reading is possible for (2).

The contrast between (1) and (2) can only be attributed to the different ways in which past and present tenses semantically interact with conditionals. Pragmatic factors, like general world knowledge, can of course affect the temporal interpretations of conditionals, ruling out some semantically possible readings as pragmatically implausible[1]. But if world knowledge allows (1) to be read in either order, the same world knowledge also applies to (2). The absence of the second explanatory reading for (2) therefore cannot due to pragmatic factors like world knowledge.

This can be illustrated even more strongly by the following pair of sentences

(3) If a person was buried in an unmarked grave, they usually died in poverty.

(4) ?If a person is buried in an unmarked grave, they usually die in poverty.

In the past tense (3) the consequent is construed as an explanation for the antecedent, giving a reading where people die first and are buried after. But (4) disallows this reversed reading, giving only an implausible interpretation (unless you are Edgar Allen Poe) where being buried alive leads to penury and then death. That is, the semantic interactions between the present tense and the conditional in (4) is sufficient to block the only pragmatically plausible reading.

The contrast between (1) and (2) and between (3) and (4) is a special instance of the more general phenomenon whereby tenses can undergo *deictic shift* in modal and conditional contexts. Deictic shift occurs when a tense locates an event as being past or present with respect to some time other than the speech time. Often this results in past and present tenses that refer to times in the future. More clear-cut examples of modal and conditional deictic shift are:

(5) If I smile when I get out, the interview went well.

(6) By 1998, everybody will know someone who died of AIDS.

In (6), reference is made via the past tense to people who died before 1998, and not just to people who have already died. In (5), the past tense consequent refers to an interview that is yet to take place, but which will precede the time at which I get out of it and (perhaps) smile[2]. Perhaps less immediately obvious, but no less important, is that fact that in (2) the present tense antecedent also has future time reference, and that the consequent tense is (to a first approximation), centred on the time at which the antecedent event occurs.

Modal and conditional contexts are not the only ones in which tenses can undergo deictic shift. However, they form a central class of shifting context. Deictic shift in other circumstances can be understood by analogy to what happens with modals and

[1] In other conditionals formally similar to the past tense (1), world knowledge can rule out one or other of the readings. For example, *If the house burned down, then something set it alight* only allows a reversed, explanatory reading; while *If someone dropped the ball, it fell* only allows a predictive reading, where antecedent precedes consequent. (1) is unusual in that the nature of bimetallic strips and thermostats pragmatically allows the conditional to be read in either direction.

[2] *If I smile when I get out, the interview will have gone well* is perhaps a preferred way of expressing the content of (5), but (5) is nevertheless an acceptable sentence.

conditionals.

Despite widespread interest in the semantics of modals and conditionals on the one hand, and tense and temporal reference on the other, there has been remarkably little work on the interactions between the two (Dudman (1983, 1991), Thomason and Gupta (1981) and Thomason (1984) are notable exceptions). What work there has been has, by and large, not paid close attention to, or been systematic in, its description of the actual patterns of temporal reference exhibited by modal and conditional sentences. The primary goals of this thesis are therefore twofold. The first is to provide a fuller description of the range of temporal interactions between tense, conditionals and modals than has hitherto been given. The second is to develop a semantic treatment of tenses, modals and conditionals that explains this range of data, accounting for such things as the contrast between (1) and (2).

In this chapter I will introduce some of the basic ideas to be developed in the rest of this thesis. These include the following. (i) The past and present tenses have both a primary and a secondary deictic centre. The existence of two deictic centres accounts for certain kinds of asymmetry in the behaviour of past and present tenses in deictically shifting contexts. (ii) The two deictic centres correspond to times at which informational operations of *assertion* and *verification* take place. The nature of these two operations is discussed below. The tenses not only serve to describe the way that the world changes over time, but also the way that information about the world changes. (iii) Modals and conditionals legislate for what happens with the future assertion and verification of information (even if the information concerns events that happened in the past). It is the futurity of assertion and verification that leads to deictic shift in modal and conditional contexts.

1.1 Tense, Conditionals and Information Change

There is currently a popular move away from a bare truth-conditional account of meaning towards a treatment that defines meaning as the potential to change states of information. According to Veltman (1990),

> ...the slogan "You know the meaning of a sentence if you know the conditions under which it is true" should be replaced by ... "You know the meaning of a sentence if you know the *change* it brings about in the information state of anyone who wants to incorporate the piece of news conveyed by it."

This change in perspective raises an interesting possibility with respect to the semantics of tense. On a truth-conditional account, linguistic devices for temporal reference are ultimately grounded in the way that they describe how the world changes over time. But if meaning is to be explicated in terms of information change, there is a second level at which temporal reference might additionally operate: namely, describing or constraining the way that information changes over time.

In particular, one might revise the standard account of how tenses are deictically centred. Typically, the tenses are taken to state a relation between the time at which some utterance event occurs (the speech time) and the time at which the event being described occurs (the event time). The speech time thus serves as the (single) deictic centre for the tenses. An new alternative is to centre tenses on the time at which an update is made to one's stock of information, where this update occurs as the result of the utterance of the sentence. In other words, some kind of update time provides the tense's deictic centre, rather than the speech time.

In a most cases, the move from speech time to update time will make no discernible difference: normally, the update occurs as soon as the utterance is made. However, modal and conditional sentences do not behave like this: they place constraints on the way that updates may be made in the future, as will be seen below. Tenses within the scope of the modal or conditional are consequently centred relative to these future update times, leading to deictic shift.

As will soon become apparent, matters do not work out properly if update is taken to be a single operation, occurring at a single time. Instead, it is necessary to decompose update into two operations: *assertion* and *verification.* Making an assertion adds a piece of information to one's information state. However, the assertion does not enjoy first class status until it becomes verified.

The difference between a verified and an unverified assertion can be illustrated by reference to the conditional (1) (the past tense bimetallic strip conditional). Suppose it is discovered at some point that the bimetallic strip did indeed bend. The conditional permits one to assert that the temperature rose. However, one has no direct evidence for this conclusion. The conditional allows one to assert the consequent, but the assertion does not yet count as verified. It is only once further investigation reveals that the temperature did indeed rise that the consequent assertion is verified.

A modal like *will* also has the effect of making unverified assertions. If I hear a sound at the door and say *That will be the postman*, I am asserting that the postman is at the door but conceding that until, say, I go to the door and pick up the letters on the doormat, I have no direct evidence to verify this assertion[3].

[3]The distinction between direct and indirect information will be known to anyone familiar with Veltman's Data Semantics (Veltman 1984, 1985; Landman 1985; see Chapter 2). However, this is not quite the same as the difference between verified and unverified information. In Data Semantics, indirect information does not hold in one's current information state, but will inevitably come to hold in some later state. The treatment of the conditional in Data Semantics says essentially that if the antecedent comes to hold in a given state, that state will inevitably grow into a later one in which the consequent holds. By

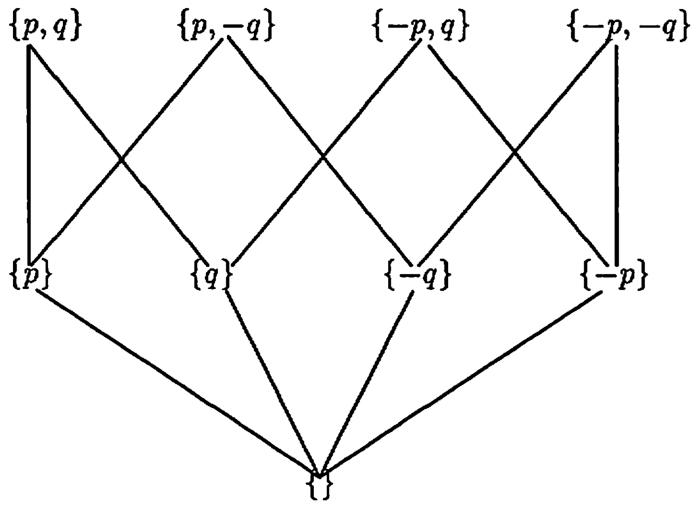

Figure 1.1: Ordered Set of Information States

To see why update needs to be decomposed into the operations of assertion and verification, we first need to consider two different but closely related approaches to the way that conditionals update information (Section 1.1.1). We then show how both these approaches are temporally unsatisfactory if update is taken as a single, undecomposed operation (Section 1.1.3). Finally we indicate how splitting update into the twin operations of assertion and verification improves matters (Section 1.1.4).

1.1.1 Two Forms of Conditional Update

To see how conditionals place constraints on future updates, we will look at a simple example illustrating how information and information states get updated. For expository purposes, let us for the time being assume that information states are simply sets of propositions or formulas in some language.

To keep things simple, we will confine the example to the following four propositions: $p, -p, q, -q$[4]. From these, we can construct nine information states, which form a lattice-like structure shown in Figure 1.1. The nine states are ordered by a relation of *information containment*, $\sqsubseteq$, such that e.g. $\{\} \sqsubseteq \{p\} \sqsubseteq \{p, q\}$. That is, the state $\{p, q\}$ contains more

contrast, I am claiming that a conditional allows one to add an unverified consequent assertion to a given information state, and that in due course this assertion will come to be verified in the *same* state. There is thus a three way distinction between different types of information: verified, unverified, and indirect, where indirect information concerns what will happen in other information states.

[4]For now, don't worry too much about what is meant by a negative proposition like $-p$. Just let $-p$ be any old proposition that is incompatible with p, e.g. p = 'John is eating his dinner' and $-p$ = 'John is dead'. Negation is more fully discussed in Chapter 5.

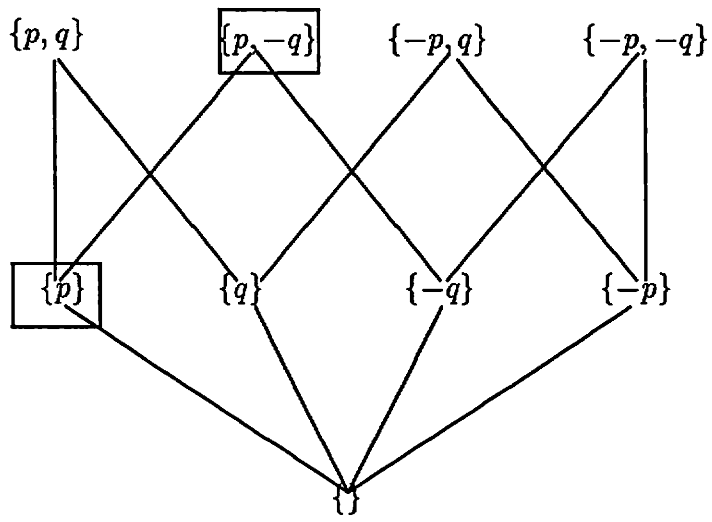

Figure 1.2: $p \rightarrow q$ in $\{\}$ prohibits boxed states

information than (is an extension of) the state $\{p\}$.

Let us suppose that we are currently in the empty information state. Let us further suppose that we process an utterance whose propositional content is p. This causes an update that has the effect of moving us from the empty state to $\{p\}$. If we were then to process an utterance of q, this would lead us into the state $\{p, q\}$. Processing an utterance of $p \wedge q$ would move us directly to $\{p, q\}$ from the empty state.

What about conditionals? There are two (nearly) equivalent ways of describing the effects of a conditional, both of which rule out certain possibilities for future update.

Conditional Update, Type 1: All Extensions

The first way of describing the conditional is borrowed from the treatment of intuitionistic implication in Kripke semantics for intuitionistic logic (see van Dalen 1984). The conditional $p \rightarrow q$ processed in a state s says that for s and all states s' extending s (i.e. possible future states for s), updating s' with p must lead to a state containing p and q. The effect of $p \rightarrow q$ on the empty state in our example is shown in Figure 1.2. Certain states that counted as possible extensions of $\{\}$ before the conditional was processed no longer count as possible extensions afterwards. The state $\{p, -q\}$ is ruled out because it can be derived by updating an extension of $\{\}$ with p, but not only does it not contain q, it contains the negation of q. The state $\{p\}$ is also ruled out because it is derivable by updating $\{\}$ with p, but does not contain q.

Conditional Update, Type 2: Minimal Extensions

The second way of treating a conditional like $p \rightarrow q$ relative to a state s is as follows. If s is updated with p, then it must lead to a state containing p and q. At first sight, this seems quite different from the previous treatment, which referred to updates in *all* states extending s. Here we consider only a minimal update to s made by adding p. However, provided that p and q are *monotonic* propositions, the two treatments turn out to be equivalent (see Chapter 3). A monotonic proposition is one that once it holds in a state s it continues to hold in all states extending s. In our example, $p, q, -p$ and $-q$ are all monotonic. So referring to Figure 1.2, $p \rightarrow q$ first rules out the state $\{p\}$. But once this is ruled out, only $\{p, q\}$ and $\{p, -q\}$ remain as states that could be derived from updating $\{\}$ with p. And $\{p, -q\}$ is in turn ruled out.

Whichever way we deal with $p \rightarrow q$, it should be apparent that information does not consist solely in the propositional contents of an information state. The range of possible extensions to a state, as given by $\sqsubseteq$, also provides a valuable source of information. One way of looking at conditionals (Landman 1985) is as a means of expressing contingent restrictions on possible extensions of an information state. Logically speaking, $\{p\}$ is a possible extension of $\{\}$, but the conditional $p \rightarrow q$, which reflects the properties of $\sqsubseteq$, contingently rules this possibility out.

Modal Update

Modals behave in a similar way to conditionals (Figure 1.3). To claim in $\{\}$ that it must be the case that q, $\mathbf{must}(q)$, is to say that all future courses of update must *eventually* lead to the addition of q. Or alternatively, that update of $\{\}$ with $-q$ would not lead to a permissible information state. This rules out states $\{-q\}, \{p-q\}$ and $\{-p, -q\}$ as possible extensions of $\{\}$. The observant reader might notice that $\mathbf{must}(q)$ has the same effect as $-q \rightarrow q$. (It should also be pointed out that the informational ordering $\sqsubseteq$ does not play the same role here as the accessibility relation familiar from Kripke semantics for modal logics (see Chellas 1980). $\mathbf{must}(q)$ does not force q to hold in *all* states extending $\{\}$ (c.f. all states accessible from $\{\}$). It only says that whatever happens, $\{\}$ gets extended in such a way as to eventually include q.) As with conditionals, there are two different approaches to modal update are possible. Either the modals legislate for all possible extensions of a state (type 1 update), or they legislate for immediate extensions of a state (type 2 update). Only the first kind of update has been described: a fuller description of the second type is deferred until Chapter 4.

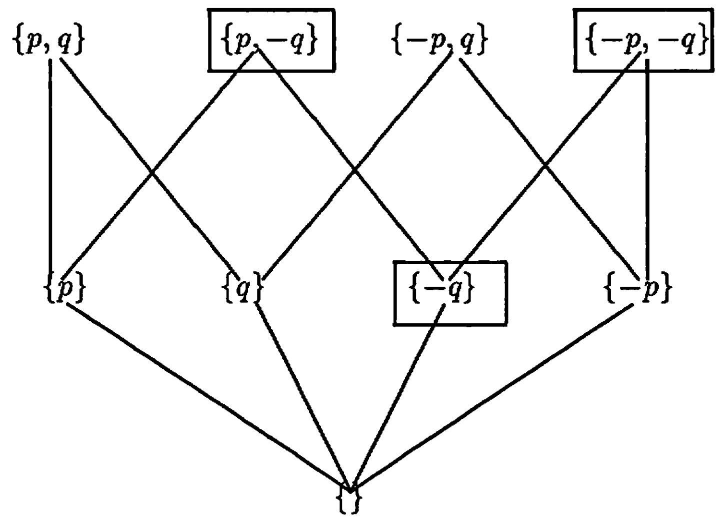

Figure 1.3: **must**(q) in $\{\}$ prohibits boxed states

Representing Information States

There are a couple of respects in which the above examples of modal and conditional update may be slightly misleading as to what will follow.

In Chapter 3, information states will be taken as primitive. That is, information states are *not* to be construed as sets of propositions or formulas. But as with possible worlds, it often does no harm to think of them in these terms, (and indeed this is necessary when constructing canonical models for completeness proofs — Chellas 1980, Appendix C). It is solely for ease of exposition that information states are represented as sets of propositions in this chapter.

It is important to forestall one possible confusion relating to the order of information extension, $\sqsubseteq$. The ordering over information states in Figure 1.1 is deliberately set up so that $\{\} \sqsubseteq \{p\} \sqsubseteq \{p, q\}$. This should not lead one into thinking that the informational ordering can be defined from the contents of information states, e.g. in terms of the subset relation such that $s_1 \sqsubseteq s_2$ iff $s_1 \subseteq s_2$. This just so happens to work in Figure 1.1. But in Figure 1.2, where the conditional $p \rightarrow q$ holds, it turns out that $\{\} \not\sqsubseteq \{p\}$, despite the fact that $\{\} \subseteq \{p\}$. It is by taking the informational ordering as primitive that we can capture contingent restrictions on information extension, as expressed by conditional and modal sentences.

Representing Update Statically

The examples above (Figures 1.1–1.3) describe update in *dynamic* terms: processing a non-modal, non-conditional sentence moves you from one information state to another, processing a modal or conditional sentence removes certain states from the set of possible extensions, and so on.

Updates can also be described more statically, in terms of their end results. For example, update via a sentence p results in a state containing the proposition p, and update by $p \rightarrow q$ results in a state that has no extensions containing p but not q. In Figure 1.1, the states $\{p\}, \{p, q\}$ and $\{p, -q\}$ are all possible results of an update by p. In Figure 1.2, all the states shown unboxed are possible results of an update by $p \rightarrow q$, provided that none of the boxed states count as possible informational extensions of them.

This kind of static representation loses some of the fine detail provided by a more dynamic representation. For example, if one processes an utterance of the sentence p in state $\{\}$ in Figure 1.1, one would expect to land up in the state $\{p\}$, and not $\{p, q\}$ or $\{p, -q\}$. This detail is lost in the static description. In compensation, static descriptions are simpler, and formally easier to deal with[5].

The level of detail provided by static descriptions of update is enough to get a workable treatment of tense and deictic shift going. It is therefore unnecessary to bring in the extra complications that would be introduced by a more dynamic account. (This is not to imply that more dynamic treatments of modality and conditionality, such as Veltman 1990, are without value. Far from it. But if things can be expressed in simpler terms, there is good reason to do so. Besides, it ought to be possible to incorporate any results arising from a static treatment into a dynamic treatment).

I will use the notation

$$s \models \phi$$

to say that the state s can be the result of an update by a sentence ϕ. For the relation $s \models \phi$ to hold, the state s will normally have to satisfy a number of properties, e.g. containing an assertion to the effect that ϕ. Another way of looking at this is that if $s \models \phi$, then updating s with ϕ will have no affect (add no further information), and leave one still in s. If $s, \models \phi$ holds, we will say that ϕ is *supported* by s. In due course, revised versions of the

[5]A static update representation would consist of a set of pairings of information states with information order, while a dynamic representation would be a set of pairs of such pairings, corresponding to a transition from one state-order pair to another.

Note that this use of the term 'static' does not correspond to that in Dynamic Montague Grammar (Groenendijk and Stokhof 1990), where it means a set of sets of states, rather than a set of states (plus orders) as here.

support relation $\models$ will be introduced, differing in that the relation will also take temporal arguments.

Despite the move to a more static representation of update, there is a trace of dynamism remaining in the informational ordering $\sqsubseteq$. In effect it determines which information states are permitted as results of future updates.

Which kind of conditional update?

Recall that our intention is to account for deictic shift by centering tenses on times at which updates are made. Which of the two kinds of conditional update presented above is better suited to this?

On the face of it, it would seem that the first kind of update, explicitly legislating for all future updates, is preferable. The second kind of update only seems to consider the effect of a single update, made at the present time. It does not appear to talk about future updates at all.

Yet it is the second type of update, legislating only for minimal information extensions, that will be employed. The reasons for this are as follows. First, there are asymmetries between the behaviour of past and present tenses in deictically shifting contexts. These are hard to account for if the tenses have just one deictic centre, based on a single, unitary operation of update. One must break the operation of update down into two components of assertion and verification, thus supplying primary and secondary deictic centres for the tenses.

Second, once update is decomposed into assertion and verification, the objections against the second form of conditional update disappear. Minimal extension conditional updates make present assertions, but allow for future verifications of those assertions. This allows for futurate deictic centres in conditionals (and modals).

Third, while an all extensions treatment of conditionals is compatible with breaking update into assertion and verification, it turns out to give the wrong temporal results.

The situation amounts to picking one possibility from the four-way choice shown in Table 1.1: Option 1 is ruled out because of the asymmetry of deictic shift. Option 2 is ruled out because of asymmetry, and because it does not plausibly legislate for future updates. Option 3 is ruled out because it gives the wrong results in certain cases (see p. 20). This leaves option 4.

The next two sections show why options 1 and 2 are unacceptable when time is incorporated into the treatment of conditional update. The section after (Section 1.1.4) shows how option 4 fares better.

	Unitary Update	Assertion & Verification
All Extensions	1	3
Minimal Extensions	2	4

Table 1.1: Four options for conditionals

1.1.2 Incorporating Time

If we are going to incorporate time and tense into the treatment of modal and conditional update, there are three points at which this introduction will make a difference. This is true whether update is taken to be a single operation, or is broken down into assertion and verification.

1. The temporal reference of the propositions contained in information states
2. The times at which states are current, and the times at which transitions are made between states as the result of an update.
3. Tenses contained in the sentences whose utterance leads to update.

We will deal with each of these in turn.

1. It is customary in tense logic to view propositions as functions from times (and perhaps possible worlds) to truth values. This is not the way that propositions in our information states behave. A proposition in an information state like p is really short-hand for saying that "p is true at t", which we might write more explicitly as $p(t)$. $p(t)$ is an eternal proposition — if p is true at t, then it is true at all times that it is true at t. To save possible confusion, we will refer to the eternal propositions contained in information states as *assertions*. An assertion states that a particular proposition holds true at a particular time.

2. It makes sense to talk about being in (possession of) a certain information state at a certain time. All being well, as time goes by this state will grow into other states that are informational extensions of the original. In reality, some of our initial information may prove to be mistaken, and so sometimes information may decrease over time. However, I will always be making the optimistic assumption of steady information growth in what follows.

Just because information grows over time, one should *not* conclude that the ordering of information extension, $\sqsubseteq$, is a temporal order. To do so would make time forward and backwards branching (i.e. alternate futures and alternate pasts). But time, I am going to assume, is linear, so that there is only one past and one future. On the basis of limited information, there may epistemically be more than one way that the past or future could be. But at most only one of these can correspond to the way that the past or the future actually is.

However, it is possible to give a loose connection between the branching structure of information states and the linear structure of times. Starting from the fact that one is, say, in a state s_i at time t_i, one can conclude the following. If $s_i \sqsubseteq s_j$, then if one is ever in state s_j it will be at a time t_j such that $t_i \leq t_j$. And likewise, if $s_h \sqsubseteq s_i$, then if one was ever in state s_h it would have been at a time t_h such that $t_h \leq t_i$. (Both these conclusions rest on the assumption of steady information growth.)

Despite this kind of connection between information states and times, ordered sets of information states are eternal in the same way that the assertions contained within the states are eternal. If $s_i \sqsubseteq s_j$, then it is always the case that $s_i \sqsubseteq s_j$. Information states and sets of them are essentially non-temporal structures. (There will be cause to qualify this slightly in Section 1.1.4, but the claim is substantially correct). Information states and their ordering into sets of possible extensions give a gods-eye, atemporal view of the world. A state can assert that p is true at t_p and that q is true at t_q and that $t_p < t_q$ regardless of whether t_p and t_q are before, after or simultaneous with some particular time taken to be the present moment. Similarly, the order over states is not dependent on the choice of a particular present moment.

3. If information states are atemporal, the same cannot be said for English sentences. The sentence *It is raining* can be used to make different assertions at different times. The role of the past and present tenses can be seen as one of tying the time varying nature of sentences down to the static assertions found in information states. What we need to do is show how, via the tenses, the temporally varying update potential of sentences maps onto the temporally fixed relations in information states and sets.

The distinction between *sentences*, *utterances* and *assertions* needs emphasis. Assertions are what are contained in information states. Utterances of sentences are (typically) what puts them there. As just noted, different utterances of the same sentence can be used to put different assertions in an information state. Likewise, different utterances of different sentences can sometimes be used to add the same assertion to an information state, e.g. *It is raining* uttered at 4:45, and *It was raining at 4:45* uttered at 5:45. It will sometimes be useful to talk of the assertion corresponding to the utterance of a sentence at a given time. But this should not lead one to confuse utterances with assertions.

1.1.3 Tenses and Conditionals with Unitary Update

Let us use the convention that Greek letters like ϕ or ψ can be used to refer to sentences / formulas of arbitrary complexity, and lower case letters like p or q to refer to atomic formulas, containing no tense, modal or conditional operators. These formulas will normally denote a function from times to assertions (eternal propositions). When this is so, $\phi(t)$ will be used to refer to an assertion, namely the assertion that ϕ is true at t.

The notation $s, t \models \phi$ will be used to say that an assertion corresponding to an utterance of ϕ made at time t is supported by a state s.

Using $\mathbf{past}(p)$ and $\mathbf{pres}(p)$ to represent simple past and present tense sentences, we can describe the update effects of utterances of simple tensed sentences as follows:

- $s, t \models \mathbf{past}(p)$ iff there is some time $t_p < t$ such that the assertion $p(t_p)$ is contained in s (i.e. $p(t_p) \in s$).
- $s, t \models \mathbf{pres}(p)$ iff $p(t)$ is contained in s

That is, an utterance of $\mathbf{past}(p)$ at a time t is supported in state s just in case there is some time t_p preceding t for which the assertion $p(t_p)$ is contained in s. And similarly with $\mathbf{pres}(p)$.

For conditional sentences, we will adapt the first account of conditional update mentioned in Section 1.1.1 (recall that this eventually proves unsatisfactory):

- $s, t \models \phi \rightarrow \psi$ iff for all $s' \sqsupseteq s$ and all $t' \geq t$, if $s', t' \models \phi$, then $s', t' \models \psi$

That is, the conditional uttered at t is supported in s just in case for any state s' extending s that supports an assertion of the antecedent ϕ made at a time $t' \geq t$, the state s' also supports an assertion of the consequent ψ made at t'. The reference to $t' \geq t$ perhaps needs some explanation. If the sentence leads one to a state s at t, then reaching any state $s' \sqsupseteq s$ can only take place at a time $t' \geq t$.

The important thing to note about conditionals is that the antecedent ϕ and the consequent ψ will be tensed clauses, e.g.

$$\mathbf{pres}(p) \rightarrow \mathbf{past}(q)$$

The tenses of these clauses will be interpreted relative to the future states and times s' and t'. As a result, the antecedent and consequent tenses are subject to deictic shift.

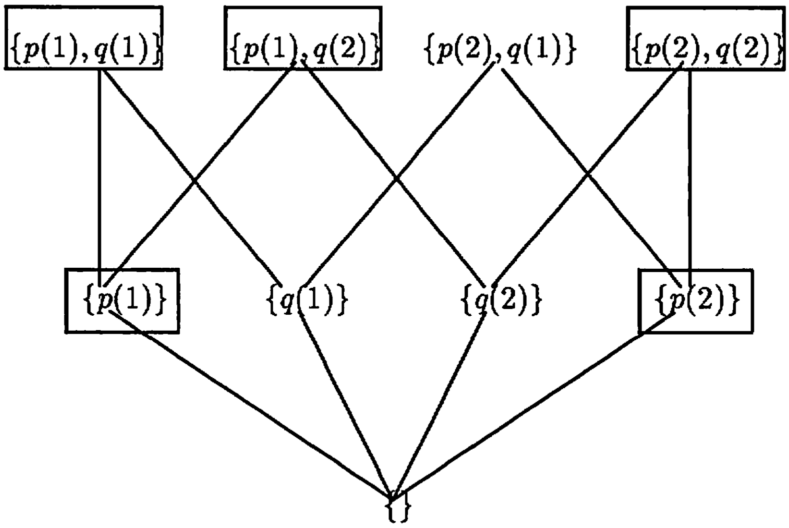

Figure 1.4: Effects of $\mathbf{pres}(p) \rightarrow \mathbf{past}(q)$ in $\{\}$ at time 1

Example

Let us see how this treatment deals with a $\mathbf{pres}(p) \rightarrow \mathbf{past}(q)$ conditional like (5) — $\mathbf{pres}(p)$ = 'I smile when I get out', $\mathbf{past}(q)$ = 'the interview went well'. Figure 1.4 shows what happens if $\mathbf{pres}(p) \rightarrow \mathbf{past}(q)$ is uttered with respect to the empty state at the time 1. This has the effect of ruling out the states shown boxed. The states ruled out are those that contain an assertion $p(t)$ but not a corresponding assertion of $q(t')$ for some $t' < t$, for all $t \geq 1$. Thus both antecedent and consequent may make assertions about what happens at times after the time of utterance, 1, but the consequent must precede the antecedent.

So, according to this treatment, the sentence (5) uttered to make an update at some time t means that if I smile at some time at or after t, the interview went well at some time before that. A past-in-the-future reading for the consequent is predicted.

Problems

While the preceding treatment seems to deal adequately with conditional sentences like (5), it suffers a number of flaws. I will mention just two.

First, the same analysis would predict that in the past tense bimetallic strip conditional (1) (formalised as $\mathbf{past}(p) \rightarrow \mathbf{past}(q)$) it should be possible to interpret both antecedent and consequent as describing events that occurred prior to some future update time t. But the antecedent and consequent may only describe events occurring prior to the present time of utterance. On the positive side, however, it does predict that no specific order between antecedent and consequent is imposed: they are both merely past with respect to some future update time.

Second, in the present tense bimetallic strip conditional (2), the analysis predicts that the strip bending and the temperature rising are simultaneous. Yet the antecedent can precede the consequent. We could loosen the treatment of the present tense so that it selects a time that is non-past with respect to the update time instead of simultaneous with it (as Lyons (1977), Nerbonne (1985), and Comrie (1986) advocate with respect to speech time). This allows antecedent to precede consequent, but at the unjustifiable expense of also allowing consequent to precede antecedent.

These two difficulties reflect the wider fact that the past and present tenses behave asymmetrically under deictic shift. This asymmetry is brought out more clearly by the following two pairs of sentences.

(7) a. If I smile when I get out, the interview went well.

b. If I smiled when I got out, the interview went well.

(8) a. I will marry someone who makes a killing on the stock exchange.

b. I will marry someone who made a killing on the stock exchange.

The present tense antecedent in (7a) is shifted forwards, but the past tense antecedent in (7b) is not. The present tense relative clause in (8a) can refer to times before or after the marriage, whereas the past tense relative clause in (8b) can only refer to times before the marriage.

The prospects for dealing with this asymmetry are slim so long as one employs only one deictic centre. This motivates splitting update into the processes of making an assertion and verifying the assertion.

1.1.4 Assertion and Verification

Verified and Unverified Information

Our original information states consisted of sets of assertions. The states, and the ordering over states, remained constant over time. We will now divide the assertions contained in an information state into two classes for any time t: those assertions that count as verified at t, and those assertions that are not yet verified at t. As time goes by, more and more of the assertions contained in a state will become verified in that state. Furthermore, no assertions will change from being verified to being unverified. This change means that information states are now (minimally) time varying.

Verification What is it to verify an assertion, and what is the difference between verified and unverified assertions? To answer this, it helps to note an obvious asymmetry between the future and the past. While we can have knowledge about both the past and the future, the origin of our information about the past is usually quite different from the origin of information about the future.

With information about the past, there is usually some causal chain linking a past event to present information about that past event. Failing backwards causation, the same is not true of information about the future. At best, one can extrapolate on the basis of what one knows about the past and the present. This is not to say that information about the past is invariably more reliable or certain than information about the future. Sometimes, information about the future can be very solid while that about the past is extremely flakey. The causal chain linking past events to present information may be long and tenuous, and subject to gross misconstrual (Evans 1982). For example, person A might see John Smith kiss his wife Mary. At some later point, he reports this fact to B by saying "I saw John kiss Mary". B then misinterprets the utterance as meaning that John Brown kissed Mary Smith, and starts spreading gossip. The false information that gets spread around nevertheless has a causal link to some past event, even though this link is widely misconstrued.

With verified information, it is normally assumed that there is some causal chain linking the information to what the information is about. When a causal chain is involved, verification of information about an event cannot precede the occurrence of the event. However, there are some cases where verification is not causal in origin. In these cases the temporal order between events and verification of information about those events may be less strict. I will mention two examples of this.

One example of where verification is not strictly causal in origin is mathematics. In (non-constructive) mathematical discourse, proving that something cannot be false (i.e. that it must be the case), is tantamount to verifying that it is the case. Consequently, there is a tendency to slip between talk of unverified and verified assertions, i.e. to switch between saying 'it must be the case that p' and 'it is the case that p'. This perhaps reflects the fact that the subject matter of mathematics is not primarily causal in origin.

Foreknowledge A more interesting case where a degree of prescience seems to apply is in talking about plans or pre-determined events. Dowty (1979) notes that when discussing planned or predetermined events, the simple present tense can appropriately be used. Thus

(9) I travel to Edinburgh next Wednesday.

describes a planned future event, despite the fact that a sentence like *I travel to Edinburgh* on its own can usually only be used to describe a present disposition towards travel.

When talking about plans, it is not really the world that we are discussing but what is laid out in the plan. Of the plan, we have a gods-eye view, so that we can survey in a glance which events are supposed to happen when. If the plan provides our source of information about these events, then we can verify that certain events will occur, according to the plan, well in advance of the time at which those events are supposed to occur.

It will be shown below that it is the possibility of *present* verification of future (planned) facts that allows the present tense to be used in sentences like (9).

The same, I suspect, is true when talking about pre-determined events. When one says

(10) The sun rises at 05:47 tomorrow.

one is in fact reporting on what is stated in a set of sunrise and sunset tables. As with plans, one can tell what will happen according to these tables in advance of the events happening.

Assertion, Verification and the Tenses

Simple, tensed sentences are normally used to convey verified information. Thus, if I say *John kissed Mary* I am urging the hearer to update his or her information state with the assertion that John kissed Mary at some time in the past, and moreover to treat the assertion as verified. Similarly, when I say *John is asleep*, I am inviting the hearer to add an assertion to his or her information state, and again count it as verified.

Updates consist of first making an assertion and then marking it as verified. For simple sentences, these two operations take place at the same time. This means that the update effect of a sentence ϕ at a time t needs to be expanded out as follows:

$$s, t \models \phi \text{ iff } s, t, t \models \phi$$

The two occurrences of t on the right hand side correspond to the *assertion* and *verification times* respectively. As we can see, to begin with the two times are assumed to be identical, and correspond to the time at which the sentence is uttered.

The effects of the tenses on atomic sentences may be described as follows (a more general definition is given in Chapter 3):

- $s, a, v \models \mathbf{past}(p)$ iff there is a time $t_p < a$ such that the assertion $p(t_p)$ is verified in s at time v.

- $s, a, v \models \mathbf{pres}(p)$ iff there is a time $t_p \geq a$ such that the assertion $p(t_p)$ is verified in s at time v.

Here, a stands for the assertion time and v for the verification time. An assertion $p(t_p)$ is verified in a state s at a time v if $p(t_p)$ is contained in the set of assertions that count as verified in s at v. (If we represent states as sets of assertions, at any one time this set will be split into two subsets of verified and unverified assertions).

The assertion time, which is the time at which an update is initiated by making an assertion, serves as the primary deictic centre for the past and present tenses.

The past tense selects an event time, t_p, that must precede the assertion time. The assertion $p(t_p)$ is the assertion corresponding to an utterance of **past**(p) relative to the assertion time a and the verification time v. The present tense selects an event time that is either simultaneous with the assertion time, or follows it. Both tenses leave the values of the assertion and verification time indices unchanged.

It might be thought that the present tense can readily be used to make futurate assertions. This is not the case. For (a) in simple sentences the assertion and verification times are identical, and (b) (causal) verification cannot precede the occurrence of the event described. That is, for a simple sentence corresponding to **pres**(p)

Initial setting of assertion and verification times:	$a = v = now$
Effect of present tense:	$t_p \geq a$
Verification:	$t_p \leq v$
Therefore:	$t_p = a = v = now$

By relaxing the constraints on verification, e.g. so as to allow foreknowledge about plans, **pres**(p) can be used to describe a future event. This explains why sentences like (9) are permissible when describing plans.

Conditionals can alter the sentence initial values of the assertion and verification times, allowing verification to occur after assertion. For this reason, present tenses in conditionals may have futurate time reference.

Minimal Information Extensions

The basic idea behind the treatment of the conditional is as follows. If $\phi \rightarrow \psi$ holds in a state s, then minimally extending s to verify an assertion of the antecedent, ϕ, results in a state that also supports an assertion of the consequent, ψ. Before elaborating on this, something needs to be said about what constitutes a minimal information extension.

Let us introduce the notation $s \sqsubseteq_v^{\phi,a} s'$. This says that s' is a state minimally extending s to verify at time v an assertion of ϕ made at time a. The time v is the earliest time not preceding a at which the state s' verifies the assertion. Furthermore, there must be no information state s'' extending s but subsumed by s' that verifies the same assertion sooner. (This notation can be defined in terms of $\sqsubseteq$: $s \sqsubseteq_v^{\phi,a} s'$ iff (i) $s', a, v \models \phi$ and (ii) there is no s'' and no v'' such that $s \sqsubseteq s'' \sqsubseteq s'$, $a \leq v'' < v$ and $s'', a, v'' \models \phi$.)

In order to make minimal information extensions do the job they need to for conditionals, a crucial extra assumption needs to be made. This is the *existence of ideally verifying information states.*

A given state s may support the assertion that p holds at some time t_p, but not verify this assertion until some time after t_p. The existence of ideally verifying states says that for any state like s, there will be another, s', supporting exactly the same set of assertions, but where the assertion $p(t_p)$ is verified as soon as it possibly could be. In the absence of foreknowledge, this means that s' verifies $p(t_p)$ as soon as the event described by $p(t_p)$ is finished, i.e. the end point of t_p. An ideally verifying information state corresponds to that of an ideal, but non-prescient, observer who always happens to be in the right place at the right time, and can verify that things occur as they occur.

Ideally verifying information states are always candidates for minimal information extensions. To see why, suppose that s' extends s only in that it supports an assertion $p(t_p)$ resulting from an utterance of the sentence **past**(p) at time a, where $t_p < a$. Suppose that s' does not verify the assertion until time v_p, where $v_p > a$. There is going to be an ideal state, s'', supporting exactly the same assertions as s', but verifying $p(t_p)$ at the end-point of t_p. This ideal state will continue to verify $p(t_p)$ up until time a and beyond. This ensures that s' is not a minimal extension of s. For while $s' \sqsubseteq s''$ and $s'' \sqsubseteq s'$, s'' verifies the assertion sooner after time a than s' does.

Given the existence of ideally verifying information states, it is possible to read the subscript v in $s \sqsubseteq_v^{\phi,a} s'$ as: the earliest time at or after a at which an assertion resulting from ϕ is *verifiable.* Being able to switch from saying when an assertion is *verified* to saying when it is *verifiable* is an important component of the treatment of tense and conditionals developed in this thesis.

Assertion, Verification and Conditionals

Having split update into two components of assertion and verification, and introduced the notion of a minimal information extension, it is now possible to resurrect the second treatment of conditional update mentioned in Section 1.1.1.

If state s is the result of an update by an utterance of $\phi \rightarrow \psi$ made at time a, the following must be true of s. Suppose s is minimally extended to include an assertion that would result from an utterance of ϕ made at time a. If v' is the first time at or after a at which that assertion is verifiable, then the extending state should support an assertion that would result from an utterance of ψ made at v'. More formally

- $s, a, v \models \phi \rightarrow \psi$ iff
  for all v' and for all s' such that $s \sqsubseteq_{v'}^{\phi,a} s'$: there is some v'' such that $s', v', v'' \models \psi$

Roughly put, the conditional acts as a way of delaying the assertion of the consequent until a (current) assertion of the antecedent is verifiable.

In other words, the conditional says the following. Suppose you were to add an assertion corresponding to an utterance of the antecedent made at the same time as the conditional is uttered. Then once this assertion is verifiable, which may be later, you should be able to utter the consequent and find that the resulting assertion is already supported. The consequent assertion need not be verified at that time, however.

Examples

We are now in a position to see how this treatment of the conditional interacts with tenses. For this purpose, let us go over the three conditional sentences (1), (2) and (5).

(1) past(p) $\rightarrow$ past(q) Suppose the conditional is uttered at a time a, resulting in a state s. Let $p(t_p)$ be an assertion corresponding to an utterance of **past**(p) made at time a. Because of the past tense, it follows that $t_p < a$.

Since $t_p < a$, the earliest time at or after a at which the assertion $p(t_p)$ is verifiable is a itself. Thus a gets passed on as the assertion time for the consequent. Since the consequent **past**(q) is also in the past tense, it follows that it will give rise to an assertion $q(t_q)$, where $t_q < a$.

We therefore arrive at a situation where p and q must both be true at times preceding a, but no relative order between the two is imposed. This corresponds to the fact that in *If the strip bent, the temperature rose* the strip bending may either precede or succeed the temperature rising, but both events must occur in the past.

(2) pres(p) $\rightarrow$ pres(q) Again suppose that the conditional is uttered at a. Let $p(t_p)$ be an assertion corresponding to an utterance of **pres**(p) made at a. Because of the present tense, it follows that $t_p \geq a$. More importantly, this assertion does not have to be verified at time a, and so it is not the case that $t_p \leq a$. Instead, we have $t_p \leq v'$ for some $v' \geq a$.

Assuming lack of foreknowledge, the earliest time, v', at which the assertion $p(t_p)$ can be verified is the end point of t_p. This gets passed on as the assertion time for the consequent. Again this is in the present tense, and so the assertion resulting, $q(t_q)$, is such that $t_q \geq v'$. And again, the consequent assertion does not have to be verified straight away, so that it is not also required that $t_q \leq v'$.

We therefore arrive at a situation where $a \leq t_p \leq v' \leq t_q$. So in *If the strip bends, the temperature rises*, it is predicted that the strip bending precedes or is simultaneous with the temperature rise, and that both events may occur now or in the future[6].

It is interesting to see what happens if foreknowledge is permitted. In this case, the antecedent may be verified before t_p, and in fact at a. We therefore have $a \leq t_p$ and $a \leq t_q$, i.e. no order is imposed between antecedent and consequent. As conditionals about plans show (like the following one discussing a rota for household chores)

(11) If I do the cooking, Valeria does the vacuuming.

this lack of relative ordering is precisely what obtains. (11) merely states that if I do one job, Valeria will do the other, but says nothing about the order in which they are done.

(5) pres(p) $\rightarrow$ past(q) It is easy to show that in the conditional corresponding to (5) the treatment here predicts a past-in-the-future reading for the consequent past tense: the present tense antecedent passes on a future verification time to be used as the assertion time for the past tense consequent.

Summary

The examples above bring out the point that there are two different kinds of deictic shift. The first kind involves shifting the primary deictic centre of a tense, i.e. the assertion time. It is this that allows past tenses in conditional consequents to refer to future time, since the deictic centre is shifted to an even more future time. It is also this that forces the antecedent precedes consequent ordering in present tense conditionals like (2).

The second kind of deictic shift arises when the verification time for a tense is no longer constrained to be identical to its assertion time. This only affects present tenses. It is this that allows present tenses in conditional antecedents to refer to future times. When the assertion and verification times are identical and foreknowledge is not possible, the present tense is forced to refer to a time point identical to the assertion time. But when the verification time is allowed to succeed the assertion time, as happens in conditional antecedents, the present tense can refer to a time that is merely non-past with respect to the assertion time.

A form of secondary deictic shift also explains the use of the futurate present when

[6]If the conditional had been given an all extensions analysis,

- $s, a, v \models \phi \rightarrow \psi$ iff
  for all $v' \geq a$ and for all s' such that $s \sqsubseteq s'$ and $s', a, v' \models \phi$: there is some v'' such that $s', v', v'' \models \psi$

we would get disastrously wrong results. For once the antecedent is verified in s', it will remain verified in s' at all times afterwards. This means that the present tense consequent can be truthfully uttered at all times after the initial verification. That is, once it is verified that the strip bends, the temperature will be rising forever after. This is why the third option of combining assertion and verification with an all extensions analysis of conditionals has to be rejected (see p. 10).

describing planned or otherwise pre-determined events. It does this without appeal to a hidden future tense operator in the logical form of the sentence (c.f. Dowty 1979). Deictic shift in conditionals also explains certain aspects of the habitual use of present (and past) tenses.

The combination of minimal information extensions with the notion of assertion and verification provides the backbone of the rest of this thesis. Obviously there are a number of details that need to be filled in, and a number of extensions that need to be discussed. The next section gives an overview of this development.

1.2 Thesis Organisation

Chapter 2 starts by summarising the temporal properties of indicative conditional sentences and other sentences containing present tense modal auxiliaries (*must, may, will, can* etc.). There is no particular virtue in novelty for its own sake; if this data can be accounted for in other, more familiar ways, there is little point in pursuing the kind of approach described in the last section. The chapter therefore considers a number of different ways in which one might explain the data, all of which turn out to be unsatisfactory.

Chapter 3 turns to a more positive account. Information models consisting of information states that can support both verified and unverified assertions are formally introduced. Following roughly the same course as this introduction, semantics are given for past and present tense operators. Two conditional connectives are then introduced. The first of these, $\rightarrow$, supports ordinary *modus ponens* inferences: from ϕ and *if* ϕ *then* ψ, conclude ψ. The second, $\Rightarrow$, supports only a weaker futurate version of *modus ponens*: from ϕ and *if* ϕ *then* ψ, conclude *it will be the case that* ψ. (The second connective is needed for the analysis of present tense habitual conditionals like *If the bimetallic strip bends, the temperature rises.* Given this conditional and the fact that the strip is bending or has just bent, one is only entitled to conclude that the temperature will rise, but not necessarily that it is already rising.) It is then shown how a habituality operator, H, can be defined in terms of the second conditional connective. It is shown how these tense and conditional operators exactly account for the range of temporal interpretations exhibited by simple (i.e. non-modal) indicative conditional sentences.

Chapter 4 deals with the present tense modal auxiliaries. These are treated as a present tense operator applied to a modal operator. It is shown how this treatment supplements that in Chapter 3 to account for the temporal properties of modalised indicative conditionals, as well as for the patterns of deictic shift found in subordinate clauses with modal superordinate clauses. It is also shown how other deictically shifted subordinate tenses can be accounted for. This includes a brief comparison of the temporal connective *when* with the conditional, and also discusses sequence of tense in clausal complements to verbs.

Chapters 2–4 are concerned almost exclusively with the temporal properties of modals and conditionals. They take it on faith that once you get the temporal properties sorted out, the logical properties will to a large degree look after themselves. Chapter 5 shows that this faith is not entirely misplaced. Intuitionistic logic is presented as logic of verified truth. It is shown how this can naturally be extended to provide a logic of verified and unverified assertion, which has a model theory that is essentially that of the information models introduced in Chapter 3. Conjunction, disjunction, implication and two forms of negation (one corresponding to an assertion of denial, and the other to a denial of assertion) are defined with respect to this model. A sound and complete proof theory is presented for these connectives. It is also shown how different combinations of the two negation operators give rise to modalities corresponding in a fairly natural way to epistemic *must, may* and *will.* The logic of verified and unverified assertion provides the internal logic of information states, dealing with assertions whose temporal reference is already fixed. As such, tenses play no role in it, since the purpose of tense is essentially to take an utterance of a sentence, and convert it into a temporally specific assertion. However, once tense is taken into account, it can be seen that the two conditional connectives of Chapter 3 are but two facets of a single underlying notion of implication. The chapter also discusses how the epistemic modals may be generalised to cover non-epistemic modals.

It is not possible discuss the temporal properties of modals and conditionals without at some point talking about counterfactual or subjunctive conditionals. Chapter 6 therefore deals with the hypothetical modals *would, could, might* and *should*, and conditionals containing them. However, this chapter is in some ways tangential to the rest of this thesis. It is argued that the hypothetical modals are not simply past tense forms of the present tense modals. Instead, they are implicitly conditional. In subjunctive conditional sentences, the explicit antecedent serves to specify the implicit conditionality of the modal. It is further argued that the tensed forms found in subjunctive conditionals reflect subjunctive (or conditional) tenses, and not the ordinary indicative tenses dealt with in the preceding chapters. Moreover, the subjunctive tense applies more to the conditional than to the antecedent and consequent. Subjunctive conditionals are treated as having present tense antecedents and consequents within the scope of an overall conditional tense. (This treatment bears some relation to Thomason and Gupta's (1981) suggestion that subjunctive conditionals are past tense indicative conditionals, though the correspondence is not exact). It is shown how a version of Kratzer's (1979, 1981; also Veltman 1976) premise semantics for counterfactual conditionals can be adapted to give an information state selection function.

Chapter 7 concludes by discussing some further directions for development. One area that needs more study is the interaction between tense, modality, conditionality and quantification. Another is the broader consequences for a theory of temporal reference of introducing the notion of verification time. One obvious question to ask is how far verification time relates to Reichenbach's (1947) reference time. Computational applications of the theory of tense and conditionals are also mentioned.

The thesis also includes three appendices. The first provides a simple unification based grammar for building up logical representations of conditional sentences. The second proves a number of claims made in Chapter 3. The third gives the soundness and completeness proof for the logic of verified and unverified assertion in Chapter 5.

Caveat Lector

There are a few aspects of this thesis that may seem rather surprising at first, especially to anyone familiar with the literature on tense or on conditionals.

First, this thesis is primarily about temporal reference in natural language (or rather English — most monolingual English speakers are over inclined to regard the two as co-extensive). Traditional problems about the logic of conditionals (such as the paradoxes of implication and the apparent invalidity of logically valid principles like transitivity, contraposition and strengthening the antecedent) receive relatively scant attention. Also, the problematic area of counterfactual conditionals is not explored in as much depth as the voluminous literature on it would suggest it merits. This is because temporally, and perhaps also logically, I am inclined to regard indicative conditionals as the more basic, and also the more interesting case.

Second, as a work on temporal reference, this thesis may seem surprisingly limited in scope. It is probably fair to say that most work on temporal reference nowadays concentrates either on aspectual problems (e.g. Dowty 1979; Taylor 1977; Krifka 1991; Verkuyl 1989; Moens and Steedman 1988), or on temporal anaphora and the temporal coherence of discourse (e.g. Kamp and Rohrer 1983; Partee 1984; Lascarides, Asher and Oberlander 1992). The indexical nature of tense has received relatively little attention recently (despite the emphasis on indexicality in Situation Theory), perhaps because it has been generally felt that there is not much more of interest to be said about it: tenses, by and large, are centred on the speech time. However, the behaviour of tenses in conditional and modal contexts indicates that the matter is far from resolved, and far from uninteresting.

Finally, for a work that takes as one of its basic premises the idea that tenses may describe how information changes as well as how the world changes, the static approach to representing information change may come as a surprise. As pointed out earlier, I believe that static descriptions of information change furnish sufficient detail to get a workable theory of tense and conditionals going. This is not to say that I believe that the current interest in dynamic descriptions of update (e.g. Veltman 1990) is of no value. Far from it. Admittedly, there are a number of problems in this area that still need to be resolved, but once they are I would hope that a more dynamic treatment could enrich the essentially static theory presented here. But in the meantime, since static treatments of update are enough to be going on with, it is easier to work from this more familiar territory.

Chapter 2

Temporal Reference in Indicative Conditionals

This chapter describes the temporal properties of indicative conditional sentences and of simple present tense modal sentences. A number of combined treatments of tense, modality and/or conditionality are discussed, and found wanting.

2.1 Patterns of Temporal Reference

2.1.1 Terminological and Methodological Preliminaries

Terminology

Conditionals are *"If Antecedent then Consequent"* sentences. There are surface variations on this basic pattern, for example *"If Antecedent, Consequent"*, *"Consequent if Antecedent"*, *"Suppose Antecedent. Then Consequent"*.

Simple conditionals are conditionals where the antecedent and consequent are past or present tense clauses containing no further conditional constructions or modal auxiliaries. Examples are

(1) a. If the baby cries, it is hungry.

b. If the baby cried, it was hungry.

c. If the baby cried, it is hungry.

d. If the baby cries, it probably soiled its nappy.

As these examples illustrate, there is no restriction on the tenses of the antecedent and consequent.

Imperative conditionals are conditionals where the consequent clause is in imperative mood, e.g.

(2) If the baby cries, feed it.

I will not deal with imperative conditionals.

Modalised conditionals are conditionals where the consequent, and possibly the antecedent, are modal clauses. Examples are

(3) a. If John comes on time, we can/could leave early.

b. If John came on time, we can leave early.

c. If John came on time, we could leave early.

d. If John can come on time, I can come on time too.

To draw further distinctions between different types of modalised conditionals, something needs to be said about the modal auxiliaries.

Following Quirk *et al* (1985) the modal auxiliaries can be divided into the following classes

- Central Modals (present tense): *may, can, will, shall, must*
- Central Modals (hypothetical): *might, could, would, should*
- Marginal Modals: *dare, need, ought to*
- Semi-Modals: *have to, have got to, be going to, be to, be bound to, be able to, be supposed to,* etc.

The central and marginal modals are all proper auxiliary verbs, satisfying Palmer's (1986) 'NICE' properties of (i) negation — the verbs can occur in negated form (*can't, couldn't, won't*) etc., (ii) inversion — the verbs can undergo subject auxiliary inversion to form interrogatives (e.g. *must I go?*, (iii) code — the verbs can stand as elliptical verb phrases (e.g. *If you won't, I will*), (iv) emphasis, or rather lack of it — the verbs do not permit emphatic positives (**I do can go*). In addition to these standard syntactic auxiliary properties, the modals have (v) no non-finite forms (**musted, *musting*), (vi) no *-s* forms (**musts*), and (vii) take infinitival verb phrase complements. The marginal modals *need* and *dare* also occur as non-auxiliary main verbs and only occur as modals either inverted or negated.

The semi-modals have finite forms, and constitute less of a closed class then the central and marginal modals. Strictly, it is the first verb in the semi-modals (*have* or *be*) that behaves like an auxiliary. But the semi-modals form a single unit (one cannot have *be going to* in non-progressive form for instance — **Several teams go to be beaten tomorrow*). The main interest in semi-modals is that they can fill slots in the modal verb paradigm where modal auxiliaries cannot occur because of their restriction to non-finite forms — *We might have to go. *We might must go.* By and large, I will concentrate on the central modals.

Subjunctive conditionals have hypothetical modal consequents and 'back-shifted' past tense antecedents. Back-shift occurs when a past tense form is used to refer to present or future times, and a past perfect form is required to refer to past times. Hypothetical modals and subjunctive conditionals are dealt with in Chapter 6.

Temporal Properties

Three temporal properties are relevant to indicative conditionals:

1. What is the temporal ordering between the eventualities described by the antecedent and consequent clauses; does the antecedent eventuality precede or follow the consequent eventuality, or are they simultaneous?

2. How are the antecedent and consequent eventualities located relative to the time at which the conditional is uttered (the speech or utterance time)? Do they describe eventualities occurring in the past, the present or the future relative to the speech time?

3. Does the conditional express a connection between two specific eventualities, or does it express a general connection between eventualities of that type?

The tense (past or present), and to a lesser extent the aspectual class (for now, stative or non-stative), of the antecedent and consequent place a number of constraints on these properties.

With (indicative) modalised conditionals, one must distinguish between the temporal reference of the modal, and that of the embedded eventuality or proposition. This difference is brought out by

(4) John may be at the party tonight.

where the present possibility of a future eventuality is described. The modal *may* has present time reference, while the embedded eventuality has future time reference. For modalised conditionals, the following temporal properties are relevant:

1. What is the relation between the modal temporal reference of antecedent (if any) and consequent?

2. What is the relation between the modal temporal reference of the consequent and the temporal reference of the antecedent eventuality (and likewise for the temporal reference of the antecedent modal, if any, and the consequent eventuality)?

3. What is the relation between the modal temporal reference of the antecedent / consequent, and the temporal reference of the antecedent / consequent eventuality?
4. What is the relation between the temporal reference of the modals and embedded eventualities and the speech time?
5. Does the conditional express a specific or general connection?

With modalised, non-conditional sentences, the relation between the temporal reference of the modal, the embedded eventuality and the speech time should be considered. When the modal sentence contains a tensed subordinate clause, as in

(5) By 1998, everybody will know someone who died of AIDS.

we also need to consider the relation between the temporal reference of the subordinate eventuality, the superordinate eventuality, the superordinate modal and the speech time.

Semantics and Context

Pragmatic factors such as world knowledge can strongly influence the way that a conditional is interpreted. In the past tense version of the bimetallic strip conditional

(6) If the bimetallic strip bent, the temperature rose.

knowledge of the way that thermostats operate allows us to read the conditional as expressing either order between antecedent and consequent. With other formally similar examples like

(7) If the ball was dropped, it fell.

pragmatics only permits one ordering. Our concern is with formal, semantic constraints on temporal reference, and so some way of factoring out pragmatic influences is required.

I will assume that pragmatic factors operate monotonically to further refine formal, semantic constraints on interpretation. For example, (6) indicates that simple conditionals with past antecedents and consequents do not semantically constrain the temporal order between antecedent and consequent eventualities. In (7), background knowledge eliminates one formally possible ordering as pragmatically implausible. The idea that contextual resolution and plausibility checking narrows down the range possible interpretations is central to the monotonic interpretation of Alshawi and Crouch (1992), and monotonicity has also been advocated elsewhere, e.g. Pollard and Sag 1987, Fenstad *et al* 1987.

To uncover semantic constraints on interpretation, it is important to look for cases where pragmatically plausible readings are blocked by the semantics, as this is a way of factoring out pragmatic effects. The unavailability of a sensible reading for a conditional like

(8) ?If the lie detector registers above ten, the suspect tells a lie.

is a case in point. Pragmatically, the read-out on the lie detector must follow the suspect lying, but semantics blocks this ordering. As Higginbotham (1985) put it (in a slightly different context), negative data is as important to semantics as positive data. Using the tension between semantics and pragmatic plausibility is a short-cut way of gathering negative data. In principle, to establish that a particular pattern of temporal reference cannot occur in a given conditional configuration one would have to look at all possible examples of that configuration, whereas to establish that it can one need only find a single instance. But if a pragmatically plausible reading is blocked, we can be reasonably confident that this is due to a negative semantic constraint.

Tense and Aspect

One cannot hope to describe the temporal properties of conditionals and modals in isolation from other mechanisms for temporal reference in English. These mechanisms are complex, and have been extensively discussed. For our current purposes a number of very general comments about tense, aspect and temporal reference suffice.

Tense Common to nearly all treatments of the past and present tenses is the idea that they select an 'event' time located relative to some deictic centre; usually, the deictic centre is taken to be the speech time. The event time is the time at which the embedded proposition or eventuality holds. Thus, a sentence like *John slept* says that there was an event of John sleeping that occurred at a time prior to the utterance of the sentence.

While there is little dispute that the past tense selects an event time preceding the deictic centre, the effect of the present tense in English is less clear. One school of thought (e.g. Dowty 1979, 1982) holds that the event time must be simultaneous with the deictic centre. Another (e.g. Lyons 1977, Nerbonne 1985) holds that the event time need only be non-past with respect to the deictic centre (or more accurately, simultaneous with or following it). Since the present tense often permits a futurate interpretation, I will follow the latter school (see below).

Perfect and Progressive For sentences in the simple present or past tense, the tense's event time is the time at which the basic eventuality occurs. But if the sentence is in the perfect or the progressive, an extra level of indirection is introduced. A sentence in the past or present perfect locates the basic eventuality as occurring prior to the tense's event time. A sentence in the past or present progressive describes an eventuality that is in progress at the tense's event time, or sometimes an eventuality whose completion is expected at this time. The perfect and progressive can be combined to describe an eventuality in progress at some time prior to the time selected by the tense.

Harper and Charniak (1986) provide one way of making this crude treatment of the perfect and progressive precise. Basically, the perfect and progressive select a new event time, located relative to the original event time set up by the tense. Unfortunately, no account is taken of such things as the 'current relevance' aspects of the perfect, or the imperfective paradox and the progressive. A variety of more sophisticated analyses of the perfect and progressive have been proposed (e.g. McCoard 1978, Lascarides 1988). What all these have in common with the cruder treatment is: (i) the perfect involves having some eventuality occurring prior to the time selected by the tense, and (ii) the progressive having some event in progress or whose completion is expected.

Temporal Adverbials The tenses and the perfect and progressive can all be seen as selecting a new value for an 'event' time. By contrast, temporal adverbials like *at 3pm* or *tomorrow* do not select a new value, but merely express properties that must be true of the value selected by the tense or perfect or progressive[1]. Consequently, adverbials cannot override any constraints on possible event times imposed by the tenses.

Temporal Anaphora Tenses often behave anaphorically (Partee 1973, 1984). Instead of selecting some arbitrary new event time, they may select a time that is either salient in context, or related in some way to such a time. (Often, the tense operates as an indefinite existential quantifier for an event time, but where the domain of quantification is anaphorically restricted. Thus Partee's famous example *I didn't turn off the stove* doesn't mean that there is some specific time at which I neglected to turn off the stove, but instead that during some specified period there were no instances of the stove being turned off). Temporal anaphora accounts for much of the temporal cohesion in a narrative. In particular, state-like sentences are often construed as referring to the start or end point of some event occurring in a narrative, due to a process of anaphoric linking.

Aspectual Class Vendler's (1967) classification of states, activities (or processes), achievements and accomplishments often forms the starting point for discussions of aspectual class (aktionsart). Aspectual class affects a variety of things, such as the kind of temporal adverbial verb phrases can co-occur with (e.g. only states and activities readily co-occur with durative *for*-adverbials), or whether a verb phrase can occur in the progressive (stative verb phrases usually cannot). I will operate with a much cruder classification into states and events (where events include Vendler's activities, achievements and accomplishments). The important thing about a state is that it can hold at a time instant, whereas an event cannot[2]. If one assumes that the deictic centres for tenses are time points, this means that

[1]Dowty (1982) makes essentially this claim. Note though that distributive adverbials like *every day* may further subdivide the event time already selected.

[2]Vendler viewed achievements as 'instantaneous' events, but Verkuyl (1989) has pointed out that even instantaneous events like blinking, dying or realising something have duration when viewed closely enough.

it is only states that can occur simultaneously with the deictic centre, e.g. in present tense sentences. Present tense event sentences, by contrast, tend to have an interpretation where they refer to a presently holding habit or propensity[3].

The Futurate Present Futurate uses of the present tense pose a problem for more traditional treatments of the present. If the present tense selects a present rather than merely non-past event time, then to account for sentences about pre-determined or pre-arranged events (e.g. *The Sun rises at 05:47 tomorrow*, or *Arsenal play Spurs at home next week*), one has to introduce a hidden future tense operator (Dowty 1979). If the present tense is viewed merely as selecting a non-past time, this problem disappears. But then it is hard to explain why futurate reference is usually only permitted when discussing pre-arranged or pre-determined events.

One quite important point to note is that present tense conditional clauses can often have a futurate interpretation, but without any element of pre-arrangement of pre-determination being involved.

Data Collection

Since negative examples are as important as positive examples, one cannot produce a descriptive survey entirely on the basis of corpora. In determining the temporal properties of conditionals, extensive use has been made of constructed examples, obtained by varying the tense, aspectual class and pragmatic connection between antecedent and consequent. However, the Lancaster-Oslo-Bergen (LOB) corpus of British English and the extracts from the Wall Street Journal contained in the Penn Treebank have proved very useful as a further check on the constructed examples. Especially in the case of modalised conditionals, corpus data has brought to light a number of permitted interpretations of conditional forms that might otherwise have been overlooked.

So while the data presented below does not pretend to be the result of a corpus based survey, the results have been checked against corpora.

The only example of a truly instantaneous event is something like the point at which a ball thrown in the air reaches the top of its trajectory, stops, and then immediately starts to fall (Galton 1984).

[3]'Sports reporter' uses of the present tense are an exception to this.

2.1.2 Simple Conditionals

1. Past Tense Precedes Present Tense

When a simple conditional contains a past tense antecedent and present tense consequent or vice versa, the event time selected by the past tense always precedes that selected by the present tense.

For example

(9) a. If the vase fell over, it is on the floor.

b. If the vase is on the floor, it fell over.

c. ??If the vase falls over, it was on the floor.

d. ?If the vase was on the floor, it falls over.

Neither of (9c) and (9d) permit the interpretation where the vase is on the floor because it fell over. They do have pragmatically possible readings where a *tendency* for the vase to topple is inferred from its having been on the floor, but in this case being on the floor precedes the presently attributed tendency.

If the present tense clause is a present perfect, the present tense eventuality can precede the past tense eventuality, e.g.

(10) If John sent flowers, then he has heard that Mary is ill.

(where John sends the flowers because he hears that Mary is ill). But this is only to be expected given the nature of the perfect as a form of non-deictic past tense. The time relative to which the present perfect expresses pastness still succeeds that of the past tense eventuality.

2. Futurate Past Tense Consequents

When a past tense consequent co-occurs with a present tense antecedent, the past tense can be interpreted as expressing pastness in the future, e.g.

(11) a. If I smile when I get out, the interview went well.

b. Usually, if I am grumpy in the morning I didn't get a good night's sleep.

In (11a) my smiling is caused by the fact that the interview I am about to go into went well, and in (11b) lack of sleep is what causes the grumpiness. Note that (11a) expresses a specific connection between two events, while (11b) expresses a general connection.

In all cases like this, a present tense antecedent with future time reference is required[4]. As dictated by the regularity noted previously, the consequent event precedes the antecedent event.

3. Futurate Past Tense Antecedents

It is also possible to have past tense antecedents that express pastness in the future, e.g.

(13) a. If I didn't get a good night's sleep, I am usually grumpy in the morning.

b. If the speaker gave an interesting talk, we usually buy him or her a drink. (But if they give a bad talk, they buy the drinks).

c. ...that has usually meant that ...I have accepted it if I was wrong. (From LOB)

Here, a futurate present tense consequent is required. Unlike the previous case, past-in-the-future readings are only possible if the conditional expresses a general regularity.

4. Distribution of Futurate Consequents

Conditional consequents (past or present tense) may only be futurate if the antecedent is futurate. As a result, conditionals with past antecedents and present non-stative consequents usually sound odd if the consequent cannot naturally be interpreted habitually or as expressing a command:

(14) a. ?If the bimetallic strip bent, the temperature rises.

b. If John had a packet of cigarettes in his pocket, then he smokes.

c. If you were at the meeting yesterday, you go again today. OK?

Likewise, conditionals with present tense stative antecedents referring to the present are odd when combined with present tense non-stative consequents

(15) ??If the bimetallic strip is bending (has bent), the temperature rises.

Futurate antecedents do not always give rise to futurate consequents:

(16) a. If John comes to the party tonight, then he got my message yesterday.

b. If the letter arrives tomorrow, it is already in the post.

---

[4]There is an apparent exception to this. Suppose an instruction manual for a lie detector says something like

(12) Ask the suspect a question, and then look at the read-out. If the lie detector registered above ten, the suspect lied.

Here, the conditional refers to events that occur in the future. However, this is not due to the conditional itself, but to the 'modally subordinating' context (Roberts 1989) set up by the preceding sentence.

5. Distribution of Futurate Antecedents

Futurate past tense antecedents are restricted to habitual conditionals with futurate present tense consequents. The same is not true for present tense antecedents. They occur very widely with futurate interpretations.

(17) a. If I shut the door, John is asleep.

b. If John comes to the party, he got my message.

c. If the letter arrives tomorrow, it is already in the post.

d. If the bimetallic strip bends, the temperature rises.

There are no restrictions on when a present tense antecedent can take on a futurate interpretation.

There are restrictions on when present tense antecedents can take on non-futurate (i.e. present) interpretations. Non-stative clauses that cannot be interpreted habitually force a futurate interpretation:

(18) If John solves that problem, he is cleverer than I thought.

(non-stative clauses cannot refer to a present time instant, unless construed habitually).

6. Past Tense Antecedent and Consequent

When both antecedent and consequent are in the past tense, no particular order is imposed on the event times selected by the two tenses, other than that they both precede the time of utterance. This is illustrated by

(19) a. If the bimetallic strip bent, the temperature rose.

b. If the temperature rose, the bimetallic strip bent.

The antecedent precedes consequent reading is perhaps more common, and reflects predictive reasoning from cause to effect. The consequent precedes antecedent reading reflects explanatory reasoning from effect to cause. Pragmatic knowledge can rule out one of the semantically possible readings as contextually implausible, e.g.

(20) a. If the lie detector registered above 10, the suspect lied.

b. (?)If the suspect lied, the lie detector registered above 10.

In these examples, telling the lie must precede the detector reading (an explanatory reading, (20b) is a little odd in this context).

Conditionals with past tense antecedents and consequents can be used to express either specific connections between events, or general regularities. When expressing a general regularity, a reading where antecedent precedes consequent is mildly preferred, but is not a necessity, e.g.

(21) Usually, if someone was buried in an unmarked grave, they died a pauper.

7. Present Tense Antecedent and Consequent

When both antecedent and consequent are present tensed, the time selected by the consequent tense usually is either simultaneous with or succeeds the end point of the time selected by the antecedent tense. This ordering between times does not translate directly into a similar ordering between events. The reason for this is that a stative eventuality can extend on either side of a time at which the state holds (e.g. if John is asleep at midnight, he is probably also asleep just before and just after midnight).

When the antecedent and consequent both describe events, which cannot extend beyond their allotted time, one usually gets a reading where the consequent event starts when the antecedent event finishes, or at some time after this. But when the antecedent and/or consequent describe states, overlap between the two is possible.

The strict ordering between events is illustrated by a sentence like

(22) If the bimetallic strip bends, the temperature rises.

and the possibility of overlap by

(23) a. If the bimetallic strip bends, it is hot.
b. If John is in a good mood, he is pleasant to be with.
c. If John sees a good film, he is happy

When both antecedent and consequent are stative, there is a fairly strong preference for taking them to be simultaneous, (23b). But this does not have to be the case

(24) If John is in a bad mood, he is usually cheerful again within an hour.

There are a number of exceptions to the ordering between the antecedent and consequent event time.

1: When the antecedent describes a pre-arranged or pre-determined event, any ordering between antecedent and consequent is permitted, as in

(25) If Valeria does the vacuuming, I do the cooking.

where a household rota is being described.

2: When the consequent is a stative clause referring to the present time

(26) If the letter arrives tomorrow, it is already in the post.

3: Explicit adverbials can also reverse the usual order.

(27) If the bimetallic strip bends, the temperature rises first / beforehand.

In rare cases, the reverse order is obtained without adverbials or pre-arrangement

(28) Usually, if the Prime Minister gives a good speech, someone else writes it for him.

There are also some more marginal, though common, cases that may or may not count as exceptions, e.g.

(29) If John goes out in the rain, he takes his umbrella.

In this case, it is probably the case that taking the umbrella occurs when John plans or arranges to go out in the rain, and so this falls under the heading of the first exception noted. This kind of ordering seems to be ruled out if pre-arrangement is implausible, e.g.

(30) ?If the machine blows up, it delivers a warning message.

(This suggests that the machine delivers a warning message once it has exploded.)

Also, if the consequent is a present perfect the consequent eventuality may precede the antecedent eventuality. But this does not show that the time selected by the consequent present tense also precedes that selected by the antecedent.

8. Generalisations

Conditionals cannot be interpretated as generalisations if the span of time between antecedent and consequent includes the time of utterance. This is not surprising since generalisations express a connection between arbitrary events and times, whereas the speech time builds in reference to a specific time.

The tense on antecedent and consequent do not need to agree (i.e. both past or both present) to get a reading as a generalisation, e.g.

(31) Usually, if I look wrecked in the morning I didn't get a good night's sleep.

(This is contrary to what has been assumed by many).

With generalisations, there is a tendency towards an ordering where antecedent precedes consequent. This is inviolable if the antecedent is present and the consequent (futurate) past, nearly inviolable if antecedent and consequent are present tense ((28) is a rare violation), and mildly preferred if antecedent and consequent are past tensed. The different ordering preferences attaching to past tense and present tense generalisations may be illustrated by the pair

(32) a. Usually, if someone was buried in an unmarked grave, they died a pauper.

b. ??Usually, if someone is buried in an unmarked grave, they die a pauper.

As (32b) shows, present tense generalisations impose a much stronger ordering constraint (one might find (32b) in an Edgar Allan Poe story, given his obsession with being buried alive, but in the normal course of events it is pragmatically unacceptable).

2.1.3 Present Tense Modals

Varieties of Interpretation

One of the major issues in the semantics of the modal auxiliaries is the variety of different interpretations that each auxiliary is open to. This issue is of lesser importance if one's concern is with the temporal properties of modals, but it does have some bearing on the matter. In the (descriptive) literature on modals, one usually finds some or all of three major distinctions being drawn between different modal interpretations.

1. Necessity, Possibility and Expectation A distinction familiar from modal logic is that between possibility and necessity. For example, *must* is used to describe different types of necessity, while *may* or *can* describe possibilities.

A less familiar category of modality is expectation[5], as illustrated by

(33) You ought to apologise, though you don't absolutely have to.

The *will* of futurity / prediction can also be cast in the form of a modal of expectation.

Expectation does not usually arise in formal treatments of modality, although something like it is to be found in the similarity ordering on possible worlds employed in Stalnaker-Lewis logics of counterfactual conditionals (see review in Nute 1984). A tense / modal operator corresponding to *will* is also sometimes treated in terms of picking out the actual course of future events in branching time tense logics (e.g. MacArthur 1976).

It might be thought that necessity, expectation and possibility lie on a continuous scale of probabilities, with necessity expressing a probability equal to 1, expectation a probability near to 1, and possibility a non-zero probability. While there are close connections between modality and probability (e.g. possible worlds semantics for probability calculi), this miscasts expectation, as the following example shows.

[5]Palmer (1990) emphasises the existence of a third linguistic category, apparently first noted by Twadell in 1960, though he does not give it a name.

Suppose there are ninety nine balls in jar A and one in jar B. They are mixed together and one is pulled out at random. There is, objectively, a high probability that the ball will be from jar A. Suppose someone claims that 'the ball will be from jar A', but on examination the ball turns out to be from jar B. This renders the original claim incorrect. But the fact the ball was from jar B does not falsify the high probability of it having been from jar A — it's just that the unlikely has happened, as it does every once in a while. If the original claim was about probabilities, finding the ball was from jar B would leave the initial probabilities unaltered and would not render the claim incorrect. (Of course, probabilities conditional on knowing where the ball came from would change, but the original claim whose truth we are interested in is not about these subsequently conditioned probabilities).

2. Epistemic and Non-Epistemic Another common distinction is between epistemic and non-epistemic uses of modals. Non-epistemic uses are often further subdivided into deontic and dynamic.

Examples illustrating epistemic, deontic and dynamic possibility are

(34) John may be in his office. (Epistemic)

(35) John can come in now. (Deontic)

(36) John can climb that tree quite easily. (Dynamic: subject oriented)

(37) Screws can be bought at an iron-mongers. (Dynamic: neutral)

Epistemic modals modals deal with what may or must be the case, given what is known or believed. Non-epistemic modals form a more mixed bag. Deontic modals concern what is obligatory, permitted or expected relative to bodies of moral, legal, cultural or other rules. Dynamic modals concern abilities, volitions, general possibilities, and so forth. Palmer (1990) further distinguishes dynamic modalities describing some property of the subject of the modal clause and those that describe some general property of the situation being described.

There is some debate about whether a continuous sliding scale connects epistemic to non-epistemic modalities, or whether there is a genuine difference in kind. Evidence apparently in favour of a difference only in degree are examples like

(38) I will go to the party tonight.

Does (38) report an epistemic prediction or a dynamic volition? But rather than a difference in degree, this just seems to indicate that the same state of affairs can sometimes be described either from an epistemic or a non-epistemic perspective. Conclusive evidence for a difference in degree would be states of affairs that cannot adequately be described either in clearly epistemic terms or clearly non-epistemic terms.

Evidence in favour of a difference in kind are paraphrase differences between epistemic and non-epistemic modals. Non-epistemic modals can be paraphrased using *for* complement constructions, but epistemic modals require *that* complement constructions:

(39) John may come in.

a. It is possible for John to come in. (Non-epistemic)

b. It is possible that John will come in. (Epistemic)

The paraphrases suggest that epistemic modality is a property of Lyons' (1977: 442) third order entities — facts, propositions, etc — while non-epistemic modality is a property of second order entities — actions, events, states, etc. However, it is far from clear what the differences between these kinds of entity are, and it is not obvious that this is an especially productive way of looking at the difference.

In Kripke style possible worlds semantics for modal logics, different types of epistemic and non-epistemic modalities can be modelled by using different accessibility relations on the set of possible worlds. With different accessibility relations, the same modal operator can be used to represent different kinds of modality. Kratzer (1977, 1981b) offers a linguistically sophisticated version of this approach where, in effect, the choice of accessibility relation is semantically underspecified and resolved by context.

3. Subjective and Objective A third distinction is between subjective and objective modalities. According to Lyons (1977) subjective modalities express the 'attitudes or opinions of the speaker'. A subjective epistemic modality reports on what is possible or necessary relative to the speaker's personal body of knowledge or belief, whereas an objective modality reports what is possible or necessary relative to what is generally known or believed. A subjective deontic modality involves the speaker conferring an obligation or permission (on the hearer), whereas an objective deontic modality reports obligations or permissions already in effect. Dynamic modalities appear to be exclusively objective, expressing facts independently of what the speaker feels or believes.

Temporal Properties of Present Tense Modals

1. Non-Past Eventualities Common to all present tense modals is the fact that they cannot be used to describe past eventualities unless used in conjunction with a perfect auxiliary. In

(40) John may/must/will be in London.

the state referred to must either be present or future (note that *John will be in London* can be used to describe a presently holding state, as in *John will be in London by now*). Past eventualities are referred to in

(41) John may/must/will have been in London.

though in this case the modals only permit an epistemic interpretation.

Use of the perfect auxiliary does not force reference to a past eventuality, e.g.

(42) You must have finished that essay by next Tuesday

which has a deontic interpretation (*contra* a claim in Coates 1983).

2. Present Modal Reference Another property common to all modals is that in simple sentences the modal predication always has present reference. Thus

(43) John may be in London tomorrow.

describes the present possibility of a future event. In general, it is necessary to distinguish the temporal reference of the modal from that of the embedded proposition / eventuality. The reference of the embedded proposition never precedes that of the modal (modulo the effects of the perfect).

3. Epistemic / Non-Epistemic Differences There are a number of temporal differences between epistemic and non-epistemic modalities. Epistemic modals can generally describe future, present, or (with the help of the perfect auxiliary) past eventualities. Non-epistemic modals mostly apply only to future eventualities.

Non-epistemic modalities concern permissions, obligations, volitions, abilities and so forth. All of these are directed to future actions (one cannot act so as to change the past, or indeed the strict present, so it is pointless to describe present obligations, abilities or volitions directed to the past). It is therefore hardly surprising that non-epistemic modalities generally describe future eventualities. The open question is whether (a) the restriction to future eventualities is a pragmatic consequence of the action based nature of non-epistemic modalities, or (b) a semantic prohibition on reference to present or past eventualities is what brings about a non-epistemic reading in the first place.

There are in any case some exceptions to this general pattern. The *can* of perception, e.g.

(44) I can see the sea!

is used to describe a present state of affairs. Objective deontic modality, usually expressed by semi-modals like *have to, be allowed to*, permits description of present states of affairs:

(45) I have to / am allowed to be here now.

(46) (?) I must / may be here now.

Unlike subjective deontic modalities, where the speaker confers a permission that subsequently permits certain actions, in an objective deontic modality the permission is already in existence and so may licence the present state of affairs.

On the epistemic side, there is a general problem in getting epistemic *must* to describe future eventualities, e.g.

(47) John must go to London (typically non-epistemic)

However, the reason for this is probably purely pragmatic. Sod's Law can usually be invoked to the effect that the completely unexpected may always happen, and stop any near certainty short of being a full necessity. Confirming the pragmatic nature of the restriction,

(48) The coin must come up heads sooner or later.

(perhaps uttered by an addicted gambler) counts as a futurate epistemic necessity.

Effect of Present Tense Modals on Subordinate Tenses

Both epistemic and non-epistemic modals can lead to deictic shift in subordinate tenses, e.g.

(49) By 1998, everybody will know someone who died from AIDS

(50) Next week, you must show me a problem that you solved on your own

In the two sentences above, the subordinate past tense has a past-in-the-future reading, similar to that found in some conditional consequents.

There is an important asymmetry between subordinate past and present tenses. Subordinate past tenses always describe an eventuality preceding that of the main clause eventuality. But with subordinate present tenses no relative order between the main and subordinate eventualities is imposed. Thus, compare

(51) One day I will marry someone who got rich quick.

(52) One day I will marry someone who gets rich quick.

In (51) the person I marry must first get rich, but in (52) the person may get rich either before or after marriage. A clearer example of a subordinate present tense eventuality preceding the main clause eventuality is given by

(53) I will give a prize to whoever solves this problem in the most elegant way.

2.1.4 Conditionals With Present Tense Modal Consequents

The temporal properties of conditionals with present tense modal consequents correspond to those of simple conditionals with present tense stative consequents. The modal behaves just like a present tense stative verb. The only point of departure is that the modal has an embedded eventuality, whose temporal reference is simultaneous with or succeeds that of the modal.

When the antecedent is present tense, the consequent describes a present or future eventuality simultaneous with or following that of the antecedent, e.g.

(54) If the bimetallic strip bends, the temperature will/may/must rise.

(55) If the bimetallic strip bent, the temperature will/may/must rise.

As before, this rule is subject to the following exceptions:
(i) When the consequent is a modal perfect e.g.

(56) If the bimetallic strip bends/bent, the temperature must have risen.

(ii) When an adverbial like *beforehand* explicitly reverses the order e.g.

(57) If the bimetallic strip bends, the temperature will rise beforehand.

(iii) When the consequent eventuality is a present state (i.e. the consequent is dependent on the time of utterance) e.g.

(58) If the letter arrives tomorrow, it must already be in the post.

(iv) When the antecedent is pre-arranged, as in a planning context, e.g.

(59) If I go to the party, I will have a shower.

(where having the shower is preparation for going out, and not what happens at the party).

When the antecedent is past tense, the temporal reference of the consequent modal is nearly always to the time of utterance, as in (55). However, there are some exceptions, as in

(60) If I didn't get a good night's sleep, I will usually be grumpy in the morning.

There are a few points that need emphasising with respect to present tense modalised conditionals. First, in simple conditionals the temporal reference of the consequent may sometimes *follow* that of the antecedent. This is never the case with modalised conditionals. Compare

(61) a. If the strip bends, the temperature rises.

b. If the strip bends, the temperature may rise.

In (61a) the consequent may refer to a time succeeding the point at which it is verified that the strip is bending. Consequently, the temperature rise follows the strip bending. In (61b) the possibility of a temperature rise comes into being as soon as the antecedent event occurs[6].

Second, with the exception of habitual conditionals like (60), past tense consequents never take on a futurate interpretation. This is just the same as for simple conditionals. However, given the ability of modals to produce past-in-the-future interpretations for subordinate past tenses, as in (51), this fact may strike one as a little surprising. It suggests that the modal only has scope over the consequent and not over the entire conditional, contrary to what most writers have assumed (e.g. Kratzer 1979, Stump 1985, Dudman 1991; see Chapter 4).

Third, modalised conditionals differ in meaning from their simple counterparts. In large measure, this is due to the modal contributing an expression of possibility, ability, permission or whatever. However, the effects of the modal *will* are more subtle.

(65) a. If the strip bends, the temperature rises. (*Rule stating*)

b. If the strip bends, the temperature will rise. (*Rule applying or stating*)

Without the modal, the simple conditional expresses a general connection or rule. But with the modal *will*, the conditional may also be used to express a particular connection between two specific events.

Fourth, whether the modal is epistemic or non-epistemic seems to make no difference to the temporal properties of the conditional. This needs to be qualified in that epistemic and non-epistemic uses of modals sometimes impose their own constraints on the temporal reference of the embedded eventuality. But these extra constraints mesh smoothly with the more general temporal properties of the conditional. It is worth pointing out, though, that in the LOB corpus the vast majority of conditionals have epistemic *will* or *may* in their consequents.

---

[6] For many conditionals with present antecedents, it seems to make little practical difference whether the modal has present or future reference. For example

(62) If Huddersfield are knocked out in their FA Cup replay with Barnsley this afternoon, they may be forced to part with their English international left-back, Ramon Wilson.

may plausibly be understood as saying either (a) that it may currently be the case that if Huddersfield lose they will be forced to part with their left-back, or (b) that if Huddersfield lose then at that point there is a possibility of them being forced to part with their left-back. However, there are other cases where the difference is clearer. Consider

(63) If I pass my test next week, I can drive any kind of vehicle I like.

The following defence would not stand up in court: Tomorrow I am stopped for driving without a proper licence, and next week I pass the test; passing the test retrospectively justifies my driving without a licence. Clearly, (63) describes an ability that comes into being next week, and not one that applies now dependent on what happens next week. By contrast,

(64) If the letter arrives tomorrow, it must already be in the post.

describes what must now be the case, though dependent on something else being verified in the future.

Conditionals with Modal Antecedents

It is sometimes claimed that epistemic modals may not occur in the antecedents of conditionals. Thus, the *will* in

(66) If John will go to the party, I'll be very surprised.

describes John's willingness to go. Lyons (1977), Haegemann and Wekker (1983) and Nieuwint (1986) have all pointed out that epistemic modals can sometimes occur in conditional antecedents:

(67) If it will rain, we should take our umbrella.

(68) If you will be in receipt of an LEA grant, complete section A of the form.

(69) If I may die tomorrow, I may as well have another drink now.

Both Lyons and Haegemann and Wekker claim that this occurs only under limited circumstances. Lyons suggests that it is only objective epistemic modals that have been absorbed into the propositional content of the antecedent clause that can occur in this configuration. Haegemann and Wekker suggest that the possibility only arise given a syntactic structure where the conditional acts as a peripheral modifier of the main consequent clause. Peripheral conditionals include such sentences as

(70) There are some biscuits in the cupboard, if you are hungry

Nieuwint argues, convincingly, that (a) conditionals with epistemic antecedents need not be peripheral, and (b) that there are no formal constraints on the occurrence of epistemic modals in conditional antecedents. At most, there is a pragmatic limitation. When an epistemic modal occurs in an antecedent, the consequent is dependent on the present possibility, necessity or expectation of a future event. It is not dependent on the future event itself. Pragmatically, it is somewhat unusual for one event to depend on the possibility of another event, rather than on the other event itself. The same argument also explains why epistemic modals do not usually occur in temporal clauses, e.g.

(71) ??After you will leave, I will leave.

For here, my leaving comes after the point at which you are predicted to leave, which precedes the time at which you actually leave.

2.2 Narrative and Conditional Temporal Reference

It has occasionally been suggested that conditionals express causal relations, where antecedent causes consequent (discussed in Comrie 1986a). This would entail an antecedent precedes consequent ordering, since causes cannot precede effects. Reversed, explanatory readings for conditionals like

(6) If the bimetallic strip bent, the temperature rose.

provide immediate and conclusive evidence that such a proposal is false.

Given the falsity of this proposal, the next most obvious treatment is to assume that conditionals merely say that it is possible to reason from antecedent to consequent. Whether the connection between antecedent and consequent licencing the inference is causal/predictive, causal/explanatory, logical, statistical or anything else is determined by context and is not part of the semantics of the conditional.

This 'null hypothesis' about the semantics of conditionals is essentially correct. However, in its bare form it poses problems about the kind of tense dependent ordering constraints between antecedent and consequent observed previously.

One possible way of explaining these ordering constraints is to say that they have nothing to do with the conditional. Instead, they may be derived from general principles determining the temporal ordering of clauses in connected discourse. Rather than try to pursue parallels between the conditional and temporal connectives like *before, after* or *when*, one treats it in parallel to a conjunction like *and*. It is well known that in a sentence like

(72) John stood up and Mary hit him.

there is a strong tendency to view the event described by the first clause as preceding that described by the second clause.

This is *not* a promising approach. While theories of temporal structure in narrative may have something to contribute to the temporal relations between antecedent and consequent clauses, it is not clear that they have anything useful to add about why and when present and past tenses in conditionals can be futurate. Worse, it does not even provide a satisfactory explanation of the difference in permitted orderings between (6) and

(73) If the bimetallic strip bends, the temperature rises.

as we will see.

Progression of Reference Time

Partee (1984), Hinrichs (1986) and Dowty (1986) invoke the progression of Reichenbach's (1947) reference time to account for narrative ordering. Reichenbach's account of tense assigns the relations between speech (S), event (E) and reference (R) times shown in Table 2.1. For simple tenses the event time equals the reference time, whereas for perfect tenses the reference time offers a perspective on an earlier event time.

Present:	$E = R = S$
Past:	$E = R < S$
Present Perfect:	$E < R = S$
Past Perfect:	$E < R < S$

Table 2.1: Reichenbachian Tenses

Dowty assumes that each succeeding clause in a narrative is interpreted relative to a new reference time occurring just after the previous one. Strictly speaking, this leads to constant temporal progression, so that each past tense eventuality temporally succeeds the previous one. The lack of apparent progression where stative clauses are involved is due to the fact that the full duration of a state can extend beyond its event and reference time. Thus, although successive states, such as

(74) John was tired. He was hungry. He was angry

have non-overlapping reference times, the states themselves may nevertheless overlap.

This account of narrative progression could serve to explain why the state overlaps with the event in

(75) If John went to bed, he was tired.

and why there is precedence between the two events in

(76) If the emperor lowered his thumb, the christian was thrown to the lions.

However, it cannot also explain the reverse ordering found in

(77) If someone was buried in an unmarked grave, then they died a pauper.

Defaults and Temporal Progression

Lascarides and Asher (1991) and Lascarides, Asher and Oberlander (1992) offer a rather more promising account of narrative structure. They start from the observation that in a discourse like

(78) Max fell. John pushed him.

the surface order is the reverse of the temporal order. It is proposed that two sets of defeasible rules operate. The first concern a Gricean pragmatic maxim of Manner, namely be orderly: by default, the order in which the clauses occur should reflect the order in which events occur. The second is a set of default causal rules, saying such things as 'if A usually causes B and A and B occur contiguously in a narrative, interpret A as preceding B'. In some cases, causal defaults will conflict with the ordering default, leading to a reversed interpretation. In other cases, such as

(79) The bimetallic strip bent. The temperature fell.

Lascarides and Asher show that two possible extensions of the default theory are possible, giving surface and reverse surface order.

This kind of account explains why reverse order past tense conditionals are possible, but generally dispreferred. The big question is whether it can account for the fact that reverse order present tense conditionals are much more strongly dispreferred, if not usually impossible. That is, can it explain why

(32a) If someone was buried in an unmarked grave, then they died a pauper.

is acceptable, but

(80) ?If someone is buried in an unmarked grave, then they die a pauper.

is not?

Present Tense Narrative Most accounts of narrative structure have only dealt with past tense narratives. While present tense narratives are uncommon, roughly the same principles apply. Consider the following discourse describing a play or film

(81) It's a great story. This wealthy Prince travels to a foreign country.... And finally the Prince is buried in an unmarked grave. He dies a pauper, you see.

In this present tense narrative, a reverse reading for the last two sentences is possible, just as it is for

(82) The Prince was buried in an unmarked grave. He died a pauper, you see.

There is thus nothing in the narrative structure that prevents an order reversing reading for present tense clauses. And so, appeals to general principles of narrative temporal structure cannot explain the unacceptability of (80) and at the same time the acceptability of (80).

2.3 Dowty and Deictic Shift

Dowty (1982) recognises the fact that a modal like *will* can give rise to deictically shifted past-in-the-future tenses in subordinate clauses. He proposes a semantic analysis intended to account for this, employing a double indexed tense logic:

1. Basic Expression: A valuation function V assigns the value 1 (true) or 0 (false) to each basic (i.e. untensed) expression α for each time e.

   $[\![\alpha]\!]^{s,e} = V(\alpha, e)$

   (i.e. the speech time index, s, plays no role in determining the interpretation of basic propositions).

2. $[\![PAST\, \phi]\!]^{s,e} = 1$ iff $[\![\phi]\!]^{s,e'} = 1$ for some $e' < s$

3. $[\![PRES\, \phi]\!]^{s,e} = 1$ iff $[\![\phi]\!]^{s,e'} = 1$ for some $e' = s$

4. $[\![FUT\, \phi]\!]^{s,e} = 1$ iff $[\![\phi]\!]^{e',e'} = 1$ for some $e' > s$

The past and present tense operators, $PAST$ and $PRES$, select a new value e' for the event time index that respectively precedes or is simultaneous with the speech time index s. The speech time index is left unaltered. By leaving the speech time index unchanged, subordinate tenses occurring within the scope of a superordinate $PAST$ or $PRES$ operator receive ordinary, unshifted interpretations.

The following equivalences therefore hold

$$\begin{aligned}
PAST(\phi \wedge PAST\, \psi) &\equiv PAST\, \phi \wedge PAST\, \psi \\
PAST(\phi \wedge PRES\, \psi) &\equiv PAST\, \phi \wedge PRES\, \psi \\
PRES(\phi \wedge PAST\, \psi) &\equiv PRES\, \phi \wedge PAST\, \psi \\
PRES(\phi \wedge PRES\, \psi) &\equiv PRES\, \phi \wedge PRES\, \psi
\end{aligned}$$

Assuming that sentences containing relative clauses can be represented schematically as shown,

(83) A woman stole the book that John bought.
$PAST(\phi \wedge PAST\, \psi)$

(84) A woman stole the book that John likes.
$PAST(\phi \wedge PRES\, \psi)$

this formulation correctly predicts that the subordinate tenses will not be deictically shifted. That is, in (83) the buying and the stealing of the book both occur in the past, but the stealing may either precede or follow the purchase. In (84), John likes the book now.

Matters are different with the future tense operator FUT, which is meant to correspond to the modal *will*. Not only is the event time index changed to be some time following the original speech time s, but the speech time index is also changed to be identical to the new event time. This means that subordinate tenses within the scope of the FUT operator receive a shifted interpretation: they are evaluated relative to the futurate value of the speech time index. In

(85) By 1998, everybody will know someone who died from AIDS.
$FUT(\phi \wedge PAST\,\psi)$

the subordinate past tense selects a time that is past relative to the future time of everybody knowing someone.

Since other modals exhibit the same shifting properties as *will*, we could define them along the same lines as FUT, though presumably we would have to include an extra possible world index to take account of the modality. In the same vein, we could also define the effects of the conditional as follows

5. $[\![IF(\phi, \psi)]\!]^{s,e} = 1$ iff $[\![\phi]\!]^{e',e'} = 1$ implies $[\![\psi]\!]^{e',e'} = 1$ for some $e' > s$

That is, the conditional acts as though it were within the scope of a FUT operator. This would predict that in

(86) If I smile when I get out, the interview went well.
$IF(PRES\,\phi,\ PAST\,\psi)$

the antecedent present tense refers to some time in the future, and the consequent past tense refers to some time preceding it.

Unfortunately, this treatment fails to predict certain asymmetries in the behaviour of deictically shifted past and present tenses. Consider the behaviour of subordinate present tenses within the scope of a modal, e.g.

(87) I will give a prize to the person who solves this problem most elegantly.
$FUT(\phi \wedge PRES\,\psi)$

According to the definitions given above, the event time for ψ must be identical to that for ϕ. That is, the presentation of the prize is exactly simultaneous with the solution of the problem. Yet it is clear that that (87) allows an interpretation where the solution of the problem precedes the award of the prize. And in other cases, such as

(88) One day, I will get to know someone who makes a killing on the stock exchange.

the subordinate present tense clause may refer to an event that succeeds the superordinate event (so that I get to know someone who then proceeds to make a killing on the stock exchange).

Thus, while Dowty's FUT operator correctly predicts what happens to subordinate past tenses, it makes incorrect predictions about subordinate present tenses.

In addition, my extension of Dowty's analysis to an IF operator fails rather more catastrophically. For conditionals with present tense antecedents and consequents, $IF(PRES\,\phi,\ PRES\,\psi)$, it incorrectly predicts that the antecedent and consequent event times must be simultaneous. For conditionals with past tense antecedents and consequents,

$IF(PAST\,\phi,\ PAST\,\psi)$, it predicts that the antecedent and consequent events may be past relative to some time in the future. We could alter the definition of IF to

5'. $[\![\ IF(\phi,\psi)\ ]\!]^{s,e} = 1$ iff $[\![\ \phi\ ]\!]^{e'',e''} = 1$ implies $[\![\ \psi\ ]\!]^{e',e'} = 1$ for some $e' \geq s$ and some $e'' \geq e'$

This would at least not force simultaneity between present antecedents and consequents, while forcing an antecedent precedes consequent order. But this would still not account for the fact that past tenses are not permitted deictically shifted readings in $IF(PAST\,\phi,\ PAST\,\psi)$ conditionals. It would also then fail to account for the consequent precedes antecedent reading of $IF(PRES\,\phi,\ PAST\,\psi)$ conditionals.

By employing just one deictic centre, Dowty falls into the problems noted in Chapter 1.

2.4 Dudman on Conditionals

Dudman (Dudman 1983, 1984a, 1984b, 1985; summarised in Dudman 1991) is one of the few people to have paid attention to the temporal properties of conditionals. He has argued that an exclusively temporal classification of different types of conditional is the only kind of classification that makes sense. While having much sympathy with Dudman on this point, his classification is nevertheless flawed in certain important respects.

2.4.1 Dudman's Temporal Classification

Dudman distinguishes three main categories of conditional, which I will call compound, habitual and real[7]. Extra, marginal categories of conditional are possible, though these are not discussed at length; e.g. conditionals like *If Mary was fat, her brother was immense* and *If you are hungry, there are some biscuits in the cupboard.*

Type 1: Compound Conditionals

The first type of conditional is compounded from two independently tensed clauses. Each clause must have exactly the same interpretation as it would do if it stood independently as a sentence (the 'stand-alone' test).

(6) If the bimetallic strip bent the temperature rose.

(89) If the bimetallic strip bent the temperature will rise.

[7]Dudman uses slightly different terminology, reserving the expression 'conditional' exclusively for what I have rebaptised as real conditionals.

Both (6) and (89) count as compound, since the antecedent and consequent clauses (*the bimetallic strip bent*, *the temperature rose* and *the temperature will rise*) have the same stand-alone meanings as they do in the conditional sentences.

A conditional like

(73) If the bimetallic strip bends, the temperature rises.

is not compound, since the futurate antecedent taken in isolation (*the bimetallic strip bends*) means that the strip has a present propensity towards bending.

Given that both antecedent and consequent clauses express exactly the same proposition standing in isolation as in the conditional, compound conditionals support instances of *modus ponens* style reasoning. Although Dudman does not explicitly say what the 'logical form' of compound conditionals should be, he does say that they are syntactically like conjunctions. It is almost certain that the intended logical structure is something like

$$tense(A) \rightarrow tense(C)$$

where $\rightarrow$ represents the conditional connective, and $tense(A)$ and $tense(C)$ the independently tensed antecedent and consequent clauses.

Since the antecedent and consequent tenses are independent, any mixture of tenses is possible. This is the only type of conditional that Dudman allows to have mixed tenses.

Type 2: Habitual Conditionals or Generalisations

In Dudman's second category, the conditional is adverbial, and describes a general connection between antecedent and consequent. Examples of generalisations are

(90) If the alarm sounds, I (usually) ignore it.

(91) If the alarm sounded, I (usually) ignored it.

(92) Up until then, if the alarm had sounded I had usually ignored it

The antecedents (and consequents) do not have the same stand-alone meanings as they display in the conditionals. This is taken as an indication that the neither the antecedent nor consequent carries a genuine and independent tense. Instead, there is a single tense with scope over the entire conditional. The tense markings on antecedent and consequent are just surface reflections of the single wide scope tense.

To the extent that one can judge the logical structure that Dudman would assign to generalisations, it is something like

$$tense(H(A \rightarrow C))$$

where H is some kind of generalisation or habitual operator, and *tense* is the tense of the entire conditional. The tense determines the time at which the generalisation is valid (i.e. a presently valid generalisation, or a generalisation true in the past). Through this, the tense indirectly determines the times at which the antecedent and consequent events may occur, but there is no direct connection.

Since the antecedent and consequent clauses are not independently tensed, but reflect the tense of the conditional as a whole, it follows that mixed tense generalisations are not permitted.

Type 3: Real Conditionals

Dudman's third category groups together such sentences as

(93) If John misses the bus, he will walk home.

(94) If John missed the bus, he would walk home.

(95) If John had missed the bus, he would have walked home.

The two defining features of this class are (i) that the consequent is modal, and (ii) that the tenses refer to times later than would normally be associated with their surface forms. Thus, in (93) the present tense antecedent refers to the future, as does the past tense antecedent in (94). It takes a past perfect antecedent (past past) to refer to a past time, as in (95).

As with generalisations, real conditionals are adverbial, and tense agreement is required between antecedent and consequent — *will* counts as a present tense modal, *would* as a past tense modal, and *would have* as a past perfect modal. Again, neither antecedent nor consequent have their own independent tenses; there is just one tense with wide scope over the entire conditional. Dudman assigns logical forms like

$$\text{Fut–Will}(A \rightarrow C)$$
$$\text{Pres–Will}(A \rightarrow C)$$
$$\text{Past–Will}(A \rightarrow C)$$

to (93), (94) and (95).

The tense on the modal refers not to the times at which the antecedent and consequent events occur, but to the time at which a 'projective fantasy' about those events begins.

With past perfect modals, the projective fantasy begins at some past point in history. That is, one goes back in time to some past state of affairs, forgetting everything you know about what happened between that past point and the present. From that past perspective, one then evaluates the conditional. This may involve going forward in history again, though quite probably not along the same path as one took in going back. One may even go beyond the present time in this projection, as happens in Dudman's example

(96) If the auditors had come tomorrow, they would have found the accounts in order.

Suppose (96) is uttered at a time when the auditors have already been and found everything in a mess. One goes back to a time before the auditors arrived, and possibly before they even said when they would arrive. From this, one works forward again, wondering what would have happened if they had arrived tomorrow, where to 'tomorrow' still refers to the day after the time of utterance rather than the day after the start of the projective fantasy.

With past modals, the fantasy begins at the the present time. The antecedent event occurs immediately and it may be necessary to drop some present facts in the history incompatible with the occurrence of the antecedent. Thus in

(97) If John was here now, he would be enjoying this

the fantasy has to ignore the fact that John is actually somewhere else right now.

With present modals, the fantasy begins at some point in the future. The course of history up to the present remains sacrosanct, but its future development up to the occurrence of the antecedent and beyond is largely the choice of the fantasist.

The claim that Dudman attaches most weight to is that his third category of conditionals cuts right across the traditional divide between indicative and subjunctive / counterfactual conditionals. Traditionally, the sentence (93) would be classified as indicative, whereas (94) and (95) would be classed as subjunctive.

2.4.2 Problems with Dudman's Classification

The problems with Dudman's classification can be traced back to his use of the stand-alone test. It is assumed that if an apparently tensed antecedent or consequent clause fails the stand-alone test, then that clause cannot be tensed in the normal way; instead, the surface tense reflects a single, wide scope conditional tense. Dudman also assumes that the present tense may only have present time reference, and cannot be used (in a stand-alone way) to refer to the future.

Taking failure of the stand-alone test to indicate the absence of genuine tensing neglects the possibility of tenses being sensitive to their semantic context; in effect, it denies from the outset the possibility of deictic shift. If one allows the semantic context (as set up by a modal or conditional) to alter the deictic centre, subordinate tenses may take on an interpretation different from the one they would exhibit in a different context. It is not the semantics of the tense that has changed, merely the context in which it is evaluated. It is premature to assume that, just because a clause is interpreted differently in a stand-alone from a conditional context, the stand-alone clause is semantically different from the conditional clause.

Mixed Tense Conditionals

According to Dudman's classification of conditionals, the following sentences should be unacceptable, but they are not:

(98) If the bimetallic strip bends, then the temperature rose.

(99) If I smile when I get out, the interview went well.

(100) Usually, if I look wrecked in the morning, I didn't get a good night's sleep.

In all of these the antecedent fails the stand-alone test, and in (99) and (100) the past tense consequents also fail the test.

According to Dudman, failure of the stand-alone test means that there is a a single conditional tense, and consequently the surface tenses of the antecedent and consequent should agree. Yet in all three cases the antecedent is marked as present tense, and the consequent as past tense.

Taking mixed tense to be a sign of a compound conditional, (98) could be accommodated if the present tense was understood as simply indicating non-pastness rather than presentness. However, it would have to be explained why only a present habitual reading is possible under normal circumstances for *The bimetallic strip bends*, when a futurate reading might be expected.

But this change will not suffice to accommodate (99). It is hardly plausible that we would wish to revise the semantics of the past tense so that it means 'past or past-in-the-future'. If (99) is to be construed as compound, we have to allow for the possibility that the antecedent and conditional may alter the context of interpretation for the consequent tense. But if this is so, then the stand-alone test ceases to be very informative about the tense structure of conditionals, or at least of conditional consequents.

Worse still, (100) is clearly a habitual conditional. As such, it casts considerable doubt on the claim that all habitual conditionals possess a single wide scope conditional tense, and have semantically untensed antecedents and consequents whose surface form reflects the wide scope tense.

Ordering in Habituals

There are other problems with Dudman's treatment of habitual conditionals. Though he is not explicit, it is clear that he assumes that there is always an antecedent precedes consequent ordering in habitual conditionals (Hornstein (1991) is explicit on this point in reviewing Dudman's classification). This raises two questions. First, how is this ordering imposed? And second, what about the exceptions to it occurring with past habituals and, rarely, with present habituals?

Since both antecedent and consequent are untensed in Dudman's treatment of habituals, one could always impose an ordering over antecedent and consequent by having the conditional itself advance the value of the event time. Unfortunately, by fixing the temporal ordering in the conditional in this brute force way, we are powerless to explain counterexamples to the antecedent precedes consequent ordering. We are still less able to explain why counterexamples are more common in past generalisations than present tense ones:

(101) If someone was buried in an unmarked grave, they usually died a pauper.

(102) ?If someone is buried in an unmarked grave, they usually die a pauper.

(103) If the Prime Minister gives a good speech, someone else usually writes it for him.

Indicative and Subjunctive

Dudman's primary reason for denying an indicative / subjunctive distinction is that the all the tenses in the conditionals (93)–(95) have non-standard interpretations, where the 'time referred to is later than the time indicated by surface form'.

From the foregoing discussion, there is no reason why (93) should not be construed as an ordinary compound conditional, where the present tense antecedent takes on exactly the same kind of futurate reading as in (98). Given the similarity in meaning between (93) and (104)

(93) If John misses the bus, he will walk home.

(104) If John misses the bus, he would walk home.

there are positive reasons for viewing (93) as compound: (104) has mixed tense according to Dudman, and so can only be a compound conditional.

In Chapter 6 I will argue that the past tense antecedents in (94) and (95) are not deictically shifted in the same way that the present tense antecedent in (93) is. That is, there *is* a distinction between indicative and subjunctive conditionals, and it can be motivated on purely temporal grounds.

Conjunctive and Adverbial Conditionals

Dudman makes some widely shared assumptions about the syntactic and logical structure of conditionals. These are (i) conditionals come in two forms, conjunctive (i.e. compound) or adverbial, and (ii) in adverbial conditionals, syntactic tense agreement between antecedent and consequent is required, and indicates the absence of any semantic tense on the antecedent and consequent clauses. Put another way, conjunctive conditionals have the logical form

$$tense(A) \rightarrow tense(C)$$

whereas adverbial conditionals have the form

$$tense(A \rightarrow C)$$

Moreover, conjunctive conditionals are taken to correspond to modus ponens supporting specific conditionals, whereas adverbial conditionals correspond to habitual or other kinds of modalised conditional.

Chapter 3 proposes a different analysis (see also Appendix A). The distinction between adverbial and conjunctive conditionals is preserved. But while conjunctive conditionals have the same logical form as above, adverbial conditionals are taken to have the following form

$$tense_c[tense_a(A) \rightarrow C]$$

Three assumptions are made here. First, conditional adverbials are formed from the word *if* and an independently tensed antecedent clause. Second, the adverbial applies at the sentential level, though below the level at which the tense applies. Third, that tense applies at the sentential level rather than the verb phrase level: following Nerbonne (1985), a sentential tense feature percolates down to the verb phrase where the tense receives its surface realisation, but semantically the tense applies to the sentence and not the verb phrase. In adverbial conditionals, the adverbial modifies a semantically untensed clause, with the sentence's tense applied to the modified clause.

2.5 Data Semantics

Data Semantics (Veltman 1984, 1985, 1986; Landman 1986) is very similar to the first treatment of conditional update described in Chapter 1. The treatment of conditionals and modals in the next chapter and beyond was greatly influenced by Data Semantics, though it differs from it in a number of respects.

2.5.1 Basic Data Semantics

Veltman (1986) starts off with a language L consisting of atomic sentences, three one-place sentence operators $\neg$, **must** and **may**, three two-place operators $\wedge$, $\vee$ and $\rightarrow$, and parentheses, with the usual formation rules. The semantics for this language is provided by an information model, which is a triple $\langle S, \sqsubseteq, V\rangle$ where S is a non empty set of information states, $\sqsubseteq$ is a partial order of information extension over S, (where each maximal chain in $\langle S, \sqsubseteq\rangle$ contains a maximal element—see below), and V is a partial valuation function from information states and atomic sentences into the truth values 1 and 0.

For a given state, s_1, the valuation will assign truth values to some atomic sentences in that state, while it will leave others undefined. When $s_1 \sqsubseteq s_2$, then where the valuation function V assigns a value to an atomic sentence in s_1, it will assign the same value to that sentence in s_2. But there may be more sentences for which V is defined in s_2 than there are in s_1. $s_1 \sqsubseteq s_2$ thus means that state s_1 can grow into s_2 by the addition of further atomic facts, while preserving all those atomic facts supported in s_1. The requirement that each maximal chain in $\langle S, \sqsubseteq\rangle$ contains a maximal element simply says that any information state should be able to grow into a state of complete information where V assigns a value 1 or 0 to every atomic sentence in L.

The language L is interpreted against an information model as follows. We let $s \vdash_M \phi$ abbreviate 'ϕ is true / supported in model M in state s, and $s \dashv_M \phi$ abbreviate 'ϕ is false in model M in state s'. Then

> If p is atomic, then
> $s \vdash_M p$ iff $V(s,p) = 1$
> $s \dashv_M p$ iff $V(s,p) = 0$
>
> $s \vdash_M \neg\phi$ iff $s \dashv_M \phi$
> $s \dashv_M \neg\phi$ iff $s \vdash_M \phi$
>
> $s \vdash_M \phi \wedge \psi$ iff $s \vdash_M \phi$ and $s \vdash_M \psi$
> $s \dashv_M \phi \wedge \psi$ iff $s \dashv_M \phi$ or $s \dashv_M \psi$
>
> $s \vdash_M$ **may**(ϕ) iff for some $s' \sqsupseteq s, s' \vdash_M \phi$
> $s \dashv_M$ **may**(ϕ) iff for no $s' \sqsupseteq s, s' \vdash_M \phi$
>
> $s \vdash_M$ **must**(ϕ) iff for no $s' \sqsupseteq s, s' \dashv_M \phi$
> $s \dashv_M$ **must**(ϕ) iff for some $s' \sqsupseteq s, s' \dashv_M \phi$
>
> $s \vdash_M \phi \rightarrow \psi$ iff for no $s' \sqsupseteq s, s' \vdash_M \phi$ and $s' \dashv_M \psi$
> $s \dashv_M \phi \rightarrow \psi$ iff for some $s' \sqsupseteq s, s' \vdash_M \phi$ and $s' \dashv_M \psi$

An atomic sentence is true in a state s if the valuation functions assigns the value 1 to the sentence in that state. It is false if the valuation function assigns 0. And otherwise, the truth value of the sentence is undefined, meaning that it is not known whether it is true or false. A conjunction is true if both its conjuncts are true, false if either of the conjuncts are false, and undefined otherwise. The negation of a sentence is true if the sentence is false, false if the sentence is true, and undefined if the sentence is undefined.

The circumstances under which modal and conditional sentences are true or false in a state s are a little more complicated. They generally involve looking at other information states that are possible extensions of s, i.e. states into which s might grow. The modal **may**(ϕ) is true in s if there is at least one state into which s can grow in which ϕ is true. Otherwise **may**(ϕ) is false. The modal **must**(ϕ) is true in s is true if there are no states into which s can grow that make ϕ false. Since any information state can be extended to a maximal information state in which all sentences are either true or false, if **must**(ϕ) is true in s, then ϕ is true in all maximal states into which s can grow.

Had data semantics included a modal operator for **will** it would need to be defined along the following lines: **will**(ϕ) is true in s if ϕ is true in one of the states into which s actually grows. Of course, there are major problems in picking out the actual course of s's development from all the possible ones. The modals **must** and **may** on the other hand do not require that the actual course of information growth be distinguished from all the possible ones.

The conditional $p \rightarrow q$ is true in state s, provided that every possible course of information growth stemming from s that contains a state where p is true, contains a subsequent state in which q is true. In other words, the conditional is true if, whenever your information grows in such a way as to make p true, it will always grow further so as to make q true.

This may not be immediately obvious from the semantic definition, which says only that $p \rightarrow q$ is true in s if s cannot grow into a state s' where p is true and q is false. Confining our attention to cases where both p and q are non-modal and non-conditional, if either p or q become true (or false) in an information state s', they will remain true (or false) in all states informationally greater than s', these being the states into which s' may grow. The valuation function V and the ordering $\sqsubseteq$ are mutually constrained to ensure that this is so [8]. So suppose that s grows into a state s' where p is true. In this state, q may be either false, true or undefined. If q is false, the conditional $p \rightarrow q$ is falsified. If q is true, both p and q will remain true in all states into which s' may grow, and so no course of growth stemming from s' will provide a counterexample to the conditional. Suppose q is undefined. The condition on maximal chains of states ensures that sooner or later s' will grow into a state where q is either true or false. Suppose that s' can grow into a state s''

[8]As Veltman observes, modal and conditional formulae do not have truth values that are stable under information growth. As we gain more information, possibilities may be eliminated, and previous possibilities may become inevitabilities. But since truth values are always defined for modals and conditionals in all states, it turns out that the alternative description of $\rightarrow$'s truth conditions also applies when either of p or q are modal or conditional.

where q is false. p will still be true in this state, and so s'' is a state into which s can grow (passing via s') in which p is true and q is false, thus falsifying $p \rightarrow q$. If the conditional is not to be falsified, therefore, any state where p is true and q is undefined. must eventually grow into a state where q is true. This means that if the conditional $p \rightarrow q$ is true in s, then any state s' into which s' may grow will either be one where q is already true, or will grow further into one where q is true.

Iterated Modalities

It turns out that the basic data semantics presented above validates some rather strange equivalences (Landman 1986, Veltman 1985). Two are particularly worthy of note:

$$\mathbf{must}(\mathbf{may}(\phi)) \equiv \mathbf{must}(\phi)$$
$$(\phi \rightarrow \mathbf{may}(\psi)) \equiv (\phi \rightarrow \psi) \equiv (\phi \rightarrow \mathbf{must}(\psi))$$

These equivalences would licence the following implausible patterns of inference: (a) that John did it has to be a possibility, therefore John must have done it; (b) if John comes then Mary may leave, therefore if John comes Mary will / is bound to leave.

The reason that these equivalences hold is that modals and conditionals dictate what happens in *all* states extending a given state. Taking $\mathbf{must}(\mathbf{may}(p))$ as an example, if this is true in state s, then there is no state extending s in which $\mathbf{may}(p)$ is false. That is, the possibility of p being true can never be ruled out. Consider what happens at a total extension of s, e.g. s'. In s', $\mathbf{may}(p)$ must not be false. Since s' is total, this means $\mathbf{may}(p)$ must be true. And the only way for this to happen is if p is true in s'. Therefore p must be true in all total extensions of s. This guarantees that $\mathbf{must}(p)$ is true.

The same happens with the conditional. If s supports $p \rightarrow \mathbf{may}(q)$, then any state s' extending s in which p is true, must grow into a state in which $\mathbf{may}(q)$ is true. Suppose that s' is a total extension of s in which p is true. Since the state cannot grow any further, $\mathbf{may}(q)$ must be true in s'. And as s' is total, the only way for this to happen is if q is true in s'. So any state extending s in which p is true must ultimately grow into a (total) state where p and q are true.

Iterated modalities and the interaction of modals with conditionals is a matter to which we will return in Chapter 5. Problematic equivalences like the above are blocked if one adopts a different treatment of conditionals, corresponding to the second type of conditional update mentioned in Chapter 1. That is, the conditional specifies what would happen if the antecedent were added directly to the current information state s. It does not legislate for what happens in all states extending s in which the antecedent holds. (This is not the only form of solution, and both Veltman (1985) and Landman (1985) suggest alternative ones).

Data Lattices, Negation and Incompatibility

Veltman (1984) offers a somewhat different semantics to the one described above. Instead of taking information states as primitive objects, they are taken to be sets of facts or propositions. The complete set of facts form a semi-lattice, known as a data lattice. This structure has an operator $\bigwedge$ corresponding to the conjunction of propositions, and an ordering over propositions $\models$. If p and q are elements of the data lattice, then the conjunction $p \bigwedge q$ is also an element of the data lattice. And if $p \models q$, then p expresses as much or more information than q. The operator and order are related such that

If $p \models q$, then there is some r such that $p = q \bigwedge r$.

Included in the data lattice is an impossible fact, $\perp$. If $p \bigwedge q = \perp$, then p and q are said to be incompatible.

A data set D, which corresponds to an information state above, is a certain kind of subset of a data lattice. It must be closed under the operation $\bigwedge$. So if $p, q \in D$ then $p \bigwedge q \in D$. Put another way, if $p \models q$ and $p \in D$, then $q \in D$. In addition, $\perp \notin D$. (In more familiar terms, D is a deductively closed, consistent set of propositions).

The informational ordering over data sets / information states now becomes set inclusion, i.e. $D1 \sqsubseteq D2$ iff $D1 \subseteq D2$.

A data model for the language L consist of a data lattice and an interpretation function $\mathcal{I}$, which assigns atomic sentence letters to propositions. An atomic sentence is true relative to a data set precisely if the proposition assigned to the sentence is contained in the dataset. It is false if the data set contains some proposition incompatible with it. That is

If p is atomic, then
$D \vdash_M p$ iff $\mathcal{I}(p) \in D$
$D \dashv_M p$ iff there is some $q \in D$ such that $\mathcal{I}(p) \bigwedge q = \perp$

Complex formulas have truth and falsity conditions similar to those given previously (substitute information states s by data sets D).

No Negative Facts By exploiting incompatibility between facts, one does not need to posit the existence of negative facts. To say that something is not the case is merely to say that something else is the case that is incompatible with it.

This being so, it would be possible to provide a more intuitionistic account of negation, i.e.

$D \vdash_M \neg\phi$ iff there is no $D' \sqsupseteq D$ such that $D' \vdash_M \phi$

That is, ϕ is not the case in D if there is no way of extending D in such a way as to include ϕ. If D cannot be (consistently) extended to include D, this means that there is some fact in D incompatible with ϕ. For non-modal and non-conditional formulas ϕ, this treatment of negation has exactly the same effect as the original.

Furthermore, if negation is treated in this way, we could dispense with the negative support relation $\dashv_M$ altogether. It is only there in the first place to allow a sensible positive definition saying when a negation holds.

Veltman argues against an intuitionistic treatment of negation. The reason for this is that it would make the support conditions for a formula like $\neg\mathbf{must}(p)$ somewhat implausible. Intuitively, if it is not the case that p must be the case, then it must be possible for p to either be true or false. But according to the intuitionistic definition of negation, $\neg\mathbf{must}(p)$ holds in D if $\mathbf{must}(p)$ holds in no extension of D. Let D' be a total extension of D. If $\mathbf{must}(p)$ fails to hold in D', it can only be because p fails to hold in D'. Since D' cannot be further extended, this means that $\neg p$ holds in D'. Thus $\neg\mathbf{must}(p)$ turns out to be equivalent to $\mathbf{must}(\neg p)$! In Chapter 5, however, it will be shown how a roughly intuitionistic treatment of negation can be made to work.

Why Data Sets? One of the motivations for for using data sets instead of information states is that it reduces the number of primitive objects. Using information states, one needs to take both the states and the ordering $\sqsubseteq$ over them as primitive. Using data sets, only facts / propositions are primitive, and both information states / data sets and the ordering $\sqsubseteq$ can be defined in terms of them.

Landman (1984) has made use of data sets in an analysis of attitude reports. If John sees it raining, under this analysis, John has a relation to a data set of facts (an interpreted situation), which includes the fact that it is raining. If John sees *that* it is raining, he has a relation to a data set containing just the fact that it is raining. One could not obtain the same kind of analysis if instead of data sets one had primitive information states.

However, taking the order $\sqsubseteq$ as defined rather than primitive leads to certain problems. The ordering becomes one essentially of logical entailment. This has the result that a conditional like *If John comes, Mary will leave* has to be construed in such a way that the antecedent together with some implicit extra premise contained in the data set entails the consequent. Landman (1985) finds this implausible, since most conditionals express a contingent connection between facts rather than a logical one, and it is often far from clear what the extra implicit premise should be.

Landman therefore re-introduces a primitive, by taking a notion of incompatibility between propositions as basic. Judgements of incompatibility underlie a more contingent information ordering on propositions, so that $p \models q$ iff p is incompatible with everything incompatible with q. Data lattices need to be more constrained than Boolean algebras for this to work. And as Landman admits, there is little difference between taking incompatibility as primitive on the one hand, and an ordering of information extension as primitive on the other. He does comment that individual incompatibilities are often easier to recognise than a general ordering, and so may be a more plausible primitive. But this observation is weakened by the fact that one has to determine all the propositions incompatible with a given proposition in order to derive an informational order.

There seems little to be gained by treating information states as data sets of propositions rather than as primitive entities, with a primitive ordering defined on them.

2.5.2 Adding Time and Tense

The most natural way of adding time and tense to data semantics proceeds, and fails, along the lines mentioned in Chapter 1. I will therefore be brief about how this happens.

We need to extend the language L to include two tense operators **past** and **pres**. To ensure that these operators stand at least some chance of producing sensible results, it will be assumed that they are double indexed. That is, a formula counts as true or false relative to an information state, an update time (t) and an event time (e). The salient semantic definitions are as follows

If p is atomic
$s, t, e \vdash p$ iff $V(s, p, e) = 1$
$s, t, e \dashv p$ iff $V(s, p, e) = 0$

$s, t, e \vdash \mathbf{past}(\phi)$ iff there is some $e' < t$ such that $s, t, e' \vdash \phi$

$s, t, e \vdash \mathbf{pres}(\phi)$ iff $s, t, t \vdash \phi$

$s, t, e \vdash \phi \to \psi$ iff for all s_1 and t_1 such that $s_1 \sqsupseteq s$, $t_1 \geq t$ and $s_1, t_1, e \vdash \phi$, it is not the case that $s_1, t_1, e \dashv \psi$.

An alternative to the definition for $\to$, which spells out what is required for a formula not to be falsified in states like s_1, would be

$s, t, e \vdash \phi \to \psi$ iff $\forall s_1, t_1$ such that $s_1 \sqsupseteq s$, $t_1 \geq t$ and $s_1, t_1, e \vdash \phi$,
then $\forall s_2 \sqsupseteq s_1$: $\exists s_3 \sqsupseteq s_2$ and $\exists t_3 \geq t_1$ such that $s_3, t_3, e \vdash \psi$

Sentences are evaluated only relative to an information state and a speech / update time. For this evaluation to proceed in terms of an evaluation relative to a state, an update time and an event time, the event time is initially set to be the same as the update / speech time for the sentence. That is $s, t \vdash S$ iff $s, t, t \vdash S$.

As observed in Chapter 1, this kind of treatment mispredicts that conditionals with past tense antecedents and consequents should be able to refer to future times, in the same ways as conditionals with present tense antecedents and consequents. Nor does it predict that past-in-the-future consequents must refer to time preceding those referred to by their present tense antecedents.

2.6 Conclusions

This chapter has presented a range of hitherto unremarked upon data, showing how the temporal interpretation of indicative conditionals is subject to a number of quite tight constraints. These constraints are not predicted by any of the semantic analyses of tenses, modals and/or conditionals surveyed in this chapter.

While a natural extension of Data Semantics fails along with the other analyses, it forms one of the starting points for the more satisfactory account presented in the next chapter. From the point of view of Data Semantics, the next chapter introduces two fundamental changes. First, a distinction between assertion and verification is introduced, so that information states can support both verified and unverified assertions, and where the assertions that count as unverified in a state diminish over time. Second, conditionals and modals are understood as specifying what will happen if a certain piece of information is added to the current state. They do not explicitly dictate what happens in all states extending the current one.

Chapter 3

A Semantics for Tense and Conditionals

This chapter presents a semantic analysis of the past and present tenses and two indicative conditional connectives. It is shown how this correctly predicts the range of temporal interpretations exhibited by simple conditionals, as described in the last chapter.

Section 3.1 introduces information states that support verified and unverified assertions. Section 3.2 gives semantic definitions for past and present tense operators, and Section 3.3 gives semantic definitions for two conditional connectives. Section 3.4 describes how the tenses and conditionals interact. Section 3.5 introduces a habituality operator that can be defined in terms of one of the conditional connectives.

3.1 Information States and Models

In this chapter, we will assume a language L consisting of atomic sentence letters (p, q etc, plus the absurd sentence $\perp$), three one place sentence operators ($\mathbf{past}, \mathbf{pres}, \neg$), four two place connectives ($\wedge, \vee, \rightarrow, \Rightarrow$), parentheses, and the usual formation rules. In Chapter 4, this language will be enlarged by the addition of three further one place sentence operators, **must**, **may** and **will**. In this section, we will describe the information models against which formulas in this language are evaluated.

3.1.1 Information Models

An information model M is a septuple

$$M = \langle S, E, I, \sqsubseteq_t, <, \trianglelefteq, V \rangle$$

where: S is a set of information states
E is a set of time points and periods
I is a subset of E containing only time points
$\sqsubseteq_t$ is a relation in $S \times I \times S$
$<$ is a linear order of temporal precedence on E
$\trianglelefteq$ is a relation of temporal inclusion on E
V is a valuation function for atomic sentence letters

The information states in S are primitive objects, ordered by the time varying relation of information extension, $\sqsubseteq_t$. The set of times E is the set of possible event times, and includes time periods as well as time points. The set of time points I provides possible values of the assertion and verification time indices. The ordering of temporal precedence, $<$ (described below) applies to times in E as well as in I. The ordering of temporal inclusion, $\trianglelefteq$, only relates times contained in E; it is not defined for the time points in I. For $<$, we will make free use of its standard counterparts, $>, \leq, \geq$ and $=$. $\sqsubseteq_t$ is a time dependent ordering of information extension, with variants in the form of $\sqsubset_t, \sqsupset_t, \sqsupseteq_t, \approx_t$. When $s_1 \sqsubset_t s_2$ we say that s_2 is an extension of s_1 at t. For any specified time index t, $\sqsubseteq_t$ is a transitive, reflexive, well-founded ordering on S. The valuation function V is a function from states, time indices, atomic letters and event times to the values 1 or 0.

The connection between information states and times is best introduced through consideration of the valuation function.

Valuation Function

The valuation function V is a *total* function from states, (verification) time points, sentence letters and (event) times to the values 1 and 0. When

$$V(s, t_v, p, t_e) = 1$$

—where $t_v \in I$, and $t_e \in E$

we say that the assertion that p is true at time t_e is verified in state s at time t_v. When $V(s, t_v, p, t_e) = 0$, the assertion that p is true at time t_e is not verified in state s at time t_v.

The valuation function takes a snapshot of a state s at a time t_v, and says which assertions count as verified and which do not. This is not the same as saying which assertions count as true and which as false. The relation between whether an assertion is true or false in a state, and whether it is verified or not in a state at a certain time is as follows. When an assertion is verified in a state, then it counts as true in that state. But when an assertion is not verified in a state at a given time it can mean one of three things (a) the assertion is false in that state, (b) the truth of the assertion is undecided by the state, or (c) the assertion is true according to the state, but its truth has not yet been verified.

It might seem perverse to run the different cases (a)–(c) together. However, by reference to different information states and verification times, related to s and t_v by the orderings $<$ and $\sqsubseteq_t$, these three cases can be distinguished. Moreover, they can be distinguished without appeal to a partial valuation function. This turns out to simplify matters when it comes to stating the evaluation conditions for formulas in the language L. It will not be necessary to give separate conditions for positive and negative evaluations, as Data Semantics does with $\dashv$ and $\vdash$.

As a piece of useful terminology, if $V(s, t_v, p, t_e) = 1$ for some time t_v, we say that s supports the assertion that p is true at t_e.

Temporal Orders

In what follows, only the temporal ordering of precedence, $<$ will be of interest to us (except for Section 3.2.3). For the most part we will be concerned with relations between two time indices or between a time index and an event time. It turns out that $<$ can be used to express all that is required.

We can liken the times in I and E to numbers (points) and convex sets of numbers (periods) on any number line you care to choose (e.g. integers, rationals, reals). When applied to time points (individual numbers), $<$ has the properties one would expect from the numerical relation of 'less than' (given that $<$ is linear, not branching). However, matters are slightly more complicated when $<$ relates a time point to a time period, or two time periods.

The easiest way of proceeding is to define time periods in terms of their start and end points. To this end, we can introduce the functions *start_point* and *end_point* that map periods onto their start and end points respectively (*start_point* and *end_point* map time points onto themselves). We can then define precedence between periods or periods and points in terms of precedence between their start and end points:

$t1 < t2$ iff $\mathit{end\_point}(t1) < \mathit{start\_point}(t2)$

$t1 \leq t2$ iff $\mathit{end\_point}(t1) \leq \mathit{start\_point}(t2)$

$t1 > t2$ iff $\mathit{start\_point}(t1) > \mathit{end\_point}(t2)$

$t1 \geq t2$ iff $\mathit{start\_point}(t1) \geq \mathit{end\_point}(t2)$

$t1 = t2$ iff $\mathit{start\_point}(t1) = \mathit{start\_point}(t2)$ and $\mathit{end\_point}(t1) = \mathit{end\_point}(t2)$

The negations of these relations also need to be considered. For example, if $t1$ and $t2$ are points, then $t1 \not< t2$ expresses the same as $t1 \geq t2$. But when they are periods, the two are different (Figure 3.1); $t1 \geq t2$ forbids overlap between the $t1$ and $t2$, whereas $t1 \not< t2$ permits it.

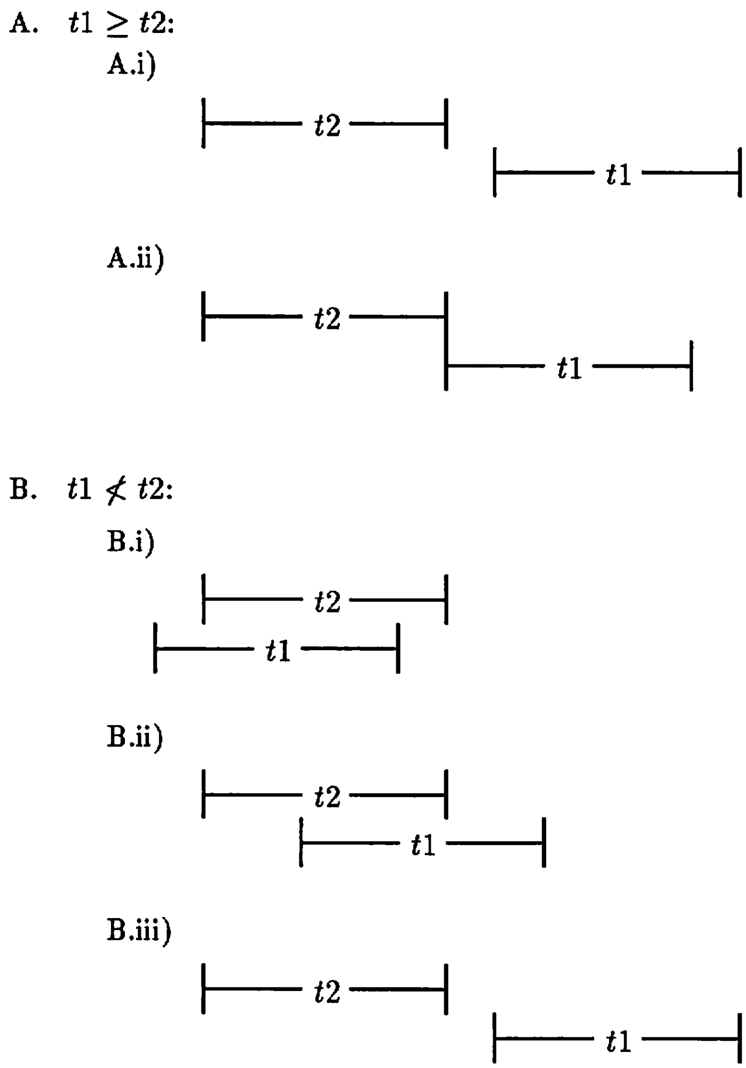

Figure 3.1: Negated Temporal Relations

Beyond saying that $<$ is transitive, irreflexive and linear it is not necessary to impose any further conditions on it. The times in E behave like numbers and convex sets of numbers on any number line you care to choose.

The instants in I selected from E must satisfy an extra property (note that I does not necessarily contain all the time points in E):

$$\forall t.\ \exists t_1.\ t_1 < t \wedge \neg(\exists t_2.\ t_1 < t_2 < t)$$

That is, for every point t in I, there must be another point t_1 in I that immediately precedes it, with no other points t_2 intervening. The discreteness of the times in I reflects a comparable discreteness of information states in S.

I will not have much to say about the relation of temporal inclusion (though see Section 3.2.3). When $t_1 \trianglelefteq t_2$, then the point or period t_1 is included / contained within the period t_2. Van Benthem (1991) discusses different properties that may be satisfied by the inclusion relation.

3.1.2 Constraints on Information Models

The valuation function V, the informational ordering $\sqsubseteq$ and the temporal ordering $<$ are mutually constrained in a number of ways.

Constraint 1: Monotonicity of Verification

Verification of (atomic) assertions is monotonic over time. That is, as time goes by, all the atomic assertions verified in a given state will remain verified, and some extra assertions may also become verified.

Monotonicity of Verification:
If $V(s, t_1, p, t_e) = 1$
then for all time indices t_2 such that $t_1 < t_2$, $V(s, t_2, p, t_e) = 1$
(for all states s, sentences letters p, time indices t_1 and event times t_e).

Constraint 2: Monotonicity of Information Extension

Given a time point t, $s_1 \sqsubseteq_t s_2$ says that at t, s_2 is an informational extension of s_1. The relation $\sqsubseteq_t$ on S is transitive and reflexive for any t.

The following monotonicity property is imposed for atomic sentence letters

Monotonicity of Information Extension:
If $s_1 \sqsubseteq_t s_2$, then
(a) if $V(s_1, t, p, t_p) = 1$, then $V(s_2, t, p, t_p) = 1$, and
(b) if $V(s_1, t', p, t_p) = 1$ for some time point t', then $V(s_2, t'', p, t_p) = 1$ for some time point t''

That is, s_2 extends s_1 at time t if (a) all the atomic assertions verified in s_1 at t are also verified in s_2 at t, and (b) if all the atomic assertions verified at some time in s_1 are also verified in some time in s_2. The state s_2 may in addition support some assertions not supported by s_1.

It should be noted that this constraint is expressed using an 'if' and not an 'if and only if'. This means that one state, s_2 can support all the same atomic assertions as s_1 and perhaps more, and yet still not count as an informational extension of s_1. Some logically possible extensions of s_1 can contingently be ruled out by the $\sqsubseteq$ ordering, which (remember) is taken as primitive.

Other Informational Orderings

Variants of the $\sqsubseteq$ ordering can be defined in obvious ways, e.g.

$s_1 \sqsubset_t s_2$ iff $s_1 \sqsubseteq_t s_2$ and not $s_2 \sqsubseteq_t s_1$

$s_1 \approx_t s_2$ iff $s_1 \sqsubseteq_t s_2$ and $s_2 \sqsubseteq_t s_1$
(informational equivalence at t)

Two states are informationally equivalent at a time t if they verify exactly the same assertions at t, and will go on to verify exactly the same additional assertions. However, there is no necessity for the two states to verify the extra assertions at the same times. Thus, two states can be informationally equivalent without being identical states.

It is also useful to have versions of the information order that abstract away from the temporal argument. That is, the ordering is determined by the assertions verified in states, regardless of when they are actually verified. This can be defined as follows (note the lack of a t subscript)

$s_1 \sqsubseteq s_2$ iff there is any time t such that $s_1 \sqsubseteq_t s_2$

If $s_1 \approx s_2$, then s_1 and s_2 make exactly the same assertions, but may verify them at completely different times. In this case we say that s_1 and s_2 are *informationally equivalent.*

Restriction 3: No Foreknowledge

A third restriction — one that we may wish to relax for certain sentence letters — concerns foreknowledge or more strictly fore-verification:

No foreknowledge:
If $V(s, t_v, p, t_e) = 1$, then $t_v \geq t_e$

In other words, foreknowledge aside, we cannot verify that some fact occurs at a particular time until it actually occurs. This reflects the intuition, mentioned in Chapter 1, that verification is typically a causal process and that it is only the occurrence of an event that can provide the causal basis for verifying that the event has occurred.

Constraint 4: Ideally Verifying Information States

A crucial fourth constraint is the assumption that ideally verifying information states always exist.

> *Existence of Ideally Verifying States:*
> If there is a state s such that $V(s, t_v, p, t_e) = 1$,
> then there exists a state s' such that $s \approx s'$ and $V(s', t'_v, p, t_e) = 1$ for some $t'_v \not> t_e$

That is, for any assertion made by a state, there is always an informationally equivalent state where the assertion is verified as soon as possible. In practice, this means that the assertion is verified as soon as it has become true, though with foreknowledge this may be earlier.

This restriction ensures that for each information state in the model there will be an informationally equivalent state in the model where all assertions are verified as soon as possible. The restriction is essentially a richness condition on models that ensures that there are enough information states.

Constraint 5: Convergence of Verification

When $s_1 \sqsubseteq_t s_2$, s_2 verifies all the same atomic assertions at t as s_1 does, and supports all the same unverified atomic assertions as s_1. It may well be the case that some of these unverified assertions come to be verified sooner after t in s_1 than they do in s_2. Other assertions may come to be verified sooner in s_2 than in s_1.

However, assuming that verification is never indefinitely postponed, eventually the two states will converge. That is, at some time after t all the assertions supported by s_1 will be verified both in s_1 and in s_2. Imagine there is another state s_3 extending s_2 for some other time t'. Eventually s_2 and s_3 will converge, and if s_1 and s_2 have also converged, s_1 and s_3 will also converge.

To capture this behaviour, the following constraint is placed on information models[1]

[1]Frank Veltman has pointed out that this constraint is actually a consequence of the following statement prohibiting indefinite postponement of verification

> If $s_1 \sqsubseteq_t s_2$, then exists t' such that
> for all $t'' \geq t'$, if $V(s_1, t'', p, t_p) = 1$ then $V(s_2, t'', p, t_p) = 1$.

Convergence of Verification:
If $s_1 \sqsubseteq_{t_1} s_2 \sqsubseteq_{t_2} s_3$,
then there is a time t_3 such that $t_3 \geq t_1$, $t_3 \geq t_2$ and $\forall t_4 \geq t_3$, $s_1 \sqsubseteq_{t_4} s_3$

Constraint 6: Absurdity

The final constraint says that the absurd sentence, $\perp$ is never verified in any information state

No Absurdities:
For all s, t_v, t_e: $V(s, t_v, \perp, t_e) = 0$

3.2 Past and Present Tenses

This section describes how formulas in the language L built up from atomic sentence letters, conjunction, disjunction and the past and present tense operators are evaluated relative to information models.

To do this, two support relations are introduced. The first,

$$s, a \models \phi$$

says that a sentence ϕ uttered at a time a is verified in state s. This relation is defined in terms of the main support relation (which is defined recursively below):

$$s, a \models \phi \text{ iff } s, a, a, a \models \phi$$

The three occurrences of a correspond to the values of the assertion, verification and event time indices respectively. When interpreting a sentence, it is assumed that these indices are initially identical the time of utterance. The past and present tense operators can alter the value of the event time index relative to which subformulas are evaluated, while conjunction and disjunction leave it unaltered. It is only in the next section that we will encounter connectives that can alter the assertion and verification time indices.

3.2.1 Conjunction and Disjunction

The evaluation of atomic sentences, conjunctions and disjunctions are as follows

Atomic sentences:
$s, a, v, e \models p$ iff $V(s, v, p, e) = 1$

Def: $\wedge$
$s, a, v, e \models \phi \wedge \psi$ iff $s, a, v, e \models \phi$ and $s, a, v, e \models \psi$

Def: $\vee$
$s, a, v, e \models \phi \vee \psi$ iff $s, a, v, e \models \phi$ or $s, a, v, e \models \psi$

Note that with atomic sentences the value of the assertion time index has no effect on the result of the evaluation.

Disjunction While the clauses above for the support relation, $\models$, are self explanatory, the treatment of disjunction may seem a little odd. If says that $\phi \vee \psi$ is verified either if ϕ is verified or if ψ is verified. The English word *or* does not behave in this way — *A or B* can be true, even if it is not known which of A or B are true. While conjunction may plausibly be likened to the English word *and*, disjunction is not the same as *or*. *Or* in fact has a more conditional air to it.

Monotonicity Properties

We can define obvious analogues of monotonicity of verification and monotonicity of information growth:

Monotonicity of Verification:
If $s, a, v, e \models \phi$ then $s, a, v', e \models \phi$ for all $v' > v$

Monotonicity of Information Growth:
If $s, a, v, e \models \phi$ then $s', a, v, e \models \phi$ for all $s' \sqsupseteq_v s$

It is quite straightforward to show that monotonicity of verification and information growth applies to all atomic sentences. It is also easy to show that they both apply to conjunctions and disjunctions of monotonic formulas.

3.2.2 The Tenses

The support relation for tensed formulas is defined as follows

Def: past
$s, a, v, e \models \textbf{past}(\phi)$ iff there is some $e' < a$ such that $s, a, v, e' \models \phi$

Def: pres
$s, a, v, e \models \textbf{pres}(\phi)$ iff there is some $e' \geq a$ such that $s, a, v, e' \models \phi$

In both cases, the assertion time provides the primary deictic centre of the tense. The past tense selects a new value for the event time preceding the assertion time. The present tense selects a new event time either simultaneous with of following the assertion time.

From these definitions, it would appear that the past and present tenses are simple duals of each other. However this neglects the contribution made by restrictions on foreknowledge.

Suppose that the sentence $\mathbf{pres}(p)$ is uttered at a time a, where p is an atomic sentence about which foreknowledge is not possible. We therefore have

$$s, a, a, a \models \mathbf{pres}(p) \text{ iff there is a time } e \geq a \text{ such that } V(s, a, p, e) = 1$$

Absence of foreknowledge means that $a \geq e$. If $a \geq e$ and $e \geq a$, it follows that $e = a$. In other words, the present tense is constrained to select an event time identical to the present point of utterance. This means that p must correspond to a stative proposition, since it has to be true at a time point, and only stative propositions can do this.

So, although the present tense looks as though it can select futurate event times, under normal circumstances it cannot do this. This explains why the present tense is unacceptable with non-stative sentences (**John builds a house*), unless the sentence can be interpreted as expressing a present habit (see Section 3.5).

As observed in Chapter 1, when foreknowledge is permitted, e.g. when discussing plans, the present tense can be used to refer to future events.

(1) I go to Edinburgh next Wednesday.

It can do this without recourse to a hidden future tense operator, as proposed by Dowty (1979).

It is important to note that a state holding at a time point may also hold for an extended period surrounding that point. Thus, although in normal circumstances a present tense stative proposition means that the state holds at the time point of utterance, there is no reason why the state should not further extend to either side of this point.

Monotonicity Provided that ϕ satisfies monotonicity of verification and information growth, it can easily be shown that $\mathbf{past}(\phi)$ and $\mathbf{pres}(\phi)$ do as well.

Refinements The definitions given for the past and present tenses are simplifications in two respects. Below, we describe what these simplifications amount to.

3.2.3 Localisation Time

A treatment of tense that deals with both time periods and negation must employ an additional localisation time index. The value of this index is a contextually determined period of time, during which all the eventualities described by a sentence take place.

The need for this extra index can be shown by consideration of Partee's (1973) classic example

(2) I didn't turn the stove off.

Although negation will not be dealt with until the next section, it should be apparent that under any reasonable treatment neither of the following formulas accurately reflect the meaning of this sentence:

(2') a. $\neg\mathbf{past}(p)$

b. $\mathbf{past}(\neg p)$

where p represents turning off the stove. The first says that there was no point preceding the time of utterance at which the stove was turned off (i.e. the stove has never been turned off). The second says that there is at least one time in the past that wasn't spent turning the stove off. Neither corresponds to what is meant when someone utters (2) to say that they forgot to do something when leaving the house that morning.

Partee argues that the tense must be interpreted anaphorically: the event time selected by the tense is some contextually salient one, and the tense does not act like an existential quantifier. However, if one's intuitions are that the contextually salient time is not a specific, stove-turning-off length of time, this solution is not sufficient. My own intuitions are that there is some general period of time, say between cooking the breakfast and leaving the house, during which no events of turning the stove off took place.

These intuitions lead to problems if it is the past tense that selects the salient period. For we have to ask what it is for a proposition to be false at a period. There are two possibilities. The first is that the event making the proposition true does not occupy exactly that period. An appropriate event may occur at some time within t, or it may occupy a larger period containing t, but it does not occupy exactly t. The second is that there is no time contained within the period at which an event of the appropriate kind occurs. Only the second possibility does justice to the intuitions about (2). But it is not a coherent possibility, for the following reason:

We are assuming that 'p is not true at t' means that there is no period within t at which p holds. If we wish to preserve bivalence[2] we then have to say that p is true at t if there is some period within t at which p holds. Suppose that p holds at exactly the time t_p.

[2]While not wishing to preserve bivalence for whether something is *known* to be true or not, we do wish to preserve it for whether something *is* true or not.

Then p will be true at all periods containing t_p, no matter how large. As a result, saying that a proposition is true at a time gives no real indication of when the event that makes the proposition true actually occurs. The notion of a proposition being true at a time becomes a hopelessly vague tool for determining the temporal location of events described by sentences.

A better account of (2) is given by Hinrichs (1988). Here, the past tense acts as an existential quantifier over event times. But as well as forcing the event time to precede the deictic centre / speech time, it is also constrained to occur within some contextually determined localisation time[3]. This is not to deny the possibility of the event time sometimes being anaphorically resolved, but in (2) the required context dependence is given by the localisation time. Provided that the negation has wide scope over the tense, Hinrichs's use of localisation time furnishes the desired truth conditions.

Use of a localisation time also answers a criticism levelled by Tichy (1985) and Galton (1987) against interval based tense logics, namely that they confuse the notion of truth *at* a time with truth *of* a time. Something is true at a time if it occurs exactly at that time, while it is true of a time if it occurs during that time.

Thus, an extra, contextually determined localisation time index is required. Given this, we might redefine the support conditions for the tenses as follows (where l is the localisation time):

- $s, a, l \models \phi$ iff $s, a, a, l, l \models \phi$
- $s, a, v, e, l \models \mathbf{past}(\phi)$ if there is some e' such that $e' < a$ and $e' \trianglelefteq l$ and $s, a, v, e', l \models \phi$
- $s, a, v, e, l \models \mathbf{pres}(\phi)$ if there is some e' such that $e' \geq a$ and $e' \trianglelefteq l$ and $s, a, v, e', l \models \phi$

Here, sentences are interpreted relative to an information state, utterance time and contextually given localisation time. The verification time is set to initially be identical to the utterance time, and the event time to be identical to the localisation time. The $\trianglelefteq$ relation in the tense definitions is temporal inclusion.

For convenience, we will suppress reference to this extra localisation index, although it can be used to good effect in defining the temporal connectives *before* and *after* (Chapter 4).

[3]Hinrichs uses the term 'reference time' in place of 'localisation time'. However, his use of reference time bears little resemblance to Reichenbach's (1947) original usage, and the term is used to mean so many different thing by different writers. Following Bertinetto (1985), 'localisation time' is a better and less ambiguous expression.

3.2.4 Temporal Quantification and Tense Operators

There is nothing sacrosanct about the use tense operators rather than (indexically restricted) temporal quantification. Following Dowty (1979, 1982) it may be advantageous to combine the two.

To see how this can be done, let us recast a Priorean treatment of the past tense operator $PAST$ using a combination of temporal quantifiers and operators. In Priorean tense logics (Prior 1967), propositions are evaluated relative to a single temporal index. The notation $[\![\phi]\!]^t$ will be used to indicate the result of evaluating a formula ϕ relative to a time t. The relevant Priorean definitions are (valuation V from atomic formulas and times to truth values):

$$[\![p]\!]^t = V(p,t)$$
$$[\![PAST(\phi)]\!]^t = 1 \text{ iff there is some time } t' < t \text{ such that } [\![\phi]\!]^{t'} = 1$$

Within Dowty's framework, there is a special indexical constant t^* in the language, which always denotes the current value of the temporal index. There are also other terms and variables referring to times, quantification over times, and most importantly a two place tense operator AT which takes a time and a formula as its arguments. Instead of writing $PAST(\phi)$, in Dowty's formalism one writes

$$\exists t'.\, t' < t^* \wedge AT(t', \phi)$$

The two significant semantic definitions are

$$[\![t^*]\!]^t = t$$
$$[\![AT(i, \phi)]\!]^t = [\![\phi]\!]^j, \text{ where } [\![i]\!]^t = j$$

That is the past tense consist of a quantifier that selects a time relative to the current value of the temporal index, and then resets the value of that index by means of the operator AT[4].

The Dowty style treatment of the past tense is not equivalent to the Priorean one; it is potentially more expressive. In particular, it allows two slightly different notions of a

[4]It is necessary to assume that the denotations of terms referring to times — constants and variables — do not vary depending on the value of the temporal index. Dowty's framework has many similarities to the reified temporal logics of Shoham (1986) and Reichgelt (1986). Shoham argues for a clear separation of the temporal part of the language from the non-temporal part. This would prohibit predications of the form $Q(a, b, t)$ where a and b are non-temporal terms and t is a temporal term. To capture the time varying of predications, one would instead have to write $AT(t, Q(a, b))$. A similar separation of the temporal and non-temporal parts of the language is also required in Dowty's framework.

wide scope tense, as compared to Prior's one. Specifically, given some operator Op, a wide scope Dowty tense could be either:

$$\exists t'.\, t' < t^* \wedge Op(AT(t', \phi))$$

or

$$\exists t'.\, t' < t^* \wedge AT(t', Op(\phi))$$

(in comparison to the narrow scope $Op(\exists t'.\, t' < t^* \wedge AT(t', \phi))$. In the first case, the $AT(t', \ldots)$ operator cancels any temporal effects of the operator Op on the formula ϕ, since whatever Op does, the temporal index gets reset to the value of t'. But in the second case, Op can temporally affect the interpretation of ϕ. It is this second case that corresponds to a Priorean wide scope tense — $PAST(Op(\phi))$.

Within the treatment of tenses being developed in this chapter, instead of the operator AT we would require something like

$$verif(v, e, \phi)$$

which says that the assertion that ϕ holds at event time e is verified at v in whatever the current state happens to be. We would also need indexical constants referring to the assertion and verification times, a^* and v^*. Instead of writing $\mathbf{past}(\phi)$, we might write

$$\exists e, v.\, (e < a^* \wedge v = v^*) \wedge verif(e, v, \phi)$$

where $s, a, v', e' \models verif(e, v, \phi)$ iff $s, a, v, e \models \phi$

For convenience we might abbreviate the quantificational part of the tense and instead write

$$\mathbf{past}_{[e,v]}\, verif(e, v, \phi)$$

where $\mathbf{past}_{[e,v]}$ is a short-hand for $\exists e, v.\, (e < a^* \wedge v = v^*) \wedge \ldots$. We can also use $\mathbf{pres}_{[e,v]}$ as the corresponding abbreviation for the quantification associated with the present tense.

The consequence of doing this is that it keeps the values of the event and verification times fixed for the formula ϕ, regardless of what other operators intervene between the temporal quantifier and the *verif* operator.

Under this revised treatment of the tenses we could do away with keeping the event time as a third temporal index. Instead, the event time is treated as an ordinary variable whose occurrence in a *verif*(...) formula must be bound by a tense. To completely do away with having a third temporal index, we would have to make sure that *verif* applying to complex formulas could be broken down into applications of *verif* to simple formulas, e.g.

$s, a, v', e' \models \textit{verif}(e, v, \phi \wedge \psi)$ iff
$s, a, v', e' \models \textit{verif}(e, v, \phi)$ and $s, a, v', e' \models \textit{verif}(e, v, \psi)$

By doing this, we also have a potential explanation of why untensed main clauses, e.g. **John swim*, are unacceptable. Without a tense to introduce a bound event time variable, tying a proposition to a time via *verif*, there is no way of even starting to evaluate the formula. This is unlike the Priorean case, where there is always an index of evaluation available.

Mixing temporal quantifiers and operators considerably increases the complexity of the language. But so long as the the temporal and non-temporal components of the language are kept strictly separate, no semantic surprises will ensue (Shoham 1986). Enforcing this separation means that it is only through the *verif* operator that times are linked to non-temporal formulas. In what follows, I will usually suppress the extra complications to formulas that would arise out of mixing temporal quantifiers and operators, though there is one occasion on which this simplification needs to be lifted (p. 94).

3.3 Conditionals

There are two forms of indicative conditional in English, giving rise to slightly different kinds of *modus ponens* inference. The first supports normal *modus ponens*:

$A \rightarrow C$, A
Therefore: C

e.g. If John is asleep, he is tired.
He is asleep.
Therefore, he is tired.

The second supports a weaker version of *modus ponens*:

$A \Rightarrow C$, A
Therefore: It will be the case that C

e.g. If the strip bends, the temperature rises.
The strip bends/is bending/has just bent.
Therefore, the temperature will rise.

This necessitates the introduction of two slightly different conditional connectives, $\rightarrow$ and $\Rightarrow$. The connective $\rightarrow$ supports immediate *modus ponens* reasoning. The connective $\Rightarrow$ supports futurate *modus ponens*.

I will sometimes refer to $\Rightarrow$ as a habitual conditional, and $\rightarrow$ as a specific conditional. This terminology should be treated with caution, though. While $\Rightarrow$ is the only candidate for expressing certain kinds generalisations about the future (e.g. *if the strip bends, the temperature rises*), not all generalisations are expressed via $\Rightarrow$. This matter is taken up in Section 3.5.

Both connectives say what happens in any information state obtained by minimally extending the current state with an antecedent assertion, namely that a consequent assertion must also be supported in this minimal extension. The two conditionals differ only in when they allow the antecedent and consequent assertions to be verified in such minimal extensions.

Before defining the connectives $\rightarrow$ and $\Rightarrow$, it is necessary to define the notion of minimal extension.

3.3.1 Minimal Information Extensions

The basic idea behind a minimal information extension is as follows. Suppose you have a state s that at a time v does not verify an assertion corresponding to a formula ϕ uttered at time a, i.e.

$$s, a, v, e \not\models \phi$$

Suppose there is a state s_1 such that

$$s_1, a, v_1, e \models \phi, \text{ where } v_1 \geq v \text{ and } s \sqsubseteq_{v_1} s_1$$

Then s_1 is a minimal extension of s with respect to ϕ at a, written as

$$s \sqsubseteq_v^{\phi,a,e} s_1$$

if there is no s' and no v' such that (a) v' precedes v_1, (b) s' precedes or is equivalent to s_1 at time v', and (c) $s', a, v', e \models \phi$. That is, s_1 is the (possibly joint) first state coming after s that verifies the assertion as soon as possible. More formally, we can say

Def: $s \sqsubseteq_v^{\phi,a,e} s_1$
$s \sqsubseteq_v^{\phi,a,e} s_1$ iff
(a) $s_1, a, v, e \models \phi$, and
(b) There is no s' or v' such that
$s \sqsubseteq_{v'} s' \sqsubseteq_{v'} s_1$, and
$a \leq v' < v$, and
$s', a, v', e \models \phi$

Note that if s already verifies the assertion at a, s counts as its own minimal extension, with $v = a$.

Another way of looking at a minimal extension is in terms of adding just the assertion to the original state. This will take you to a possibly new state, where in due course the assertion will be verified. Minimal extensions will be those states that result where the assertion becomes verified soonest. The subscript v in $s \sqsubseteq_v^{\phi,a,e} s_1$ says what this soonest verification time is.

In Appendix B.1 the following theorem about minimal extensions is established:

> *Theorem: All extensions extend minimal extensions:*
> If $s_2 \sqsupseteq_{v_2} s_0$ and $s_2, a, v_2, e \models \phi$,
> then $\exists s_1, v_1$ such that $a \leq v_1 \leq v_2$, $s_2 \sqsupseteq_{v_2} s_1$, and $s_0 \sqsubseteq_{v_1}^{\phi,a,e} s_1$

That is, if one state (s_2) extends another (s_0) and verifies some assertion, then s_2 also extends some minimal extension of s_0 that verifies the assertion. As a direct consequence of this we have that if there are any states like s_2, then such things as minimal extensions exist.

Examples of Minimal Extension

A minimal extension to a state must verify the extending assertion at least as soon as any other informationally equivalent state. This means that ideally verifying information states are always prime candidates for minimal extensions. They are not invariably the only candidates, however. In most cases, there will be a set of informationally equivalent states that are all minimal extensions of the original.

This can best be illustrated by an example. Suppose that an utterance of some sentence ϕ relative to an assertion time 4 corresponds to an assertion that 'p is true at time 2'. Consider an information state s that verifies just one assertion, e.g. 'q is true at time 1'. What are the minimal extensions of s with respect to ϕ uttered at time 4?

Suppose that we have at least the states shown in Table 3.1 all of which are informational extensions of s (and assume that s_1–s_5 are related to each other by information extension as the assertions they verify would suggest). Assuming monotonicity of verification for all the assertions, states s_1–s_4 are all minimal extensions of s at time 4. Of these, s_1 is a completely ideally verifying state, while s_2 ideally verifies the assertion about p but not the one about q. States s_3 and s_4 are not ideally verifying at all. But they both verify the assertion about p at time 4. This satisfies the definition of minimal extension, since there is no time not preceding the assertion time 4 but preceding the verification time 4 at which the assertion is also verified. State s_5 is ruled out because there are other states

s:	Assertion: q holds at t=1	Verified at: t=3
s_1:	Assertion: p holds at t=2	Verified at: t=2
	q holds at t=1	1
s_2:	Assertion: p holds at t=2	Verified at: t=2
	q holds at t=1	2
s_3:	Assertion: p holds at t=2	Verified at: t=3
	q holds at t=1	2
s_4:	Assertion: p holds at t=2	Verified at: t=4
	q holds at t=1	2
s_5:	Assertion: p holds at t=2	Verified at: t=5
	q holds at t=1	2
s_6:	Assertion: p holds at t=2	Verified at: t=5
	r holds at t=1	3
	q holds at t=1	2

Table 3.1: States Extending s

that verify the assertion before time 5 and at or after time 4. State s_6 is not a minimal extension, since it contains another assertion altogether.

Well-Foundedness

For there to be such things as minimal information extensions, it is vital that both the informational ordering $\sqsubseteq_t$ and the temporal ordering over indices $<$ are well-founded. An ordering is well-founded if for any object, there is always another object that immediately succeeds it in the order without a third object intervening. (The numerical order 'greater than' is well-founded on the integers, but not on the real numbers). Without well-foundedness, we cannot claim that there is such a thing as a *minimal* information extension, since for any extension we choose, we could always find a smaller extension.

Lewis (1973) questions the limit assumption, which is a form of well-foundedness applied to a similarity relation between possible worlds. He asks us to consider a line just under an inch long, and contemplate the most similar world in which the line is longer. There is in fact no minimal enlargement one can make to the length of the line, since its length corresponds to a real number. For this reason, there is no possible world *most* similar to the actual one in which the line is longer — for any world in which the line is longer, there is another one in which it is slightly less long, but still longer.

Does the same argument apply to information states and information extensions? I think not. First, the limit assumption concerns similarity between worlds or states, and not extensions. Suppose there is an information state that merely asserts that there exists some line. This state can be extended to give information about the length of the line. One extension might say that the line is 1.1 inches long, another might say that it is 1.09 long, and so on. These all count as different (minimal) extensions of the original state. The state asserting the line is 1.09 inches long may be more similar to a state asserting it is 1.0 inches long than the state asserting the line is 1.1 inches long, but this is not the point at issue.

Second, while Lewis views possible worlds as language independent objects, the same assumption does not have to be made about information states. Information states are being used to give an account of the semantics of natural language utterances. It is plausible to suppose that no natural language utterance will be fine-grained enough for the non-well foundedness of, say, line lengths to make any practical difference.

Stalnaker (1981) makes roughly the same point with regard to the use of possible worlds in evaluating actual, natural language conditionals. In this context, worries about the limit assumption amount to a pointless emphasis on detail. Imagine a situation where there is a line two inches long, and a box one inch across by half an inch deep. Is it true that *if the line were shorter, it would fit in the box*? The answer is clearly 'yes', but it involves the line being shorter by a considerable amount. This indicates that comparatives like *shorter* involves reference to some (contextually salient?) measure of length by which the object is shorter. Once this is taken into account, the inclination to worry about the non-well foundedness of line lengths diminishes considerably.

3.3.2 Conditional Definitions

Two conditional connectives can be defined in terms of minimal information extensions:

Def: $\Rightarrow$
$s, a, v, e \models \phi \Rightarrow \psi$ iff
$\forall s_1, v_1$ such that $v_1 > a$ and $s \sqsubseteq_{v_1}^{\phi,a,e} s_1$,
there exists a $v_2 \geq v_1$ such that $s_1, v_1, v_2, e \models \psi$

Def: $\rightarrow$
$s, a, v, e \models \phi \rightarrow \psi$ iff
$\forall s_1, v_1$ such that $v_1 \geq a$ and $s \sqsubseteq_{v_1}^{\phi,a,e} s_1$,
there exists an s_2 such that $s_2 \approx s_1$ and $s_2, v_1, v_1, e \models \psi$

Explanation

Both conditionals make a minimal extension to the initial s state to add an assertion corresponding to the antecedent. The time at which the assertion is first verified in this state is passed on as the assertion time for the consequent. The same state should then support an assertion of the consequent made at this time.

With $\rightarrow$, the minimal extension must verify the antecedent assertion either at or after the time the assertion is made. With $\Rightarrow$, the antecedent verification must occur only *after* the assertion is made. (Hence $v_1 \geq a$ in the definition of $\rightarrow$, but $v_1 > a$ in the definition of $\Rightarrow$.)

Since minimal extensions always include ideally verifying states, it follows that the verification time passed onto the consequent will be either (a) the time at which the antecedent assertion can first be verified, if this is after the assertion time, or (b) the assertion time itself, if the antecedent can already be verified before this. With $\Rightarrow$ conditionals, the second possibility is ruled out, since the verification time is constrained to succeed the assertion time for the antecedent.

As a result, past tense antecedents are not appropriate with $\Rightarrow$ conditionals: the antecedent describes some event preceding the assertion time a, and so all the minimal extensions of s will have verified the assertion by a. But $\Rightarrow$ demands that the assertion only be verified after a.

A second difference between $\rightarrow$ and $\Rightarrow$ is that $\rightarrow$ demands that the consequent assertion be verifiable as soon as it made. With $\Rightarrow$, the consequent assertion does not have to be verifiable straight away.

This is why $\rightarrow$ quantifies over informationally equivalent states. A given minimal extension verifying the antecedent, may not verify the consequent straight away. However, if there is an equivalent state that does verify the consequent immediately, the consequent still counts as immediately verifiable. An example will illustrate.

Suppose that $\mathbf{past}(p) \rightarrow \mathbf{past}(q)$ holds in a state s at time 3, and let us suppose for ease that s makes no other assertions. Consider one possible minimal extension of s for the antecedent:

s_1: Assertion:	p holds at t=1	Verified at: t=3
	q holds at t=2	4

In this state, the assertion about q is not verified straight away; it is first verified only at time 4, and not at time 3. However, there will be an ideally verifying state informationally equivalent to s_1:

s_2: Assertion:	p holds at t=1	Verified at: t=1
	q holds at t=2	2

Provided both assertions satisfy monotonicity of verification, they will continue to be verified at time 3.

The existence of a state like s_1 would falsify the conditional according to a simplified definition of $\rightarrow$:

$$s, a, v, e \models \phi \rightarrow \psi \text{ iff}$$
$$\forall s_1 \text{ such that } s \sqsubseteq_a^{\phi,a,e} s_1,\ s_1, a, a, e \models \psi$$

There is a minimal extension of s that does not verify the consequent assertion immediately. However, s_1 does not falsify the conditional under the definition that is actually used. Although s_1 does not verify the consequent immediately, there is some informationally equivalent state, s_2, that does.

3.3.3 Monotonicity Properties

Appendix B.2 shows that formulas constructed using $\Rightarrow$ and $\rightarrow$ satisfy monotonicity of verification. More importantly, it also establishes monotonicity of information growth, *provided that the consequent is monotonic.*

The fact that monotonicity of information growth holds means that nothing is lost by moving away from an account of conditionals that explicitly legislates for what happens in *all* information states extending the current one ('all extension conditionals'). The present treatment only legislates for what happens in states minimally extending the current one to make the antecedent hold ('minimal extension conditionals'). But monotonicity ensures that the constraints imposed here will be propagated through to all states extending those minimal extensions. Any other state extending the current one in which the antecedent holds will also extend one of these immediate extensions. So, provided that the consequent behaves monotonically, it will also hold in any of these other states.

The point of departure between the minimal extension treatments and all extension treatments (e.g. Data Semantics, Intuitionistic Logic) is where non-monotonic formulas are involved. Monotonicity breaks down for minimal extension conditionals when the consequent behaves non-monotonically. But for all extension conditionals, monotonicity continues to hold. This is significant when it comes to dealing with conditionals like $\phi \rightarrow \mathbf{may}(\psi)$, where **may** is a non-monotonic modal introduced in the next chapter, since this does not entail $\phi \rightarrow \psi$ (c.f. Chapter 2.5).

Intuitively, minimal extension conditionals are preferable to all extension conditionals. To judge whether an all extension conditional holds, one has to survey the entire sweep of information states extending the current one. There may well be an infinity of such states. With a minimal extension conditional, one need only consider the effects of immediately adding the antecedent to the current information state.

3.3.4 Negation

It is common to define negation as implication into absurdity. The definition for negation given below has essentially this effect:

> **Def:** $\neg$
> $s, a, v, e \models \neg(\phi)$ iff $\forall s_1$ such that $s \sqsubseteq_a^{\phi,a,e} s_1$, $s_1, a, a, e \models \bot$

This does not quite justify the equivalence $\neg(\phi) =_{df} \phi \rightarrow \bot$, since negation demands that the earliest verification time for the (absurd) minimally extending state be identical to the assertion time (without this, negated present tenses could have futurate time reference).

All negations of this form are monotonic. If all minimal extensions of a state s with respect to ϕ support absurdity, and no states do support absurdity, it follows that there are no minimal extensions of s with respect to ϕ. Since any extension of s supporting ϕ must extend a minimal extension with respect to ϕ, it follows that no state extending s supports ϕ. This reasoning applies whether ϕ itself is monotonic or not.

If ϕ is an atomic sentence, we might well ask what it is about a state s that prevents it from having any extension supporting ϕ. One answer might be that s already supports a set of atomic assertions somehow incompatible with ϕ. However, this is not a very satisfying answer if you take the view that all the atomic sentences in the language L are independent of each other, since this precludes incompatibilities of this kind[5]. A better answer is that there are (conditional) restrictions placed on the ordering of information extension. For example, take two atomic sentences corresponding to *John is dead* and *John is eating his dinner*. These sentences are independent of one another — there is no need to decompose the predicates 'being dead' and 'eating dinner' into sets of semantic primitives that are incompatible (e.g. ALIVE and NOT-ALIVE amongst others). Instead, knowledge about the meaning of English imposes a conditional restriction on information extension saying that not state supports John both being dead and eating his dinner.

Double Negation When $\neg$ is applied to non-atomic sentences, a few surprises are in store (though they are not surprising if one is familiar with the behaviour of negation in intuitionistic logic). For example, $\neg\neg\phi$ does not mean the same thing as ϕ. If a state s

[5]By independence, I mean that the valuation function is not constrained in such a way that if one atomic sentence counts as verified in any state, some other sentence must never be verified. The information ordering is required to introduce (contingent) incompatibilities and dependences between atomic sentences.

supports $\neg\neg\phi$ it means that for any state s' extending s, there is a state s'' extending s' that supports ϕ. That is, $\neg\neg\phi$ says that ϕ must eventually come to be supported. It does not necessarily say that ϕ is already supported (though if ϕ is already supported, it follows that $\neg\neg\phi$ is as well).

The failure of equivalence between a doubly negated formula and an unnegated formula is defensible on linguistic grounds. Were the equivalence to hold, we would have to explain the difference between *I went* and *I didn't not go* in terms of a violation of a Gricean maxim to be brief. But without the equivalence, there is in fact a slight difference in meaning (so if anything, it is the maxim to be informative that is violated).

Negated Conditionals Rather less acceptably $\neg(\phi \rightarrow \psi)$ entails $\phi \rightarrow \neg\psi$. On the face of it this seems too strong. To say that it is not the case that *if* ϕ *then* ψ is to admit the possibility of ϕ being true and ψ false. It does not say that whenever ϕ is true ψ is false. What counts as the negation of a conditional has caused a certain amount of controversy (see Veltman 1985). The controversy can be resolved once it is realised that there is a second form of negation.

Non-Assertion

Given information states sensitive to verification times, it is possible to define a second form of negation:

Def: $\sim$
$s, a, v, e \models \sim \phi$ iff $\forall v_1$ such that $v_1 \geq v$, $s, a, v_1, e \not\models \phi$

The negation $\sim$ effectively denies that something is asserted in the current information state. This denial does not extend to saying that it will never be asserted in any other information state, as $\neg$ does. Consequently, $\sim \phi$ is generally non-monotonic.

Negating a conditional is usually a matter of denying an assertion, rather than asserting a denial. It can be established from the definition of $\sim$ that $\sim (\phi \rightarrow \psi)$ does allow for the possibility of ϕ being true and ψ being either true or false.

I do not intend to say much about negation here. In Chapter 5 $\sim$ plays a major role in the logic of modals like **may**.

3.4 Support Conditions for Simple Conditionals

We now turn to the patterns of temporal reference predicted for simple conditionals under the treatment of **past**, **pres**, $\rightarrow$ and $\Rightarrow$. With the exception of two rare kinds of example, covered in Section 3.5 below, the analysis predicts all the different temporal properties listed in Chapter 2, but does not over-predict properties not listed there.

Appendix B.3 presents one example worked through in explicit and gory detail for the benefit of sceptical readers. This section summarises the results and discusses the general reasons why they hold.

3.4.1 Summary of Temporal Properties

As mentioned in Chapter 2, conditional sentences come in two forms: conjunctive and adverbial. Appendix A gives a simple grammar for conditional sentences that makes explicit the difference between conjunctive and adverbial conditional sentences. Conjunctive conditionals conjoin two already tensed clauses. Adverbial conditionals modify the untensed consequent by a conditional construction. The tense of the superordinate consequent clause is a feature of the sentence / clause rather than the verb phrase. As a result, the consequent tense gets wide scope over the adverbial conditional. That is,

$$\begin{array}{ll} \text{Conjunctive:} & \textit{tense}(A) \rightarrow \textit{tense}(C) \\ & \textit{tense}(A) \Rightarrow \textit{tense}(C) \\ \text{Adverbial:} & \textit{tense}(\textit{tense}(A) \rightarrow C) \\ & \textit{tense}(\textit{tense}(A) \Rightarrow C) \end{array}$$

where *tense* can be either **past** or **pres**. By placing the consequent tense outside of the deictically shifting influence of the conditional, adverbial $\Rightarrow$ conditionals can exhibit different temporal properties from conjunctive $\Rightarrow$ conditionals.

The patterns of temporal reference for different configurations of conditionals and tenses are given in Tables 3.2 and 3.3, where n (for 'now') is used to refer to the time of utterance, and A and C to refer both to the antecedent and consequent formulas and the times at which those formulas are asserted to hold.

3.4.2 Discussion of Examples

The major facts about the conditionals tabulated are:

a. The antecedent tense has a primary deictic centre on the time at which the conditional is uttered / asserted.

b. In conjunctive conditionals, the consequent tense is primarily centred on the first time at which the antecedent becomes verifiable (at or after the time of utterance).

$\mathbf{past}(A) \rightarrow \mathbf{past}(C)$	$A < n, C < n$
$\mathbf{past}(A) \rightarrow \mathbf{pres}(C)$	$A < n, C = n$
(with foreknowledge)	$A < n, C \geq n$
$\mathbf{pres}(A) \rightarrow \mathbf{past}(C)$	$A \geq n, C < A$
(with foreknowledge)	$A \geq n, C < n$
$\mathbf{pres}(A) \rightarrow \mathbf{pres}(C)$	$A \geq n, C = A$
(with foreknowledge)	$A \geq n, C \geq n$
$\mathbf{pres}(A) \Rightarrow \mathbf{past}(C)$	$A \geq n, C < A$
$\mathbf{pres}(A) \Rightarrow \mathbf{pres}(C)$	$A \geq n, C \geq A$
(with foreknowledge)	$A \geq n, C \geq n$

Table 3.2: Conjunctive Conditionals

$\mathbf{past}(\mathbf{past}(A) \rightarrow C)$	$A < n, C < n$
$\mathbf{pres}(\mathbf{past}(A) \rightarrow C)$	$A < n, C = n$
(with foreknowledge)	$A < n, C \geq n$
$\mathbf{past}(\mathbf{pres}(A) \rightarrow C)$	$A \geq n, C < n$
$\mathbf{pres}(\mathbf{pres}(A) \rightarrow C)$	$A \geq n, C = n$
(with foreknowledge)	$A \geq n, C \geq n$
$\mathbf{past}(\mathbf{pres}(A) \Rightarrow C)$	$A \geq n, C < n$
(with foreknowledge)	$A \geq n, C \geq n$
$\mathbf{pres}(\mathbf{pres}(A) \Rightarrow C)$	$A \geq n, C = n$
(with foreknowledge)	$A \geq n, C \geq n$

Table 3.3: Adverbial Conditionals

c. In adverbial conditionals, the consequent tense is has a primary deictic centre on the time at which the conditional is uttered / asserted.

Different antecedent and consequent tenses, different conditional connectives used, and different assumptions about the possibility of foreknowledge all affect the patterns of temporal reference that result from these three basic facts.

Verifiability of Past Tense Antecedents

What is the time at which a past tense antecedent first becomes verifiable? Answer: the time at which the conditional is uttered.

Suppose that a conditional with a past tense antecedent, $\mathbf{past}(A)$, is uttered relative to an information state s at a time t. Consider all the states minimally extending s to verify the antecedent assertion.

As the antecedent is a simple past tense clause, it will describe an eventuality occurring prior to t. Call the time at which this eventuality occurs t_A. Given the existence of ideally verifying states, there is bound to be at least one minimally extending state that verifies the assertion from the end point of t_A onwards. It follows that the soonest that $\mathbf{past}(A)$ can be verified at or after t is t itself.

Verifiability of Present Tense Antecedents

What is the time at which a present tense antecedent first becomes verifiable? Answer: the time, at or after the time of utterance, at which the eventuality described has occurred (subject to lack of foreknowledge).

A simple present tense clause uttered at t describes an eventuality occurring at or after the time of utterance. In simple present tense *sentences*, the eventuality is normally constrained to occur at the time of utterance. This is because for sentences, any assertion made must be verified when it is made. However, for conditional antecedent clauses this constraint is lifted. The assertion made by the antecedent clause can be verified after the assertion is made. This allows the antecedent assertion to describe future eventualities.

Without foreknowledge, the assertion cannot be verified until the eventuality it describes takes place. But given the existence of ideally verifying states, once the eventuality has occurred, there will be at least one information state that verifies the occurrence from that time on. (Note that with stative eventualities, the eventuality may continue to hold beyond the point at which it is first verified.) At least one of these ideally verifying states will be a minimal extending state for the antecedent assertion.

With foreknowledge, it may be possible to verify the antecedent assertion before the eventuality it describes takes place.

Distribution of Past Tense Antecedents

Past tense antecedents are inappropriate with $\Rightarrow$ conditionals. This is because $\Rightarrow$ demands that the first time at which the antecedent can be verified must strictly succeed the assertion / utterance time. But past tense antecedents are verifiable at the assertion / utterance time.

Strictly speaking, this means that $\Rightarrow$ conditionals with past tense antecedents are vacuously supported. For there are no minimal antecedent extensions that first verify the antecedent only after the assertion time. And so, all minimal antecedent extensions support the consequent assertion, because there are none.

The $\Rightarrow$ conditional could be strengthened to ensure that there must be at least one minimal antecedent extension satisfying the desired properties. This would ensure that $\Rightarrow$ conditionals with past tense antecedents invariably come out unsupported. Rather than do this, it is better to assume that any sentence that is vacuously true on the basis of form alone should be counted as semantically and pragmatically anomalous[6].

Distribution of Futurate Consequents

There are two ways in which a conditional consequent (in conjunctive conditionals) may come to have futurate reference. The first is where the antecedent passes on a future verification time for use as the deictic centre of the consequent tense. The second is where the consequent is in the present tense, and is not constrained to be verifiable as soon as it is asserted. The first possibility arises in either $\rightarrow$ or $\Rightarrow$ conditionals with present tense antecedents. The second possibility only arises in $\Rightarrow$ conditionals. This is because $\rightarrow$ conditionals demand that the consequent be verifiable as soon as it is asserted, while $\Rightarrow$ permits verification afterwards.

Some Examples

1. Past tense conditionals:

[6] *Bachelors are unmarried men* is trivially true on the basis of what words mean, and not form alone. *All men are men* has more of an air of vacuous truth. While not unacceptable, outside of an elementary logic text it would normally be used in a context where the two occurrences of the word *men* meant something slightly different, e.g. 'male humans' and 'beneficiaries of a patriarchal society'. Conditionals with past tense antecedents formed using $\Rightarrow$ are yet more unacceptable because there is a non-vacuous, syntactically parallel alternative to be constructed using the $\rightarrow$ conditional.

(3) If the bimetallic strip bent, the temperature rose.
$\mathbf{past}(A) \rightarrow \mathbf{past}(C)$

If (3) is uttered at time t the then antecedent event precedes t, and the earliest time that this gets verified at or after t is t itself. Hence, the deictic centre used for the consequent is also t. Consequently, the consequent event also precedes t. That is, the strip bends before t, the temperature rises before t, but nothing is said about the order in which the events occur.

2. Futurate past tense consequents:

(4) If I smile when I get out, the interview went well.
$\mathbf{pres}(A) \rightarrow \mathbf{past}(C)$

The first time at which the antecedent can be verified is when I smile, and this may occur after the time of utterance, t. The consequent describes describes an event occurring before this future time point.

3. Unacceptable futurate antecedents:

(5) ?If John went to the party, he drives home over the limit.
$\mathbf{past}(A) \rightarrow \mathbf{pres}(C)$

The antecedent is verifiable at the time of utterance t. This provides the deictic centre for the present tense consequent. Because the connective used is $\rightarrow$ ($\Rightarrow$ does not permit past tense antecedents), the consequent assertion must be verifiable as soon as it is made. Assuming there is no foreknowledge, this means that the consequent must describe a (stative) eventuality occurring at t. Unfortunately, the consequent describes an event, and events don't occur at time points. As a result, the conditional is semantically unacceptable (though a conditional of the same syntactic form with a stative consequent is acceptable). That (5) is not pragmatically unacceptable is shown by the acceptability of

(6) If John went to the party, he will drive home over the limit.

However, (5) *is* acceptable if construed as expressing proof of the fact that John habitually drives when drunk. Here, the consequent refers to a habit or propensity holding at the time of utterance.

4. For very similar reasons, only one analysis is possible for the following

(7) If the bimetallic strip bends, the temperature rises.
$\mathbf{pres}(A) \Rightarrow \mathbf{pres}(C)$

The antecedent describes an event that might occur in the future. The verification time at which the strip has just bent is passed on as the deictic centre for the consequent. The consequent describes an event occurring either at or after this time. Thus, both events occur in the future, and the antecedent event precedes the consequent event.

Had the conditional been analysed as

(7') $\mathbf{pres}(A) \rightarrow \mathbf{pres}(C)$

we would run into difficulties. The $\rightarrow$ connective demands that the consequent assertion be verifiable when it is made. This would restrict the consequent eventuality to being a stative one that can hold at a time point. But the consequent describes an extended event, and this is not permitted.

5. The possibility of foreknowledge can affect the interpretation of present tense conditionals.

(8) If I do the washing up, Valeria does the vacuuming.
$\mathbf{pres}(A) \rightarrow \mathbf{pres}(C)$

When (8) is used in the context of describing a pre-arranged rota for household chores, it is possible to verify who will do what, according to the rota, long before they actually do it. This allows the present tense antecedent and consequent to have present assertion and verification times, but still describe future events. That is, the antecedent describes an event occurring after now, but verifiable now, as does the consequent. No relative ordering between antecedent and consequent is imposed.

6. Present tense stative consequent:

(9) If the bimetallic strip bends, it is hot.
$\mathbf{pres}(A) \rightarrow \mathbf{pres}(C)$
$\mathbf{pres}(A) \Rightarrow \mathbf{pres}(C)$

The future state of being hot must hold at the time that the strip bends. There is nothing, however, to prevent this state of affairs from extending further on either side of this event.

Since the consequent is stative, it can readily occur in a $\rightarrow$ conditional, but there is nothing to prevent it from occurring in a $\Rightarrow$ conditional either. In this connection, it is worth noting that (9) appears to have two slightly different interpretations open to it. The first describes a general regularity connecting the strip bending and it being hot. The second refers to the possibility of the strip bending at some fairly specific time in the future, and from this being able to conclude that it is then hot. Perhaps these two slightly different readings mark the difference between the use of a habitual $\Rightarrow$ conditional and a specific $\rightarrow$ conditional? This matter is further discussed in Section 3.5 below.

There is a third, adverbial, interpretation for (9) not covered by either of the formulas above. In this interpretation, the consequent refers to a present state of being hot, and not a future state holding at the time the bimetallic strip bends.

Adverbial Conditionals

In an adverbial conditional the consequent tense does not fall under the influence of the conditional connective. This means that it is interpreted as though it were a tense in an ordinary, non-conditional sentence. As a result, adverbial consequent tenses do not take on futurate interpretations. This is because the tense will be interpreted relative to the assertion and verification times set for the sentence as a whole, and these are identical to one another and the time of utterance.

If a sentence like (10) is analysed adverbially

(10) If I smile, the interview went well.
$\mathbf{past}(\mathbf{pres}(A) \rightarrow C)$

the interview precedes the time of utterance. The consequent past tense selects an event time preceding the time of utterance, but it leaves the assertion and verification times unchanged. The conditional plus antecedent alters the values of the assertion and verification times for the untensed consequent formula, but leaves the event time unaffected. Thus, the untensed consequent formula describes an event occurring at the event time selected by the consequent tense.

Likewise with the adverbial interpretation of (9)

(9') $\mathbf{pres}(\mathbf{pres}(A) \rightarrow C)$

The consequent present tense is evaluated with respect to assertion and verification time identical to the time of utterance. In the absence of foreknowledge, this means that the consequent event time must itself be identical to the time of utterance. That is, the consequent refers to a stative eventuality holding at the time of utterance.

Adverbial Verification

In (9') I have glossed over a slight problem in applying constraints arising from the absence of foreknowledge. Absence of foreknowledge means that if t_C is the time at which the consequent eventuality occurs, then the consequent formula C cannot be verified until (the end-point of) t_C.

While the antecedent and conditional does not alter the event time relative to which the consequent formula C is evaluated, it does alter the verification time. The verification time is re-set to be the time at which the (future) antecedent event occurs, t_A. This means that there is nothing to prevent the wide scope consequent tense from selecting an event time the succeeds the time of utterance but precedes t_A. The consequent formula C could still be verified for such an event time without violating constraints on foreknowledge. But it is an empirical fact that in conditionals like (9), when the consequent eventuality precedes the antecedent eventuality, it is because the consequent describes a presently occurring state of affairs. The formula (9') appears to over-predict the temporal properties open to (9).

This over-prediction is the result of an over-simplified representation of the logical form of the conditional. If, as suggested in Section 3.2.4 above, the conditional is represented as (9"), the correct predictions result.

(9") $\mathbf{past}_{[t_c,v_c]}(\mathbf{pres}_{[t_a,v_a]}(\mathit{verif}(t_a, v_a, A)) \rightarrow \mathit{verif}(t_c, v_c, C))$

Although the conditional alters the value of the verification time index current at the point where the consequent formula is to be evaluated, the *verif* operator ensures that the original value, $v_c = n$ is used for verifying the formula.

For the sake of accuracy, I should really make consistent use of the *verif* operator when presenting the logical forms of conditional sentences. But since this is the only kind of example where doing so would make any significant difference, I have simplified matters elsewhere.

3.4.3 Verifiability and Non-Ideal Verification

The following thought may have occurred to some readers. Take

(3) If the bimetallic strip bent, the temperature rose.
$\mathbf{past}(A) \rightarrow \mathbf{past}(C)$

Here, the antecedent is verifiable at the time of utterance. This time is used as the assertion time for the consequent, which is also verifiable at that time.

Suppose (3) is truthfully uttered at time t. However, it is not until some time later that it is discovered that the strip did bend before t. Given this non-ideal verification of the antecedent, does (3) still allow us to draw the conclusion that the temperature rose before t? If not, there is something seriously amiss with the treatment of conditionals proposed here.

Fortunately, the inference does follow. Suppose that (3) is uttered with respect to state s at time t. At some time later, t', we arrive at a state of information, s', that verifies that the strip bent at a time t_A, where $t_A < t$. Corresponding to s' there will be an ideally verifying state s^i. In this state, the fact that the strip bent at t_A will be verified at t_A and (by monotonicity of verification) at all times afterwards. This ideal state s^i will be an extension of one of the states minimally extending s with respect to the antecedent. In this minimally extending state, the fact that the temperature rose before t will also be asserted, given the truth of the conditional. By monotonicity of information growth, this assertion will continue to be made in the state s^i. Furthermore, s^i is informationally equivalent to s'. Therefore, s' also asserts that the temperature rose before t.

Thus, given that s' verifies the antecedent of (3), we can be sure that it will assert the consequent. It may well be that that s' will not verify the consequent assertion until some while after t'. This means that verifying the antecedent only gives indirect, unverified evidence for the consequent. But in due course, if the conditional is true, the consequent will also be verified.

This illustrates a basic point about the way that conditionals have been treated. Their semantics has been given in terms of constraints on ideally verifying information states. Ideally verifying information states are what omnipresent but not omniscient deities would possess — they are always in the right place at the right time to immediately verify what is taking place. The semantics for conditionals can be seen as describing the actual course of information growth for such beings. Lesser mortals will inevitably accumulate information in a more piecemeal, less timely fashion. However, so long as we do not make mistakes about the information gathered, our less ideal course of information development will track that of the more privileged deities. Put another way, what happens in ideally verifying states constrains what happens in informationally equivalent but non-ideal state. By specifying the semantics for conditionals in idealised terms, the constraints on ideal information development get propagated through to all non-ideal courses of development.

Conclusions

The majority of the temporal properties of simple conditionals listed in Chapter 2 have been accounted for, without any over-prediction.

However, there are two kinds of example that have not been covered. The first is exemplified by sentences like

(11) If I didn't get a good night's sleep, I am usually grumpy in the morning.

where there is a futurate past tense antecedent. The second exception is exemplified by

(12) If the Prime Minister gives a good speech, someone else usually writes it for him.

where the consequent eventuality precedes the antecedent one, but it is neither a presently holding state of affairs, nor the kind of eventuality about which we might plausibly enjoy foreknowledge. What is striking about both these examples is that they occur in conditionals expressing a habitual or general connection between events.

3.5 Habituals

3.5.1 Habitual Conditionals

Some conditionals express a general, habitual connection between events of a certain kind, e.g.

(13) If ice is heated, it melts.

Other conditionals express a particular connection between specific events, e.g.

(14) If that ice melted, it was heated.

It is a mistake to look for a single characteristic that distinguishes between habitual and specific readings of conditional sentences. A host of factors contribute to the way that a conditional is interpreted, including (a) the choice of connective, $\rightarrow$ or $\Rightarrow$ (b) whether the conditional is analysed adverbially or conjunctively, (c) whether the antecedent and/or consequent are (anaphorically or otherwise) interpreted as referring to specific events, and (d) the temporal relation of the antecedent and consequent events to each other and to the speech time. These factors can interact or cross-cut in a variety of ways. For example, anaphoric reference to specific events will always block a habitual reading, no matter what other factors work in its favour.

There is one unexpected constraint on when a conditional can be interpreted habitually. This can be brought out by comparing the following three conditional sentences

(15) a. If John goes to a party, he smokes and drinks too much.

b. ?If John goes to the party, he smokes and drinks too much.

c. If John goes to the party, he will smoke and drink too much.

The sentence (15a) is acceptable, and describes what John habitually does at parties. Reference to a definite party in (15b) renders the sentence unacceptable (unless it is interpreted as saying that John's going to the party is evidence for his general heavy smoking and drinking). However, inserting the modal *will* in (15c) makes the non-habitual reading acceptable.

The striking thing about (15a) and (15b) is that they may only be analysed in terms of the $\Rightarrow$ connective, and not $\rightarrow$. This is because they have present tense consequents referring to future events (rather than states), something not permitted by the $\rightarrow$ connective. Perhaps, in light of the unacceptability of (15b), the connective $\Rightarrow$ is confined exclusively to expressing present/futurate generalisations?

There is reasonable evidence in favour of this position, though it is far from conclusive. And if true, this is not to say that the connective $\rightarrow$ cannot also be used to express present and past generalisations. Indeed, $\rightarrow$ is the only available candidate for expressing past generalisations.

Conditionals and Quantification

Why it is that conditionals can give rise to habitual readings in the first place? Semantically, the conditional behaves as a universal quantifier over information states and verification times. The past and present tense operators act as existential quantifiers over event times, where there is a functional relation between the event time selected and its (earliest possible) verification time. If you place an existential tense operator within the scope of a universal conditional connective, the tense can be used to refer to a multiplicity of events.

It is important to note that both $\Rightarrow$ and $\rightarrow$ behave like universal quantifiers in this respect. Therefore, it is only to be expected that under the appropriate circumstances both connectives can be used to formalise habitual conditionals.

Adverbial conditionals will not normally express generalisations. This is because the consequent tense is outside the scope of the universalising conditional. The consequent tense therefore selects a single event time for the consequent prior to evaluating the conditional.

It would be wrong to infer from this that all adverbial conditionals are specific and all conjunctive conditionals are habitual. First, even though a conditional is conjunctive, there is nothing to prevent the tenses within the scope of the conditional from being anaphorically resolved to refer to specific times. Second, as will be seen in Section 3.5.2, a wide scope consequent tense can also be interpreted as expressing habituality in its own right, independent of any effects of the conditional.

Past Generalisations

Since past tense antecedents are not compatible with the $\Rightarrow$ connective, it follows that $\rightarrow$ is the only candidate for expressing past generalisations like

(16) If someone died a pauper, they were buried in an unmarked grave.

Habitual interpretations are not usually open to conditionals that have been analysed adverbially, so past generalisations will typically correspond to conjunctive conditionals:

$$\mathbf{past}(A) \rightarrow \mathbf{past}(C)$$

Past specific conditionals will typically be adverbial:

$$\mathbf{past}(\mathbf{past}(A) \rightarrow C)$$

However, other things besides the bare logical form of the conditional can determine whether it is interpreted habitually or specifically, such as reference to specific events via temporal anaphora, adverbials, and so forth.

Present Generalisations

Some types of present generalisation can only be formalised in terms of $\Rightarrow$, e.g. (15a). But there are other cases in which either $\Rightarrow$ or $\rightarrow$ could be used equally well, such as

(17) a. If the bimetallic strip bends, it is hot.

b. Usually if I am grumpy, I didn't get a good night's sleep.

On the face of it, there are also non-habitual conditionals where either $\Rightarrow$ or $\rightarrow$ could be used, e.g.

(18) a. If I smile when I get out, the interview went well.

b. If the light comes on, that fuse is OK.

It is conceivable that there is something about $\Rightarrow$ conditionals that precludes the possibility of resolving tenses to specific times. This is suggested by the previously noted unacceptability of (15a). If this is so, while $\rightarrow$ could still be used for present generalisations with stative or past tense consequents, $\Rightarrow$ could not be used for specific conditionals.

However, the evidence on this point is not entirely conclusive. For example,

(19) If, at the end of this quarter, you write off the money we may lose as a development expense, the Inland Revenue refunds 50% of the tax.

has a futurate, present tense, non-stative consequent, but refers to a specific refund of tax. It is conceivable in this case that the accountant enjoys foreknowledge with respect to the operations of the Inland Revenue (which would allow $\rightarrow$ to be used instead of $\Rightarrow$). But it is far from obvious that this has to be the case.

The best one can say is that there is a preference against specific temporal reference within the scope of $\Rightarrow$ conditionals, though whether this is a hard restriction or a soft preference is difficult to judge. This preference biases $\Rightarrow$ conditionals to habitual interpretations.

Spurious Ambiguity?

Given two different conditional connectives ($\rightarrow$ and $\Rightarrow$), and two broad forms of conditional sentence (adverbial or conjunctive), there would seem to be considerable systematic ambiguity in the syntactic and semantic analysis of conditional sentences. This ambiguity is unwelcome in those case where either $\Rightarrow$ or $\rightarrow$ could be used in a semantic analysis without change of meaning. One such case would be

(20) If water steams, it is hot.

(a present habitual conditional with a stative consequent).

The (computational) consequences of this kind of ambiguity can be considerably reduced if one adopts something like a monotonic approach to semantic interpretation (Alshawi and Crouch 1992). In the CLARE system, for example (Alshawi et al. 1992), prepositions are treated as having vague senses, rather than multiplying out individual sense entries for each preposition. During reference resolution, a more precise sense can be selected for the preposition if necessary. As a result of this, there is not an explosion in the number of parses for sentences involving preposition phrases.

The same techniques can be applied to conditionals. That is, to begin with *if* is given a vague sense that can be further resolved to either $\rightarrow$ or $\Rightarrow$. Moreover, given the way that monotonic interpretation works, if the choice of $\rightarrow$ or $\Rightarrow$ makes no difference to the meaning of the sentence, as in (20), and $\Rightarrow$ and $\rightarrow$ are the only possibilities for resolving *if*, then there is no need to make the choice. The grammar and semantics for conditionals listed in Appendix A employs this kind of technique.

Thus, although introducing two slightly different but occasionally overlapping meanings for *if* is less appealing than having a single meaning covering all cases, given the right framework the ill effects of this can be reduced.

3.5.2 A Habitual Operator

It is instructive to compare the treatment of habitual conditionals advocated in this chapter with the one proposed by Dudman (1983, 1991). Dudman gives habitual conditionals the following kind of analysis

Present Habitual: $Pr(H(\phi \rightarrow \psi))$

Past Habitual: $Pa(H(\phi \rightarrow \psi))$

where Pr and Pa are past and present tense operators, and H is a habituality operator. Aside from the problems noted in Chapter 2 about mixed tense habituals, Dudman's analysis presupposes the existence of some kind of habituality operator. In contrast, the treatment of habitual conditionals that I have suggested makes no appeal to an extra habituality operator — habitual interpretations arise naturally from the semantics of the conditional itself.

There are times when some form of habituality operator appears to be required in non-conditional sentences, e.g.

(21) a. John smokes.

b. John smoked. (But now he only chews tobacco.)

One way of glossing a habitual sentence like *John smokes* is to say that if certain unspecified or implicitly specified conditions arise (e.g. John is offered a cigarette, he has been drinking, etc.), then an event of John smoking takes place. This suggests that habituals might be definable in terms of habitual conditionals, which themselves make no use of a habitual operator.

More specifically, this suggests that wrapping a habitual operator around a proposition should have a similar effect to placing the proposition in the consequent of a habitual conditional. That is, the formula

$$\mathbf{pres}(H(smoke))$$

bears some relation to

$$\mathbf{pres}(?) \Rightarrow \mathbf{pres}(smoke)$$

where ? represents whatever implicit conditions give rise to instances of the habitual behaviour.

The precise form of this relationship depends on a number of subsidiary assumptions. For example, is there a habituality operator H distinct from the past and present tense operators, so that *John smokes* is indeed to be analysed as $\mathbf{pres}(H(smoke))$? Or are there different habitual and non-habitual tenses, e.g. **pres**, $\mathbf{pres}_{hab}$, **past** and $\mathbf{past}_{hab}$, so that *John smokes* is to be analysed as $\mathbf{pres}_{hab}(smoke)$? If there are separate habituality operators, is it the same operator that combines with both the past and present tenses, or are there different versions of the operator?

The most conservative assumption to make is that there is just one habitual operator that combines with both past and present tenses. Bearing this in mind, the following suggests itself as the semantics for this operator H, given the correspondence noted above

Def: H
$s, a, v, e \models H(\phi)$ iff
$\forall s', v'$ such that $s' \sqsupseteq_{v'}^{\psi,e,e} s$, $\exists v'', e'$ such that $v'' \geq v'$, $v' \leq e' \leq v''$ and $s', v', v'', e' \models \phi$

Here, ψ represents an implicit condition furnished by context. The semantics for $H(\phi)$ is entirely dependent on the value of the event time index e and the state s. It is assumed that the event time is set up by a wider scope past or present tense, so that one may talk of habits or propensities existing now, or habits and propensities existing in the past. The event time is used as the assertion time for an implicit habitual conditional.

The habitual operator defined above bears a number of similarities to the analysis of habituals proposed by Schubert and Pelletier (1989), which involves quantification over cases. They suggest that a habitual sentence like

(22) John usually beats Marvin at ping pong.

is to be interpreted as meaning roughly: in most of the cases when John and Marvin play ping pong, John wins. The implicit use of a conditional in the definition of H achieves this effect, though for 'cases' one should read 'information states'. There is of course another difference in that the conditional involves universal quantification over information states, while Schubert and Pelletier only quantify over most cases. However, following van Benthem (1984), we could posit a range of different habitual conditionals, differing in the quantifier over minimal information extensions. At one extreme, the universal quantifier is used, but generalised quantifiers like 'most' are also possible (this possibility is discussed further in Chapter 7).

Temporal Reference of Habits

A habitual formula like $H(\phi)$ is stative inasmuch as it can be predicated of a time point. This means that in conjunction with the present tense there is a contextually salient value that may be selected for the tense's event time, namely the assertion / speech time. This predicts that a present tense habitual, like *John smokes*, will normally be construed as saying that at the time of utterance John has a propensity to smoke.

Intuitions are unclear as to whether *John smokes* can ever describe a future propensity to smoke. There is nothing in the definition of H to completely rule this possibility out. Recall that with ordinary present tense sentences, constraints on foreknowledge dictate that the event time is identical to the speech time. But the formula $tense(H(\phi))$ holds independently of the initial value of the tense's verification time, and there is no obvious way of imposing foreknowledge constraints on ϕ. However, should it be felt necessary to do this, the definition of H above can be revised to include the condition that the initial event and verification times be such that $e \leq v$.

Duration of Habits

Although John may be a smoker at the moment, he always has the option of giving up; habits or propensities may be of limited duration. At first sight, this fact does not seem to be sanctioned by the definition given for the habitual operator H. If a habit holds at a time e, then all events of the appropriate kind occurring at or after e will lead to an occurrence of the habitual behaviour.

This, however, is to neglect the role played by the *localisation time* (Section 3.2.3) in the semantics of tense. If we include a localisation time index in the semantics of H as follows, habits may be of limited duration:

$s, a, v, e, l \models H(\phi)$ iff
$\forall s', v'$ such that $s' \sqsupseteq_{v'}^{\psi,e,e,l} s$, $\exists v'', e'$ such that $v'' \geq v'$, $v' \leq e' \leq v''$, $e' \trianglelefteq l$ and
$s', v', v'', e', l \models \phi$

Here, all the event referred to are constrained to occur within some (probably contextually given) localisation period. Whether the habit extends beyond this localisation period is a matter that is left entirely open. Exactly the same comments apply to habitual conditionals, where the localisation time limits the temporal application of the conditional.

This also serves to correct what might seem like a problem with past habituals. Without the localisation time, a sentence like *John smoked* would say that from some point in the past onwards, whenever appropriate conditions arose (e.g. John was offered a cigarette), he smoked. This habit would continue to the present and beyond, whether or not John has given up. But with a past tense habitual, it is plausible to suppose that the chosen localisation period refers to some past stretch of time. There is thus no reason to infer that John's bad habits extend to the present and beyond[7].

3.5.3 Habitual Adverbial Conditionals

I am not going to deal with the raft of problems concerning interactions between tense, quantification and habituality. My aim is the more limited one of accounting for conditionals like

(23) Usually, if the Prime Minister gives a good speech, someone else writes it for him.

(24) If I didn't get a good night's sleep, I am usually grumpy in the morning.

where the consequent eventuality precedes the antecedent one.

The proposed analysis is quite simple. These are just adverbial conditionals, but where the consequent tense is interpreted habitually. That is, the conditionals correspond to the formulas

(23') $\mathbf{pres}(H(\mathbf{pres}\phi \to \psi)$

(24') $\mathbf{pres}(H(\mathbf{past}\phi \to \psi)$

Spelling out the effects of the wide scope habitual consequent tense on the embedded conditional for (23),we get

[7]There is nothing to prevent localisation periods covering the past and future from being contextually chosen. This is an implausible choice in respect of a sentence like *John smoked*, but were the choice made, it would follow that John still smokes.

$s, a, v, e \models \mathbf{pres}(H(\mathbf{pres}\phi \rightarrow \psi)$ iff
$\forall s', v'$ such that $s' \sqsupseteq_{v'}^{X,e,e} s$, $\exists v'', e'$ such that $v'' \geq v'$, $v' \leq e' \leq v''$ and
$s', v', v'', e' \models \mathbf{pres}\phi \rightarrow \psi)$

That is, ψ holds at some time $e' \leq v''$. The antecedent ϕ holds at some time $e'' \leq v''$. Thus consequent may precede antecedent. With (24), note that the past tense antecedent is evaluated with respect to a *futurate* assertion time, v', set up by the habitual operator. This allows the antecedent to take on a deictically shifted past-in-the-future reading.

$s, a, v, e \models \mathbf{pres}(H(\mathbf{past}\phi \rightarrow \psi)$ iff
$\forall s', v'$ such that $s' \sqsupseteq_{v'}^{X,e,e} s$, $\exists v'', e'$ such that $v'' \geq v'$, $v' \leq e' \leq v''$ and
$s', v', v'', e' \models \mathbf{past}\phi \rightarrow \psi)$

We also have $e' \geq v'$, and by the past tense of the antecedent, $e'' < v'$, where e'' is the event time of the antecedent. Thus antecedent precedes consequent.

A similar analysis may also be given for conditionals like

(25) If the bimetallic strip bends, the temperature rises beforehand.

However, an alternative explanation for the consequent precedes antecedent reading is possible here. It is arguable that *beforehand* is an elliptical adverbial, so that in full the sentence says *If the bimetallic strip bends, the temperature rises before the strip bends.* Here, the complex event of the temperature rising before the strip bends is in fact simultaneous with that of the strip bending.

If the habitual operator can be freely used in conditionals, other forms of habitual conditional are also possible. However, apart from adverbial ones like the two above, most of these will not be straightforwardly habitual. For example, there are at least two possible readings for (26)

(26) If Fido has a wet nose, he is usually healthy.

On the normal habitual reading of (26), we can take it to mean that whenever Fido has a wet nose it is safe to conclude that he is healthy. But there is another reading where from the fact that Fido currently has a wet nose we can conclude that he is the sort of dog who enjoys generally good health, even if he is not too well right now. On this second reading, we have a specific (or rule applying) conditional with a habitual consequent. It does not specify a habitual connection between antecedent and consequent, but a specific connection between an antecedent and a consequent habit.

3.6 Conclusions

This chapter has introduced a class of information models and given definitions for tense and conditional operators relative to those models. From a temporal standpoint, these definitions successfully predict the patterns of temporal reference exhibited by simple conditional sentences. No other theory of tense and conditionals I am aware of even begins to give the right results.

The information models make use of information states with an internal, temporal structure. A given information state can support both verified and unverified assertions, and as time goes by more and more of the unverified assertions will become verified.

Utterances of simple tensed sentence result in the addition of verified assertions to information states — part of the semantic content of such an utterance is precisely that a verified assertion is being made. This is different from conditional and, as we will see, modal utterances, where unverified assertions may be added to an information states.

By using both the assertion and verification times as temporal indices, asymmetries in deictically shifted past and present tenses are accounted for. The assertion time acts as a primary deictic centre for both past and present tenses. In conditional consequents, the assertion time may be shifted forward, allowing for the possibility of futurate past and present consequent tenses. The verification time acts as a secondary deictic centre, and only affects the interpretation of present tenses. When the verification time is shifted to a time succeeding a (present) assertion time, as happens in conditional antecedents, present tenses may be used to make futurate assertions. Past tenses continue to make past assertions under these circumstances. That is, there are two different kinds of deictic shift: primary shift, which affects both past and present tenses, and secondary shift, which only affects present tenses.

The use of the verification time as a secondary deictic centre also explains why the present tense can be used in simple sentences to described planned or pre-determined future events, without recourse to a hidden future tense operator.

The treatment of habitual conditionals also allows a habituality operator to be defined, which may account for habitual uses of the present and past tenses.

The next chapter turns to present tense modal auxiliaries and modalised indicative conditionals.

Chapter 4

Present Tense Modals

The majority of conditionals in English have a consequent clause whose main verb phrase contains a modal auxiliary, e.g.

(1) If the temperature changes, the bimetallic strip *will/may* bend.

(2) If John is in trouble, we *must* go to help him.

(3) If it was less windy, I *could* climb that tree

Conditionals with modalised antecedents are also possible:

(4) If John *can* do that, then so *can* I.

This chapter extends the treatment of Chapter 3 to cover the temporal properties of conditionals containing present tense modal auxiliaries (*will, may, can, shall, must*, etc.). Conditionals with hypothetical modals (*would, might, could, should*) are dealt with in Chapter 6.

Present tense modal auxiliaries are treated as a combination of a present tense operator plus a modal operator. Thus, the auxiliary *will* is analysed as **pres(will...)**. Section 4.1 introduces three (epistemic) modal operators, **must**, **may** and **will**. Section 4.2 describes how these operators interact with the present tense in simple sentences containing no subordinate clauses. Section 4.3 accounts for the temporal properties of modalised indicative conditionals, and Section 4.4 describes the effects that modals have on subordinate relative clause tenses. Section 4.5 deals with other subordinate tense constructions, and also discusses the connections between the temporal connective *when* and conditionals.

4.1 Epistemic Modal Operators

If the temporal properties of the epistemic modals were of no concern, an obvious way of defining possibility (**may**), necessity (**must**) and expectation (**will**) would be along the following lines. *Possibility*: a formula ϕ is possible with respect to a given state iff the state can be minimally extended to support ϕ. *Necessity*: a formula ϕ is necessary with respect to a given state s iff there is no way of minimally extending the state to support $\neg\phi$ (which means that any course of extension to s will eventually lead to a state supporting ϕ). Necessity thus counts as a dual of possibility: not possible not. *Expectation*: a formula ϕ is expected / predicted by a given state s iff any minimal extension of s relative to ϕ leads to a state that is informationally equivalent to s.

Unfortunately, once the temporal reference of modals is taken seriously this simple picture breaks down, especially for necessity. To say something like *John may go* is to say that it is possible that John will go at some time in the future. In terms of minimal information extensions, this means that there is at least one future time, t, such that the current information state can be minimally extended to support the assertion that John goes at t. To say that *John must go* is to say that it is necessary that John will go at some time in the future. In terms of minimal extensions, taking the dual of possibility amounts to the following. There is no time future t such that the current state can be minimally extended to support the assertion that John does not go at t. In other words, John must go at *all* future times, and not just one. This is clearly not what is meant by *John must go*.

An alternative way of forging a duality between temporalised possibility and necessity would be to ensure that the future time referred to is taken outside the scope of the dualising negations. That is, there is some time in the future, t, such that there is no minimal extension of the current state supporting the assertion that John does not go at t. This is clearly better than before, but still does not capture the intuitive meaning of *John must go.* This (usually) says only that John must go sooner or later, and not that there is some specific time at which John must go.

Apart from these temporal problems, it also transpires that this kind of dualised treatment brings in its wake the problematic inference from $\mathbf{must}(\mathbf{may}(\phi))$ to $\mathbf{must}(\phi)$ (Chapter 2). For, if there is no minimal extension of a state s by $\neg\mathbf{may}(\phi)$, then every extension of s must be extendable to support $\mathbf{may}(\phi)$. That is, whatever happens, ϕ must remain a possibility. And if ϕ always has to remain a possibility, no matter what else happens, then ϕ is in fact a necessity.

The definitions given below for the modal operators **must**, **may** and **will** avoid these problems, and predict other patterns of temporal reference exhibited by modal sentences, as will be seen in due course.

Def: may
$s, a, v, e \models \mathbf{may}(\phi)$ iff
$\exists s', a', v', e'$ such that $e' \geq e, v' \geq e$ and $s' \sqsupseteq_{v'}^{\phi, a', e'} s$.

Def: must
$s, a, v, e \models \mathbf{must}(\phi)$ iff
$\forall s'$ if $s' \sqsupseteq_v s$, then $\exists s_0, a', v', e'$ such that $e' \geq e, v' \geq e$ and $s_0 \sqsupseteq_{v'}^{\phi, a', e'} s$ and $\exists s''$ such that $s'' \sqsupseteq_v s_0$ and $s'' \sqsupseteq_v s'$.

Def: will
$s, a, v, e \models \mathbf{will}(\phi)$ iff
$\exists s', a', v', e'$ such that $e' \geq e, v' \geq e$, $s' \sqsupseteq_{v'}^{\phi, a', e'} s$ and $s' \approx s$.

may

The modal formula $\mathbf{may}(\phi)$ corresponds closely to the treatment first suggested. That is, there is at least one time e' that is present or future with respect to the original event time e, such that the state s can be minimally extended to support the verified assertion that ϕ holds at e'. The original event time e will have been previously set up by the present tense associated with the modal.

The way that the assertion and verification times a' and v' are set up is significant when it comes to the effect of the modal on subordinate tenses (Section 4.4).

An important point to note about $\mathbf{may}(\phi)$ is that it is not in general monotonic with respect to information growth. The fact that a state s can be minimally extended to support ϕ does not mean that all states extending s can be similarly extended. This is as one would expect: suppose one claims that John may be at work tomorrow, but subsequently finds out he has booked a day's leave. The possibility that John will be at work disappears on discovering this extra information.

will

The semantics of **will** also corresponds fairly closely to what was at first suggested. Suppose a state s supports $\mathbf{will}(\phi)$ relative to an initial event time e (set up by the associated present tense). Then there will be some time $e' \geq e$ such that minimally extending s to verify that ϕ holds at e' leads to a state that is informationally equivalent to s (i.e. the extension s' supports exactly the same assertions as s but may verify them at different times). This means that s itself will verify that ϕ holds at e' at some point, though it does not have to be the case that s verifies this as soon as it could possibly be verified.

must

The semantics for **must** differs quite widely from that first suggested. A state s supports the formula $\mathbf{must}(\phi)$ relative to an event time e and a verification time v as follows. Let $\mathcal{S}$ be the set of states obtained by minimally extending s with the assertion that ϕ holds at some time $e' \geq e$. Then, any state s' extending s must inevitably grow into a state, s'', that extends one of the minimal extensions in $\mathcal{S}$.

This definition allows for the fact that the assertion that ϕ holds at e' may not be monotonic. It does not demand that the assertion holds in all states s' extending s, and even allows the assertion to be false in some such states. However, when the assertion is false in a state s', it must extend an extension of s in which the assertion is true.

Modality and Negation

The three modal operators defined above appear to form a pretty mixed collection. However, Chapter 5 shows that there is an underlying connection between

$\mathbf{must}(\phi)$ and $\phi \vee \neg\neg\phi$,

$\mathbf{may}(\phi)$ and $\sim \neg\phi$,

$\mathbf{will}(\phi)$ and $\sim\sim \phi$

4.2 Simple Present Tense Modal Sentences

How do these definitions work out in the case of simple present tense modal sentences, e.g.

(5) John will/must/may sleep.

a. $\mathbf{pres}(\mathbf{will}(\phi))$

b. $\mathbf{pres}(\mathbf{may}(\phi))$

c. $\mathbf{pres}(\mathbf{must}(\phi))$

(where $\phi = sleep(j)$).

The present tenses in (5a–c) will all select event times either simultaneous with the time of utterance or following it. The modals will in turn select an event time simultaneous with or following that. Thus the eventualities described by (5a–c) all occur either at or after the time of utterance. This is as required.

This is not quite the end of the story, however. The temporal reference of the modal predication in (5a–c) is to the time of utterance. That is, the present tenses should select event times simultaneous with the time of utterance, and not succeeding it. That is, (5b) says it is *now* possibly the case that ϕ, and not that at some future time it will possibly be the case that ϕ.

This behaviour is predicted by the following: it can be shown that an (epistemic) modal evaluated relative to a time $t' > t$ entails what is expressed by the modal evaluated relative time t. So whatever event time the present tense selects, the same state of affairs could be described by selecting the speech time. And since the speech time is always contextually salient, it makes sense to assume that by preference the present tense will select the speech time as its event time.

To show that this entailment relation holds, consider **may**(ϕ) evaluated relative to some event time e, where e is later than the time of utterance n. By the semantics of **may**, there is at least one $e' \geq e$ such that the current state can be extended to support the assertion that ϕ holds at e'. But if $e' \geq e$ and $e > n$ it follows that $e' \geq n$. So if **may**(ϕ) is supported at e, it is also supported at n. Similarly for **will** and **must**.

Put more informally, if you can show that something is going to be an epistemic possibility in the future, it is thereby an epistemic possibility right now. Likewise, if you can show that later on something will happen in the future, then you can show now that it will happen in the future. And if you can show that in the future something is an epistemic necessity, then you can show now that it is an epistemic necessity.

At first sight, this argument may seem implausible in connection with necessity. Epistemic possibilities may only diminish with time and the acquisition of further information — no new possibilities come into being. Therefore, if something is a possibility in the future, it has to be a possibility now. By contrast, epistemic necessities increase over time — the more information one acquires, the more possibilities are eliminated and the more thing become necessary. In other words, it is quite possible for a certain fact not to be necessary right now, but in due course to become a necessity.

However, this is not the point at issue here. Necessities increase and possibilities diminish only as one moves from one information state to the next. Within a single state, they remain fixed. The argument given previously only concerns present and future modalities within a single information state. That is, we have shown that if ϕ is a necessity / possibility at some time t in a state s, it is also a necessity / possibility at all times preceding t in the same state s. We have not shown that if ϕ is a necessity in a later state s', it is also a necessity in an earlier state s; this is, in general, false.

4.3 Modalised Conditionals

Let us now turn to the behaviour of present tense modal auxiliaries in conditional sentences.

Lack of Antecedent Deictic Shift

It has been observed that present tense modals can lead to deictic shift in subordinate past (and present) tenses. In Section 4.4 it is shown how modal operators can give rise to deictic shift in past tenses within their scope.

One of the striking facts about modalised conditionals with past tense antecedents, such as

(6) If John was at the meeting, he will know what happened.

is that a deictically shifted, past-in-the-future reading for the antecedent is not possible.

This indicates that the consequent modal operator cannot have scope over the antecedent tense. On the face of it, this would seem to rule out the possibility of adverbial modalised conditionals.

However, this presupposes that in adverbial modalised conditionals both the consequent tense *and* the consequent modal are given wide scope, i.e.

$$\mathbf{pres(will(pres(}\phi\mathbf{)} \rightarrow \psi\mathbf{))}$$

But there is no reason to assume this is so. Present tense modal auxiliaries are analysed as a present tense operator plus a modal operator. In adverbial conditionals, it is only the tense, and not the modal operator, that is given wide scope. This means that

(7) If the letter arrives tomorrow, it must/will already be in the post.

can be given an adverbial analysis of

(7') $\mathbf{pres(pres(}\phi\mathbf{)} \rightarrow \mathbf{will(}\psi\mathbf{))}$

The wide scope present tense is not, on its own, sufficient to give rise to antecedent deictic shift. And (7') states that, provided the letter does arrive tomorrow, it *now* must or will be the case that it is in the post.

Kratzer (1979), Stump (1985) and others uniformly analyse all modalised conditionals as having the modal with wide scope:

$$modal(\phi \rightarrow \psi)$$

This would erroneously predict the possibility of antecedent deictic shift.

Admittedly, neither Kratzer nor Stump pay attention to the temporal consequences of their analysis. If pushed, they could perhaps respond as follows. They assume that all modals are inherently conditional, whether they occur in a conditional sentence or not. Modals quantify over possible worlds, and some restriction on the range of possible worlds quantified over always needs to be imposed. In non-conditional modal sentences, the range of quantification is supplied by context. In conditional sentences, the antecedent further constraints the contextually supplied domain of quantification. This being so, they could argue that it can only be the consequent of the conditional that is subject to deictic shift from the modal, since the antecedent is in some sense part of the modal itself.

There are a number of reasons to be unhappy about this line of argument. First, it becomes rather unclear what simple, unmodalised conditionals do, since there is no modal for the antecedent to restrict. It is not at all obvious how or whether simple conditionals relate to modalised conditionals.

Second, there are cases where it is implausible to assign the modal wide scope. One such example was noted in Chapter 2

(8) If I pass the test next week, I can drive any car I like

Assigning the modal wide scope would tend to suggest that I have a present ability or licence to drive, but dependent on passing a future test. But clearly if I drove before the test, even if I subsequently passed the test, I would be in the wrong.

Separating the modality from the tense in auxiliaries like *will, must* or *may* provides a much more satisfactory way of accounting for temporal reference in modalised conditionals.

Futurate Consequents

One of the differences between simple and modalised conditionals is the acceptability of (9), which refers to a definite future consequent eventuality compared to the unacceptability of (10).

(9) If John goes to the party tonight, he will meet Mary.

(10) ?If John goes to the party tonight, he meets Mary.

In Chapter 3 it was suggested that (10) is unacceptable because of a specific time reference combined with a 'habitual' $\Rightarrow$ conditional (only the $\Rightarrow$ conditional permits futurate event denoting consequents).

The consequent *he will meet Mary* in (9) is stative inasmuch as it describes a modal predication holding at a time point. We thus predict that (9) can be analysed in terms of a specific, $\rightarrow$ conditional, compatible with definite time reference in the consequent.

Modal Antecedents

It is the stativeness of modal formulas that explains the general lack of epistemic modals antecedents to conditionals. As Nieuwint (1986) suggests (Chapter 2), conditionals with an epistemic *will* in the antecedent are odd because the consequent then depends on the present predictability of the antecedent event, and not on the occurrence / verification of the antecedent event itself. With non-epistemic modals, it is more plausible to suppose that a consequent may depend on the ability, permission or obligation to do something, as opposed to the performance of the act itself.

4.4 Modally Subordinate Relative Clause Tenses

In Chapter 2 it was noted that past and present relative clause tenses behave differently when the superordinate clause is modal. There is a definite temporal ordering between main and subordinate eventualities in (11) but not in (12)

(11) One day I may marry someone who got rich by honest means.
$\mathbf{pres}(\mathbf{may}(\phi \wedge \mathbf{past}(\psi)))$

(12) One day I may marry someone who gets rich by honest means.
$\mathbf{pres}(\mathbf{may}(\phi \wedge \mathbf{pres}(\psi)))$

In (11) getting rich has to precede the marriage, but in (12) it may either precede or succeed the marriage.

First consider (11), or more specifically the sub-formula $\mathbf{may}(\phi \wedge \mathbf{past}(\psi))$ evaluated relative to an event time identical to the time of utterance. That is,

$$s, n, n, n \models \mathbf{may}(\phi \wedge \mathbf{past}(\psi)) \text{ iff } \exists s', a', v', e' \text{ such that } e' \geq n,\ v' \geq n \text{ and } s' \sqsupseteq_{v'}^{\phi \wedge \mathrm{past}(\psi), a', e'} s$$

This means that in s', ϕ is asserted to hold at e'. Since the value of a' is not explicitly specified, it is permissible for the subordinate past tense to be evaluated relative to a future assertion time. This allows a deictically shifted, past-in-the-future reading. However, it needs to be shown that the time at which the past tense ψ is asserted to hold, e'', precedes e'.

Suppose then, to show a contradiction, that $e'' > e'$. Because of the existence of ideally verifying states, we can be sure that the first time v' at which s' verifies that ϕ holds at e'' will be the end point of e''. (If v' is any later, there will be another state and time verifying $\phi \wedge \mathbf{past}(\psi)$ sooner, and so s' will not be a minimal extension). Given the way that ideally verifying states are defined, the verification time v' may not precede the assertion time

a'. But by the definition of the past tense, e'' must wholly precede the assertion time a'. That is, the end point of e'' must not precede a', but e'' must wholly precede a'. This is a contradiction, so it cannot be the case that $e'' > e'$. An exactly similar argument goes through if the end points of e' and e'' coincide. Therefore, (the end point of) e'' must precede (the end point of) e'.

Now consider (12). Here, there are no constraints on the relative orders of the event times for ϕ and ψ (e' and e'' respectively).

$$s, n, n, n \models \mathbf{may}(\phi \wedge \mathbf{pres}(\psi)) \text{ iff } \exists s', a', v', e' \text{ such that } e' \geq n,\ v' \geq n \text{ and } s' \sqsupseteq_{v'}^{\phi \wedge \mathbf{pres}(\psi), a', e'} s$$

Given that s' is a minimal extension of s, the only constraint on the ordering between v' and a' is that $v' \geq a'$ Suppose that a' strictly precedes v', $a' < v'$. Then ψ holds at a time $e'' \geq a'$, where $e'' \leq v'$. It will also be the case that $e' \leq v'$. (As s' is a minimal extension, v' will in fact be identical to whichever is the later of e' or e''.) No ordering between e' and e'' is imposed by this.

Exactly parallel considerations apply to past and present subordinate tenses within the scope of **must** and **will**. Note that in the case of **will** it would not be sufficient merely to say that s itself will sooner or later verify the embedded proposition. This would not place appropriate constraints on the relation between the assertion time a' and the verification time v'.

Subject Relative Clauses

Richards (1987) observes that relative clause tenses in the subject of a modalised sentence tend not to undergo deictic shift:

(13) (*) Someone who ate an apple now will regret it.

(Richards sees this as a hard constraint on temporal reference, though personally I am not convinced that (13) is entirely unacceptable.) Richards' explanation for the putative unacceptability of (13) relies on assuming (following Montague — see Dowty, Wall and Peters 1981) that the main clause tense applies only to the verb phrase, and that the subject noun phrase corresponds to a quantifier with wide scope over the tense. This means that the past tense in the subject relative clause is not within the scope of the deictically shifting modal.

I have been assuming throughout that tenses apply at the sentential / clausal level rather than at the verb phrase level. This, however, does not necessarily mean that the deictically shifting modal receives wide scope over the subject noun phrase and relative clause. By splitting the modal's tense from the modal operator, the subject noun phrase can fall within the scope of the modal's tense without falling within the scope of the modal itself. Such a configuration would ensure that the relative clause past tense does not receive a deictically shifted interpretation.

However, it is far from clear that subject noun phrases can never fall within the scope of modals. A case in point is

(14) Someone must go and apologise to Mary.

under the interpretation where there is no particular person who must apologise, but where Mary should be apologised to nevertheless. Under this kind of interpretation, deictically shifted subject relative clauses do not strike me as unacceptable

(15) Someone who was at tomorrow's meeting must come and tell me what happened when it is over.

A present tense relative clause is much less acceptable

(16) ?Someone who is at tomorrow's meeting must come and tell me what happened when it is over.

Richards's is therefore probably hasty in adopting the hypothesis that 'tenses which occur in relative clauses of subject noun phrases are not semantically subordinate to other tenses.'

4.5 Other Subordinate Tense Constructions

4.5.1 Futurate Present Deictic Shift

It is not only modals that can lead to deictic shift in subordinate tenses. The same phenomenon occurs if the superordinate clause is in the futurate present tense, as might occur when discussing plans:

(17) John and Mary both place a book on the table.

a. Then Harry picks up the book that Mary puts down.

b. Then Harry picks up the book that Mary put down.

As with modal deictic shift, a subordinate past tense refers to a time preceding that of the superordinate tense, while a subordinate present tense may refer to a time either preceding or succeeding it.

In order to account for this kind of deictic shift it is necessary to revise slightly the proposed semantics for the present tense. Up until now it has been assumed that the present tense, like the past tense, does not affect the values of the assertion and verification time indices. But in fact the present tense sets up a new assertion time lying somewhere between the original assertion time and the event time selected by the tense. That is

- $s, a, v, e \models \mathbf{pres}(\phi)$ iff $\exists e', a'$ such that $e' \geq a$, $a \leq a' \leq e'$, and $s, a', v, e' \models \phi$

Note that when the event time e' is constrained by verification to be identical to the original assertion time a, then the new assertion time a' will also be identical to a.

A past tense within the scope of a futurate present tense will be evaluated relative to an assertion time a' preceding the superordinate event time e'. The subordinate past tense will therefore refer to an event time preceding e'. A present tense within the scope of a futurate present will also be evaluated relative to and assertion time preceding the superordinate event time e'. However, assuming that the subordinate clause is unconstrained by verification in the same way that the superordinate clause is, this means that the subordinate event time may either precede, be simultaneous with, or succeed the superordinate event time.

The same considerations apply to relative clause tenses in present tense conditional antecedents, e.g.

(18) a. If I (ever) meet someone who inherits a lot of money, I will marry them.
b. If I (ever) meet someone who inherited a lot of money, I will marry them.

Here it is not the possibility of foreknowledge that permits a shifted subordinate tense. Instead it is the way that the identity between assertion and verification time is lifted in conditional antecedents. The temporal reference of the consequent will depend on whichever of the antecedent events is verifiable last.

The habitual present can also give rise to deictic shift in subordinate tenses

(19) (After Christmas) John usually throws away the presents he didn't like

Here, however, the deictic shift is due to the fact that the habitual operator defined in Chapter 3 is deictically shifting in the same way that conditionals are. The subordinate past tense is within the scope of such an operator. The shifted reading of the tense thus has nothing to do with the proposed revision to the semantics of **pres**.

4.5.2 Temporal Connectives

Present tenses in temporal conjunctions often have a futurate interpretation, e.g. *leaves* in

(20) I will go when/before/after John leaves.

It is tempting to suppose that the present tense in the temporal clause is futurate for the same reason that it is in conditional antecedents. In other words, temporal connectives like *before*, *after* and *when* can be analysed along similar lines to the conditional.

This temptation is best resisted. It is more plausible to suppose that the futurate present tense is brought about by a wide scope, deictically shifting modal. That is, the appropriate logical form for a temporally conjoined sentence like (20) is

(20') $\mathbf{pres}(\mathbf{will}(\mathbf{after}(\mathbf{pres}(leave), go)))$

That is, the temporal clause *after John leaves* forms an adverbial modifier that attaches before the modal applies. This structure differs from that supposed for adverbial conditionals, where it is only the tense and not the modal of the consequent clause that receives wide scope. More simply put, temporal clauses act as verb phrase modifiers, while conditional clauses act as sentential modifiers.

There are two sources of evidence in favour of analysing conditionals and temporally connected sentences differently. The first is that temporal clauses only allow a futurate present tense if (a) the main clause contains a modal, or (b) the sentence is interpreted habitually, or (c) the sentence discusses plans or pre-determined actions. Conditionals differ in this respect, so that

(21) If the letter arrives tomorrow, it is already in the post.

permits a futurate antecedent without modality, habituality or predetermination. Second, there are a number of anaphoric differences between conditionals and temporally connected sentences, e.g.

(22) a. When/before/after each man$_i$ left, Mary kissed him$_i$.

b. *If each man$_i$ left, Mary kissed him$_i$.

Such anaphoric differences are suggestive of structural differences.

Before and After

It is possible to give a satisfactory analysis of connectives like *before* and *after* without appeal to minimal information extensions, though with reference to the localisation time (Chapter 3.2.3).

- $s, a, v, e, l \models \mathbf{after}(\phi, \psi)$ iff
  $s, a, v, e, l \models \psi$ and $s, a, v, e, l' \models \phi$

  where l' is the initial period of the localisation time terminated by by the start point of e.

- $s, a, v, e, l \models$ **before**(ϕ, ψ) iff
  $s, a, v, e, l \models \psi$ and it is not the case that $s, a, v, e, l' \models \phi$

  where l' is the initial period of the localisation time terminated by by the end point of e.

Recall that the event time selected by tenses must fall within the localisation time period. The connectives **after** and **before** chop off the localisation time at the event time of the superordinate clause (ψ). In a sentence like *John left after Mary arrived* — **past**(**after**(**past**($arrive$), $leave$)), John leaves at a past time e. The localisation time for the subordinate clause, *Mary arrived* terminates at e, which forces its event time to precede e. In a sentence like *John left before Mary arrived*, John leaves at a time e, and no event of Mary arriving must occur in the period from the beginning of the localisation time up until e.

The effects of aspect on the interpretation of *before* and *after* clauses have been noted by Ritchie (1979). For example, *before* prohibits any kind of overlap between the main and subordinate eventualities. *After* by contrast prohibits overlap if the two eventualities are non-stative, but permits it if one or more of them are stative. This prompted Ritchie to claim that *after* was ambiguous, depending on whether it allows overlap or not. Another symptom of the difference between *before* and *after* is exemplified by the non-equivalence between (23) and (24)

(23) John left before Bill was asleep.

(24) Bill was asleep after John left.

In (23) Bill cannot be asleep at any time before John leaves. In (24) although Bill must be asleep after John leaves, there is nothing to prevent him from falling asleep before John goes.

These differences are neatly accounted for by the definitions given for **before** and **after**. With (24), there is some time e at which Bill is asleep, and John's leaving precedes this time. However, Bill may be asleep for longer than just the time e, and perhaps even before John left. When two events are connected, e.g.

(25) John left after Bill left.

it is not possible for the main event to extend beyond the time e at which the event occurs, and so overlap between the main and subordinate events is ruled out[1].

In (23) John leaves at some time e. There must be no time preceding e at which Bill was asleep. Consequently, overlap is ruled out. However, when both main and subordinate eventualities are stative, e.g.

[1]Stump (1985) gives essentially the same explanation for when overlap is and is not permitted by *after*, though he does not deal adequately with the non-parallel situation in the case of *before*.

(26) John was asleep before Bill was asleep.

overlap is possible, but John must fall asleep first. This is because there is a time e at which John is asleep, and where there is no time preceding it at which Bill is asleep. However, John's being asleep can extend beyond e into a period at which Bill is asleep.

The treatment of *before* also captures its counterfactual uses, e.g.

(27) The car stopped before it hit the tree.

All the sentence says is that no event of the car hitting the tree occurred before the car stopped. Whether or not there was an event of the car hitting the tree afterwards is left open. Poznanski (1990) argues that the connective **before** should have its semantics supplemented with a default pragmatic directive saying that if it is possible to do so, assume the the subordinate event does occur (i.e. by default, *before* is not counterfactual). In the case of a car stopping before it hits the tree, other contextual knowledge serves to overrule this pragmatic default[2].

When

When is a rather more problematic temporal connective. The state of affairs typically denoted by a sentence like A *when* B can be diagrammatically represented as follows

```
(a)   - - -[---B---]- - -
        - - -[---A---]- - -

(b)  - - -[---B---]- - -
                - - -[---A---]- - -
```

If A and B are events, the start of A must come sometime after the start of B, and may well occur after the end of B, as in (b) above. When either of A or B are stative, the state of affairs may extend beyond the times enclosed by square brackets. That A may not normally start before B when both denote events is illustrated by the difference between (28a) and (28b)

[2]Hornstein (1991) also suggests a unified account of factual and counterfactual before. However, his account has nothing whatsoever to say about aspectual effects on the connective, and nor does it give a remotely plausible treatment of counterfactual *before*. Essentially, he claims that counterfactual *before* arises when the event times for the main and subordinate clause can be construed, within some formal, largely uninterpreted tense representation, as 'coinciding'. Hornstein claims that his theory predicts that *John had left the party before he hit someone* cannot have a counterfactual reading, since the main event can only 'precede' the subordinate event (where 'precede' refers to some kind of relation in his uninterpreted tense structures). This prediction clearly falsifies Hornstein's account — the sentence can have a counterfactual reading.

(28) a. When John came in, Bill cooked the dinner.

b. When Bill cooked the dinner, John came in.

The two sentences are not synonymous. In (28a) Bill starts to cook once John comes in, and in (28b) John comes in either when Bill starts to cook, or more probably, when Bill has finished cooking the dinner.

By analogy to the treatment of *before* and *after*, it seems to be necessary to chop off the localisation time during which the A event can occur to be that part of the period falling after the start of the B event. However, this characterisation faces two problems.

First, it is not empirically correct. As Moens and Steedman (1988) have pointed out, in some A *when* B sentences, the A event can precede the start of the B event, e.g.

(29) When they built that bridge, they hired a competent architect.

Here, the architect is hired before any building starts on the bridge.

Second, it is not possible to chop off the localisation time in this way. With **before** and **after**, the main event time was used to restrict the subordinate localisation time. But here, we need to use the subordinate event time to restrict the main clause localisation time. If the subordinate clause it tensed (which it is), this is not possible. The subordinate tense will select some new event time. But unless there is some device for dynamically passing the value of new event time back up into the main clause (c.f. Groenendijk and Stokhof's (1991) dynamic quantifiers), the subordinate event time cannot be used to restrict the main localisation time.

These two problems have a common solution, which moreover serves to explain the frequently observed similarity between *if* and *when*. Unlike *before* and *after*, *when* does not relate the times at which two event occur, but instead the times at which it is known or verified that the two events occur. In (29), hiring of the architect succeeds the time at which it is arranged that the bridge will be built. That is, foreknowledge about the subordinate event is possible, and the main event succeeds the first time of verification, which precedes the subordinate event.

How can we get hold of the earliest verification time for the tensed subordinate clause? In conditionals, minimal information extensions were used to identify the earliest verification times. Something similar will do the job here, though as we will see, it differs slightly from the minimal extensions used for conditionals.

Let us define a new form of minimal extension:

- $s' \sqsupseteq^{\phi,a,e,l}_{v,l'} s$ iff
  (a) $s' \sqsupseteq_v s$,
  (b) l' is a final subinterval of l that contains v
  (c) $v \leq a$
  (d) $s', a, v, e, l' \models \phi$, and
  (e) there is no s'', v'' or l'' such that
    (i) $v'' \leq v$
    (ii) $s \sqsubseteq_{v''} s'' \sqsubseteq_{v''} s'$
    (iii) l'' is a final subinterval of l containing v'', and is strictly contained in l'

That is, s' is a minimal extension of s verifying that the assertion ϕ holds, v is the first time at which this assertion is verified (and may precede the assertion time), and l' is the smallest final subinterval of l during which the assertion holds. Typically, l' will start either at v or the beginning of the event referred to by ϕ, whichever is earlier. This revised notion of a minimal extension in effect returns both an earliest verification time and a latest localisation time for the assertion.

We can now describe the effects of *when* as follows:

- $s, a, v, e, l \models \mathbf{when}(\phi, \psi)$ iff
  $\exists s', v', l'$ such that $s' \sqsupseteq^{\phi,a,e,l}_{v,l'} s$, $s' \approx s$, and $s, a, v, e, l' \models \psi$

If foreknowledge applies to ϕ, then ψ holds at some time after it is first verified that ψ will hold. If foreknowledge is not possible, then ψ holds at some time after ϕ begins to hold. Since the minimal extension, s' has to be informationally equivalent to s, it also follows that **when** expresses a non-conditional connection between the two events.

Although **when** is non-conditional, since it says what holds in the current information state rather than what will happen in extensions of the state, it is more closely related to the conditional than either of **before** or **after**. Just as conditionals make the assertion of the consequent dependent on the verification of the antecedent, so **when** makes the main clause assertion dependent on the verification of the subordinate assertion. Unlike conditionals, **when** on its own does not lead to deictic shift in the subordinate clause[3].

A second difference is that **when** does not universally quantify over information extensions. This means that **when** has to be used in conjunction with a habitual main clause tense to get habitual conditional-like sentences such as

(30) a. When a dog has a wet nose it is healthy.
     b. Lizards like it when it is hot.

[3]Condition (c) in the definition of $s' \sqsupseteq^{\phi,a,e,l}_{v,l'} s$ ensures that present tenses do not get futurate readings. This is, admittedly, something of an *ad hoc* restriction, and could be lifted.

The same is true of habituals with the connectives *before* and *after*

(31) John (always) takes a shower before/after he swims.

In these examples, the futurate subordinate tenses are the result of the main habitual tense, and not of the connectives.

It is worth comparing the treatment of *when* above with Moens and Steedman's (1988). They suggest that in A *when* B sentences the main A event coincides with either the preparatory process of the subordinate B event or a state resulting from the culmination of that event. In a sentence like (29), hiring the architect takes place during the initial preparation for building the bridge.

The notion of a preparatory process and a resultant state is in some ways vague. With accomplishments like building a bridge, the preparatory process is normally taken to consist of such things as digging the foundations, constructing the bridge supports and so on. This reaches a culmination point at which the bridge is finally built, and thereafter there is a resultant state in which a bridge is available for use. In other words, the event of building a bridge has a temporal duration, and what goes on during this period counts as a preparatory process leading up to the culmination.

Under this construal of what a preparatory process is, it is far from clear whether hiring an architect before work starts on the bridge really constitutes part of the bridge building event. When foundations are being laid one can truly say that a bridge is in the process of being built, but this is less obviously true when the bridge is just in the planning stage. To make the hiring occur during the preparatory process, the notion of what it is to be a preparatory process has to be more widely construed. But once this is done, it can be difficult to determine what does and what does not count as a preparatory process.

Appealing to the possibility of foreknowledge in (29) avoids some of this vagueness. Hiring the architect does not form part of the actual bridge building event, although it does occur during some kind of preparatory planning phase which does not constitute part of the actual bridge building.

As it stands, the account I have given for *when* does suffer one apparent defect. In A *when* B sentences where A occurs after B, it is invariably the case that A occurs immediately after B, although what counts as immediate succession depends on context — compare

(32) When the dinosaurs became extinct, the mammals became dominant.

(33) When John came in, Bill looked up.

Under Moens and Steedman's treatment, immediate succession is the result of the A event occurring during the resultant state of the B event. Again, a certain amount of vagueness attaches to what counts as a resultant state. Some events do not have an obvious resultant state, e.g. sneezing or blinking, other than perhaps the state of the event having taken place. Yet such events can be referred to in *when* sentences (e.g. *When John blinked, his contact lens fell out*). Presumably, therefore, one does not want to treat the resultant state as necessarily being lexicalised (though in some cases, e.g. putting up a tent, it will be). Following up on Poznanski's suggestion for *before* one might therefore associate a pragmatic directive with **when** saying that: in the current context, the subordinate B event should allow one to infer the existence of some resulting state, and this state should persist at least until the time at which the A event occurs.

4.5.3 Oblique Contexts and Sequence of Tense

Tensed sentential complements to certain verbs often exhibit something like deictic shift. For example

(34) I bet that Mary wins the prize.

(35) John said that Bill cut down the tree.

In (34) the present tense subordinate clause refers to a future event. In (35), the past tense subordinate clause can only refer to an event that is past with respect to John's reported utterance.

The futurate present in (34) is similar to futurate presents found in conditional antecedents. The verb *bet* sets up a new verification time that is not identical to the assertion time, and the subordinate clause is evaluated relative to the original assertion time and the new verification time.

Perhaps a verb like *say* sets up a new assertion and verification time for the subordinate clause identical to the main clause event time. This would ensure that a subordinate past tense has a past-in-the-past interpretation. Unfortunately, this cannot be right for *say*. As Enç (1987) points out, present tenses within these construction still refer to the present time:

(36) John said that Mary is pregnant.

But if *said* shifts the deictic to a point in the past, it should be possible to interpret (36) as meaning that John said that Mary was pregnant at some point in the past (and has subsequently given birth, perhaps).

There are three possible reactions to sentences like (35) and (36), only one of which is remotely tenable.

1. Wide scope subordinate tense The present tense in (36) shows that subordinate tense is not within the scope of the sentential complement verb. This can be either because just the tense is given wide scope, or because the whole clause is given wide scope.

If it is the whole clause that has wide scope, the utterer of (36) concurs with John in that Mary is pregnant. Rigter (1985) claims that this is what a sentence like (36) actually means: that Mary is pregnant, and that John said so. It is hard to maintain this position in light of examples like

(37) John said that he is hungry, but I know he's lying because he has just eaten a huge meal.

This leaves the option of just giving the subordinate tense wide scope, but leaving the rest of the clause within the scope of the sentential complement verb. While mechanisms have frequently been proposed for raising the scope of quantified noun phrases (e.g. Cooper (1983), Lewin (1990), Pereira (1990), and Alshawi and Crouch (1992), to mention just a few), it is not clear how they would apply to tenses. The usual assumption is in fact that the relative scopes of tenses and other intensional operators are fixed, and it is only the scope of noun phrases that may be altered. Moreover, supposing that some suitable mechanism could be found to raise the scope of the subordinate present tense, some way must be found to block its application to the subordinate past tense in (35). Otherwise, Bill cutting down the tree and John saying that Bill did so will both happen in the past, but John can be interpreted as making the claim before the tree is cut down. It is almost certain that no non *ad hoc* principle could be found to block this possibility.

2. 'Holes' in the localisation time Given the way that **before** and **after** manipulated the localisation time to control the temporal reference of subordinate tense, perhaps sentential complement verbs do something similar. That is, instead of resetting the deictic centre of the subordinate tense, they merely limit the localisation time to be used in conjunction with it. Suppose that *say* chops out the portion of the localisation time lying between the main clause event time and the assertion (speech) time. This would limit subordinate past tenses to referring to times prior to the main clause event, while still allowing subordinate present tenses to refer to times in the present. This suggestion is in fact quite similar to a proposal made by Enç (1987), though using a rather different framework.

Unfortunately, this will not work for sentences like

(38) John thought that the lecture was tomorrow, not today.

Here, the subordinate past tense refers to a (supposedly planned) future event. This is just not compatible with an unshifted past tense, no matter how the localisation period is hacked about. (Comrie (1986) and Hornstein (1991) also comment on the problems raised by this kind of sentence).

3. Sequence of Tense The best solution is in some ways the oldest. A verb like *say* or *think* subcategorises for complement clauses displaying one of three tenses: present, sequenced present, and sequenced past. The sequenced present and past tenses have a morphological realisation on the verb that is identical to that of the ordinary past tense. Semantically, they differ in that they are deictically centred on the event time rather than the assertion time. More specifically, letting **s-pres** and **s-past** be tense operators corresponding to the sequenced present and past tenses:

- $s, a, v, e, l \models \textbf{s-pres}(\phi)$ iff $e < a$ and $s, e, e, e, l \models \textbf{pres}(\phi)$.
- $s, a, v, e, l \models \textbf{s-past}(\phi)$ iff $e < a$ and $s, e, e, e, l \models \textbf{past}(\phi)$.

A sentence like *John said Mary was pregnant* is ambiguous under this proposal, depending on whether the tensed verb *was* is in the sequenced present or the sequenced past. The ambiguity corresponds to two different utterances that might be being reported:

John said "Mary is pregnant". (Sequenced Present)
$\textbf{past}(say(j, \textbf{s-pres}(pregnant(m))))$
John said "Mary was pregnant". (Sequenced Past)
$\textbf{past}(say(j, \textbf{s-past}(pregnant(m))))$

This solution is still not completely ideal. As Comrie (1986) observes, sequence of tense only applies when the main verb has past time reference. When the main verb is present tensed, the tense in the subordinate clause is subject to its usual interpretation. This is why the semantics of the **s-pres** and **s-past** operators specify that the initial event e time precedes the initial verification time a. Without this restriction, a sentence like

(39) John says that Mary was pregnant.

could be interpreted as having a sequenced present tense complement clause, meaning that what John says is that Mary *is* pregnant. But even then, the trick of stating that e must precede a only means that the sequenced interpretation of (39) only reports John saying something that is trivially false. While we might employ a principle of interpretive charity saying that peoples' utterances should not be construed in a way that is obviously false, it is not so obvious that the same principle can be applied to utterances reported by other utterances.

Another way out of this difficulty would be to make the subcategorisation of a verb like *say* depend on the (syntactic) tense of the verb. However, it is to say the least unusual to make verb subcategorisations dependent on the tense of the verb.

Deictic Shift

Although the discussion above is ultimately inconclusive, one thing is clear from it. The interpretation of tenses in complement clauses does not arise through the process of deictic shift. (Comrie (1986) comes to the same conclusion.) A slight exception to this is a verb like *bet*, where a secondary form of deictic shift in present tense complements is permitted by sundering the identity between the complement's assertion and verification times.

There are, however, cases where deictic shift does occur in complement clauses, e.g.

(40) John is not going to the lecture. But tomorrow he will say that he did. You wait and see.

Here, the past tense *did* in the second sentence refers to a future time. This has nothing to do with the sentential complement verb *say*. The deictic shift is caused by the modal *will*.

4.6 Conclusions

This chapter has shown how three (epistemic) modal operators, **must**, **may** and **will**, can be defined within the semantic framework presented in Chapter 3. On the assumption that the modal auxiliaries *must, may* and *will* correspond to these operators plus a sentential present tense, the patterns of temporal reference in modalised indicative conditionals can be accounted for. The effect that these modals have on subordinate tenses falls out of the treatment automatically. Other subordinate tense constructions have also been discussed. Some of these involve deictic shift of a similar kind to the caused by modals and conditionals (e.g. the futurate present). Others exhibit deictic shift only in the presence of a wide scope modal. And others are not due to deictic shift at all, but to a form of sequence of tense. Sequence of tense will play an important role in the discussion of subjunctive conditionals in the chapter after next. But before moving onto this, the next chapter takes a step back to review the logical properties of the conditional and modal operators that have been defined in the last two chapters.

Chapter 5

A Logic of Verified and Unverified Assertion

The preceding chapters have dealt almost exclusively with the temporal properties of modals and conditionals. In this chapter, we take a step back and comment on the logical properties that result from the approach that has been followed.

The chapter starts (Section 5.1) by presenting intuitionistic logic as a logic of verified truth. Section 5.2 shows how intuitionistic logic can be extended to give a logic of verified and *un*verified truth or assertion. The model theory that results from this extension corresponds to the information models introduced in Chapter 3. A proof theory is presented, which is shown (in Appendix C) to be sound and complete with respect to the intended semantics. The proof theory does not take account of tense operators. This is because the logic of verified and unverified assertion is in effect the internal logic of assertion and of information states. The role of tense is to take utterances of sentences and convert them into temporally specific assertions. It is only once this is done that the internal logic kicks into action.

Section 5.3 describes how three forms of epistemic modality, corresponding roughly to *must*, *may* and *will* can be defined in terms of two forms of negation present in the logic. Iterations of modalities with negation and conditionals are discussed. Section 5.4 suggests how non-epistemic modalities may be dealt with.

5.1 Intuitionism

Intuitionistic logic[1] is obtained from classical dropping the inference rules of *reductio ad absurdum* (or $\neg$-elimination, $\neg E$) from the classical natural deduction system.

$$\neg E \frac{\Gamma, \neg\phi \vdash \bot}{\Gamma \vdash \phi}$$

(i.e. If Γ combined with $\neg\phi$ proves an absurdity, $\bot$, then Γ proves ϕ). All the other introduction and elimination rules are retained.

Intuitionistic logic is a subset of classical logical: every intuitionistic theorem is classically valid, although the converse does not always hold. Intuitionism is notorious for its rejection of the law of the excluded middle ($\nvdash \phi \vee \neg\phi$), and the equivalence between a proposition and its double negation ($\nvdash \phi \leftrightarrow \neg\neg\phi$).

The loss of these theorems can be motivated on epistemic grounds. The assertion sign $\vdash$ in intuitionistic logic is not to be paraphrased, as in classical logic, as saying merely that the proposition following it is true (given the truth of any propositions that may occur on the left hand side of the sign). Instead, the paraphrase is that the proposition is 'known' to be true, that there is a concrete way of demonstrating its truth.

Intuitionistically, $\vdash \neg\neg\phi$ can be read as saying that one can demonstrate that there are no demonstrations that ϕ is false, i.e. that ϕ can only be true. But this is subtly weaker than giving a demonstration that ϕ *is* true. A similar weakness between modal and non-modal assertions in natural language can be observed: compare someone saying *I must have met you somewhere before* and *I have met you somewhere before*. The non-modal sentence suggests a direct memory of the meeting, whereas the modal sentence suggests more tenuous, indirect evidence for the assertion.

Classically, $\vdash \phi \vee \neg\phi$ can be read as saying that either ϕ is true or it is false. Intuitionistically, it can be read as saying that either ϕ is known to be true or that ϕ is known to be false. Clearly, there are some things like Goldbach's conjecture or Fermat's last theorem which are neither known to be true or known to be false. Such propositions invalidate the intuitionistic excluded middle. Classically, however, one could argue that even though their truth or falsity is not currently known, nevertheless they are either true or false, and the classical excluded middle holds.

A militant intuitionist who rejects 'the whole notion of objectively determined truth-values, independent of our capacity for recognising them' (Dummett 1977), would reject this last, classical, response. We need not follow in this. It is reasonably accurate to characterise intuitionism as being concerned with knowledge-values and verification conditions

[1]Most of the material in this section is taken from van Dalen's (1986) survey of the subject, plus elements from (Gabbay 1981) and (Fitting 1983).

rather than with truth-values and truth conditions. The militant intuitionist denies that there is anything more to truth than what is furnished by verification, and thus identifies truth and verification conditions. But one does not have to make this identification to see intuitionism as providing a useful logic of verification.

5.1.1 The Creative Subject

One way of bringing out the epistemic nature of intuitionism, and of introducing the ideas behind Kripke semantics for intuitionism, is in terms of a 'creative subject'.

The creative subject is an agent that extends its knowledge and the universe of objects it knows about over the course of time. At each moment, t the subject has a stock of sentences, Σ_t, it has established as true and a stock of objects, D_t, it has encountered or otherwise established as existent. The stock of sentences and objects at a time t constitute the subject's information state at time t.

As time goes by, the subject finds out more, and adds further sentences and further objects to its information state. At any time, the subject will have a variety of choices open to it about where to look for new information next (including the choice of not looking anywhere, but not of forgetting). Depending on what the subject decides to do at time t, the information state may grow into one of a number of possible state at time $t+1$. Thus there is a natural (partial) order imposed over the subject's possible information states, reflecting the ways in which the subject's information can accumulate.

In any information state, s, the subject can be said to know $\phi \wedge \psi$ in s if it knows both ϕ and ψ in s. It knows $\phi \vee \psi$ in s either if it knows ϕ in s or if it knows ψ in s. An implication $\phi \rightarrow \psi$ can be known in s without either of ϕ or ψ being known there, but it must be known in s that in any later state r in which ϕ is known, ψ will also be known there. $\forall x.\, \phi(x)$ is known in s if it is known that in any future state (as well as the current one), $\phi(o)$ is known for all objects o existing in that state. $\exists x.\, \phi(x)$ is known at s if there is an object o in s for which it is known in s that $\phi(o)$.

Obviously, there is much in common between an intuitionistic creative subject and the kind of information gatherer implicit in the previous chapters.

5.1.2 Kripke Models for Intuitionistic Logic

A Kripke model for the intuitionistic propositional calculus (IPC) is a triple M

$M = \langle S, \sqsubseteq, \Vdash \rangle$

where S is a set of information states
$\sqsubseteq$ is a partial order over S, and
$\Vdash$ is a forcing relation

The forcing relation relates formulas and information states, and obeys the following properties for all states s

- If $s \Vdash p$ and $s \sqsubseteq s'$ then $s' \Vdash p$ for atomic sentences p.
- $s \Vdash \neg\phi$ iff for all $s' \sqsupseteq s$, $s \nVdash \phi$
- $s \nVdash \bot$
- $s \Vdash \phi \wedge \psi$ iff $s \Vdash \phi$ and $s \Vdash \psi$
- $s \Vdash \phi \vee \psi$ iff $s \Vdash \phi$ or $s \Vdash \psi$
- $s \Vdash \phi \rightarrow \psi$ iff for all $s' \sqsupseteq s$, $s \nVdash \phi$ or $s \Vdash \psi$

We can read $s \Vdash \phi$ as saying that s forces ϕ, or alternatively as ϕ is true or is known in state s. The first condition says that if an atomic sentence is forced in one state, it is forced in all states extending it. The second says that a falsity in a state s isn't just a matter of not being forced in s, but of not being forced in s or any state extending it. If a sentence ϕ is not forced in s but is forced in some state extending s, it follows that ϕ is neither true nor false in s. The third condition says that the absurd proposition, $\bot$, is never forced in any state, and hence is false in all states. The final condition says that for an implication $\phi \rightarrow \psi$ to be true in s it is not necessary for ϕ or ψ to be true or false in s. Instead it is necessary that for s or any state extending s, if ϕ is forced in that state, then ψ is also forced in that state. It can be shown that the forcing of complex sentences is determined entirely by the way that atomic sentences are forced in information states.

Comparison to Modal Logic

For anyone familiar with modal logic, $\langle S, \sqsubseteq \rangle$ is just a Kripke frame for the system S4. However, the forcing relation is rather different from the evaluation relation familiar from modal logics. The modal evaluation relation, $\models$, is a relation on possible worlds (the modal analogue of information states) and sentences, just as the forcing relation is. And as with the forcing relation, the valuation of complex sentences is determined entirely by the valuation of atomic sentences. But unlike intuitionistic logic, in modal logic a sentence is false in a possible world simply if it is not true in that world. So in every possible world, every sentence is either true or false. Another difference is that if ϕ is true in a world w, and $w \sqsubseteq w'$, it does not follow that ϕ is true in w'. A third difference is between the languages of modal and intuitionistic logic; intuitionistic logic does not contain the modal operators $\Box$ and $\Diamond$.

To facilitate a more precise comparison between intuitionistic propositional logic and propositional S4, the properties of the modal evaluation relation are spelled out in more detail:

- $w \not\models \bot$
- $w \models \neg\phi$ iff $w \not\models \phi$
- $w \models \phi \wedge \psi$ iff $w \models \phi$ and $w \models \psi$
- $w \models \phi \vee \psi$ iff $w \models \phi$ or $w \models \psi$
- $w \models \phi \supset \psi$ iff $w \not\models \phi$ or $w \models \psi$
- $w \models \Box\phi$ iff for all $w' \sqsupseteq w$, $w' \models \phi$
- $w \models \Diamond\phi$ iff for some $w' \sqsupseteq w$, $w' \models \phi$

From a modal point of view, it is apparent that intuitionistic negation and implication have a modal feel to them; they depend on what is forced in other states, just as $\Box$ and $\Diamond$ depend on what is true in other possible worlds. Fitting (1983) refers to conjunction and disjunction as 'regular' connectives, and negation and implication as 'special' connectives in recognition of this fact.

Comparison to Partial Logic

The failure of the law of the excluded middle should not mislead one into thinking that intuitionistic logic is a partial or three-valued logic in the sense of Langholm (1987) or Fenstad *at al* (1987). Partial logics make use of partial valuation functions. That is, any formula is assigned either the value 1 (true), 0 (false) or nothing at all (undefined / unknown). By contrast, the intuitionistic forcing relation is an all or nothing affair. A formula is either forced by an information state (1), or it isn't (0). There is nothing in between.

It is clear, however, that intuitionistic logic can nevertheless deal with partial states of information. It is precisely the semi-modal nature of negation that allows intuitionistic logic to do this without a partial forcing relation. Intuitionistic logic would seem to provide a promising (though as far as I am aware, largely neglected) tool for the analysis of partial information. (It would be interesting to see whether an intuitionistic reformulation of Situation Theory could be made to work.)

5.1.3 Proof Theory For IPC

The axiomatisation of the intuitionistic propositional calculus is more complicated than that for the classical calculus, but the natural deduction / sequent calculus formulation is simpler. It simply involves dropping a single classical inference rule.

Sequent Calculus for IPC

A sequent calculus formulation of IPC has the following introduction and elimination rules:

$$\wedge I\,\frac{\Gamma \vdash \phi;\ \Gamma \vdash \psi}{\Gamma \vdash \phi \wedge \psi} \qquad \wedge E\,\frac{\Gamma \vdash \phi \wedge \psi}{\Gamma \vdash \phi}$$

$$\vee I\,\frac{\Gamma \vdash \phi}{\Gamma \vdash \phi \vee \psi} \qquad \vee E\,\frac{\Gamma \vdash \phi \vee \psi;\ \Gamma, \phi \vdash \chi;\ \Gamma\psi \vdash \chi}{\Gamma \vdash \chi}$$

$$\rightarrow I\,\frac{\Gamma, \phi \vdash \psi}{\Gamma \vdash \phi \rightarrow \psi} \qquad \rightarrow E\,\frac{\Gamma \vdash \phi;\ \Gamma \vdash \phi \rightarrow \psi}{\psi}$$

$$\neg I\,\frac{\Gamma, \phi \vdash \bot}{\Gamma \vdash \neg\phi} \qquad \bot\,\frac{\Gamma \vdash \bot}{\Gamma \vdash \phi}$$

The rules are to be read as saying, for example in $\wedge$-elimination, that if Γ proves $\phi \wedge \psi$, then Γ proves ϕ. Assumptions are introduced and eliminated in, for example, $\rightarrow$-introduction. This says that if Γ combined with ϕ proves ψ, then Γ on its own proves $\phi \rightarrow \psi$. The rules for $\wedge$ and $\vee$ are symmetric over conjuncts / disjuncts in the usual way.

The introduction rule for negation is tantamount to defining $\neg\phi$ as $\phi \rightarrow \bot$. The classical propositional calculus is obtained by adding the rule of *reductio ad absurdum*, or $\neg$ elimination:

$$\neg E\,\frac{\Gamma, \neg\phi \vdash \bot}{\Gamma \vdash \phi}$$

Note that the rule for $\neg$ introduction allows us to infer $\neg\neg\phi$ under these circumstances, but not ϕ.

The rule $\bot$ may be dropped from IPC to give *minimal* logic. In minimal logic, a falsehood does not entail everything, which is an intuitively desirable result. Less desirable is the fact that to model minimal logic, one must allow some information states to force the absurd proposition $\bot$.

Axiomatisation for IPC

The axiomatisation of IPC is not as revealing as its natural deduction formulation. The axiomatisation is given here for the sake of form, and we rapidly move on to list some important theorems and non-theorems of IPC.

IPC has one rule of inference, *modus ponens* or $\rightarrow$ elimination, and the following axioms

$$\begin{array}{l}
\phi \rightarrow (\psi \rightarrow \phi) \\
(\phi \rightarrow \psi) \rightarrow ((\phi \rightarrow (\psi \rightarrow \chi)) \rightarrow (\phi \rightarrow \chi)) \\
\phi \rightarrow (\psi \rightarrow \phi \wedge \psi) \\
\phi \wedge \psi \rightarrow \phi \\
\phi \rightarrow \phi \vee \psi \\
(\phi \rightarrow \chi) \rightarrow ((\psi \rightarrow \chi) \rightarrow (\phi \vee \psi \rightarrow \chi)) \\
(\phi \rightarrow \psi) \rightarrow ((\phi \rightarrow \neg\psi) \rightarrow \neg\phi)
\end{array}$$

More interestingly, the following theorems of CPC are *not* theorems of IPC

$$\begin{array}{l}
\phi \vee \neg\phi \\
\neg\phi \vee \neg\neg\phi \\
\neg\neg\phi \rightarrow \phi \\
\neg(\phi \wedge \psi) \rightarrow \neg\phi \vee \neg\psi \\
(\neg\psi \rightarrow \neg\phi) \rightarrow (\phi \rightarrow \psi) \\
(\phi \rightarrow \psi) \rightarrow (\neg\phi \vee \psi)
\end{array}$$

The following, however, *are* theorems of IPC:

$$\begin{array}{l}
\neg\neg(\phi \vee \neg\phi) \\
\neg(\phi \wedge \neg\phi) \\
\phi \rightarrow \neg\neg\phi \\
\neg\phi \leftrightarrow \neg\neg\neg\phi \\
\neg(\phi \vee \psi) \leftrightarrow \neg\phi \wedge \neg\psi \\
\neg\phi \vee \neg\psi \rightarrow \neg(\phi \wedge \psi) \\
(\phi \rightarrow \psi) \rightarrow (\neg\psi \rightarrow \neg\phi) \\
(\neg\phi \vee \psi) \rightarrow (\phi \rightarrow \psi)
\end{array}$$

A number of morals can be drawn from these theorems and non-theorems.

First, a number of classical equivalences work in one direction but not the other. The most striking is contraposition: $(\phi \rightarrow \psi) \leftrightarrow (\neg\psi \rightarrow \neg\phi)$. Contraposition works in the left to right direction, but not right to left. Similarly with the classical equivalence between $\phi \rightarrow \psi$ and $\neg\phi \vee \psi$, and with double negation A rule of thumb for remembering in which directions the classical implications still hold is that disjunction and positive propositions (regular propositions) are stronger than negated and implicational propositions (special propositions). A disjunction will support an implication, and a positive proposition a double negation, but not vice versa.

Second, the intuitionistic connectives are independent. In classical logic, one can start from implication and negation, or implication and $\bot$, or implication and conjunction, and define the remaining connectives. In intuitionistic logic this is not possible.

The Gödel Translation

The third moral is that all classical principles are retained in an attenuated, doubly negated form: all of classical logic can be embedded within intuitionistic logic by means of the Gödel translation.

The Gödel translation of a formula, $G(\phi)$ is defined as

$G(\bot) = \bot$
$G(p) = \neg\neg p$ for all other atomic formulas
$G(\phi \wedge \psi) = G(\phi) \wedge G(\psi)$
$G(\phi \vee \psi) = \neg(\neg G(\phi) \wedge \neg G(\psi))$
$G(\phi \to \psi) = G(\phi) \to G(\psi)$
$G(\neg \phi) = \neg\neg\neg G(\phi)$

It can be shown that

$$\Gamma \vdash_C \phi \text{ iff } \Gamma \vdash_I G(\phi)$$

Intuitionistic double negation acts as a form of necessity: $\neg\neg\phi$ means that it cannot be the case that ϕ is false, i.e. that ϕ is necessarily true. Classical propositions are intuitionistic propositions that are necessarily true in this sense; i.e. classical truth corresponds to what cannot intuitionistically be false.

This line of thought suggests defining a kind of necessity operator in terms of double negation.

Modal Intuitionistic Logic

A state s forces an intuitionistic double negation, $\neg\neg\phi$, under the following circumstances:

$$s \Vdash \neg\neg\phi \text{ iff } \forall s' \sqsupseteq s : \exists s'' \sqsupseteq s' : s'' \Vdash \phi$$

That is, it does not say that ϕ holds in all states extending s, but rather that for all states extending s there is yet a further state extending that in which ϕ does hold. In other words, ϕ must eventually turn out to hold, though it does not have to hold straight away.

This is rather different from the notion of necessity familiar from modal logics, where a proposition is necessarily true in a world w if it is true in all worlds accessible from w. But were we to take $\sqsubseteq$ as the accessibility relation, and states as worlds, adopting this definition of necessity would make $\Box\phi$ and ϕ equivalent — necessity would collapse. Consider:

$$s \Vdash \Box\phi \text{ iff } \forall s' \sqsupseteq s : s' \Vdash \phi$$

Since $s \sqsupseteq s$, this would mean that

$$s \Vdash \Box\phi \Rightarrow s \Vdash \phi$$

Since the forcing relation is monotonic, if $s \Vdash \phi$ and $s' \sqsupseteq s$, then $s' \Vdash \phi$, meaning that

$$s \Vdash \phi \Rightarrow s \Vdash \Box\phi$$

Double negation is the nearest, non-trivial, semantic notion of necessity we can get for intuitionistic logic.

Perhaps one can treat natural language necessity in terms of intuitionistic double negation. That is

$$\Box\phi =_{df} \neg\neg\phi$$

This would have the possible advantage of allowing one to handle necessity without recourse to a modal logic.

Unfortunately, matters do not proceed simply. The next question to ask is how can *possibility* be represented in non-modal intuitionistic logic? And the answer is that it can't. The dual of necessity, $\neg\Box\neg\phi$, is $\neg\neg\neg\neg\phi$ which is equivalent to $\neg\neg\phi$, i.e. $\Box\phi$. So one cannot get possibilities that way.

Semantically, the conditions for possibility are fairly obvious

$$s \Vdash \Diamond\phi \text{ iff } \exists s' \sqsupseteq s : s' \Vdash \phi$$

But no such semantic conditions can be constructed from the connectives $\neg, \wedge, \vee$ or $\rightarrow$. Possibility is an 'out of state' connective, and so would have to be constructed using the out of state connectives $\rightarrow$ and $\neg$. But the semantic definitions of these connectives involve universal quantification over states. No matter how we pile these connectives up, one cannot obtain an expression whose semantics corresponds to something with an outer existential quantifier over states.

One solution is to introduce a new modal operator, $\Diamond$, with the semantics suggested above. But if we do this, the dual of $\Diamond$ would have the following semantics

$$s \Vdash \Box\phi \text{ iff not } \exists s' \sqsupseteq s : s' \not\Vdash \phi$$
$$\text{i.e. } s \Vdash \Box\phi \text{ iff } \forall s' \sqsupseteq s : s' \Vdash \phi$$

But as we have just seen, $\Box$ defined in this way collapses into triviality[2].

Modal extensions of intuitionistic logic have been proposed (e.g. Gabbay 1981, Plotkin and Stirling 1986, Bozic and Dosen 1984 — see Section 5.4) , but escape the collapse of $\Box$ into triviality by distinguishing between the accessibility relation between states and the informational ordering over states. Thus a state s' accessible from s need not be an extension of s, and hence the monotonicity of forcing does not carry across the accessibility relation.

Another solution, pursued below, is to introduce a new form of essentially classical negation, $\sim$. This is defined as $s \Vdash \sim \phi$ iff $s \not\Vdash \phi$. We can then define

$$\Diamond\phi =_{df} \sim \neg\phi$$

But unless one is careful, it turns out that the obvious dual of $\Diamond$ thus defined ($\sim\sim \neg \sim \phi$ $= \neg \sim \phi$) collapses into triviality in exactly the same way as before.

5.1.4 Negation and Intuitionistic Logic

It is possible to introduce other forms of negation besides $\neg$. I will briefly mention three cases: strong negation, $-$, classical negation, $\sim$, and a revision to $\neg$ brought about in Gabbay's intuitionistic logic **KC**.

IPC with Strong Negation

Intuitionistic negation can tell us if something will never become true, but does not distinguish between the cases where this is so because (a) the matter is never determined, or (b) it is already false. Strong negation allows one to state that something has already been shown to be false.

To accommodate this, we need positive and negative versions of the forcing relation, and consequently enter the realm of three-valued intuitionistic logic. The positive forcing

[2]Note that we are using classical negation in the semantic definitions.

relation says when something is established as true (verified), and the negative relation when it is established as false (falsified).

Kripke models for intuitionistic logic with strong negation are quadruples

$$M = \langle S, \sqsubseteq, \Vdash \dashv\!| \rangle$$

where S is a set of information states
$\sqsubseteq$ is a transitive, reflexive) order over S
$\Vdash$ is a positive forcing relation
$\dashv\!|$ is a negative forcing relation

The positive and negative forcing relations behave as follows

- If $s \Vdash p$ and $s \sqsubseteq s'$ then $s' \Vdash p$
  for atomic sentences p.
  If $s \dashv\!| \ p$ and $s \sqsubseteq s'$ then $s' \dashv\!| \ p$
  for atomic sentences p.
- $s \nVdash \bot$
  $s \dashv\!| \ \bot$
- $s \Vdash \neg\phi$ iff for all $s' \sqsupseteq s$, $s \nVdash \phi$
  $s \dashv\!| \ \neg\phi$ iff $s \Vdash \phi$
- $s \Vdash \phi \wedge \psi$ iff $s \Vdash \phi$ and $s \Vdash \psi$
  $s \dashv\!| \ \phi \wedge \psi$ iff $s \dashv\!| \ \phi$ or $s \dashv\!| \ \psi$
- $s \Vdash \phi \vee \psi$ iff $s \Vdash \phi$ or $s \Vdash \psi$
  $s \dashv\!| \ \phi \vee \psi$ iff $s \dashv\!| \ \phi$ and $s \dashv\!| \ \psi$
- $s \Vdash \phi \rightarrow \psi$ iff for all $s' \sqsupseteq s$, $s \nVdash \phi$ or $s \Vdash \psi$
  $s \dashv\!| \ \phi \rightarrow \psi$ iff $s \Vdash \phi$ and $s \dashv\!| \ \psi$
- $s \Vdash -\phi$ iff $s \dashv\!| \ \phi$
  $s \dashv\!| \ -\phi$ iff $s \Vdash \phi$

Note that the falsifying conditions for the two special connectives, $\neg$ and $\rightarrow$, are determined by what obtains in the current information state, whereas their positive conditions, as before, depend on what happens in later states.

One can provide an axiomatisation for IPC with strong negation by adding the following to the axioms for IPC (Gabbay 1981):

$-(\phi \rightarrow \psi) \leftrightarrow \phi \wedge -\phi$
$-(\phi \wedge \psi) \leftrightarrow -\phi \vee -\phi$
$-(\phi \vee \psi) \leftrightarrow -\phi \wedge -\phi$
$(-\phi \wedge \phi) \rightarrow \psi$
$-\neg\phi \leftrightarrow \phi$
$- - \phi \leftrightarrow \phi$
$-\phi \rightarrow \neg\phi$

In natural deduction terms, one does need a introduction rule for $-$. This is because it is possible to push $-$ inwards until it has scope over only atomic sentences or strongly negated atomic sentences. If p is atomic, then $-p$ can also be treated as atomic according to negative forcing. That is, strongly negative atomic facts are taken as given and do not need to be introduced. There is a $-$ elimination rule that just removes double strong negations however.

IPC with Classical Negation

I am not aware of work on the introduction of classical negation into intuitionism, which would be defined as the complement of the forcing relation:

- $s \Vdash \sim \phi$ iff $s \nVdash \phi$

Veltman's Data Semantics introduces a mixture of classical and strong negation, however. Applied to conjunctions and disjunctions, data semantical negation behaves like strong negation, but applied to implications and to itself, it behaves like classical negation.

The Logic KC

If we add one extra axiom governing negation to ordinary intuitionistic logic, we obtain Gabbay's logic, **KC**. The extra axiom is

$$\vdash \neg\phi \vee \neg\neg\phi$$

It has been shown that the extended axiomatisation is sound and complete with respect to intuitionistic frames where the informational ordering over states is linear as well as transitive and reflexive. The significance of this will become apparent shortly.

5.2 Verified and Unverified Assertions

We now return to the information models introduced in Chapter 3, but ignore tense while still taking account of the times at which formulas are verified. In these information models, each information state can be seen as a linearly ordered sequence of temporal 'snapshots' of the state, where different formulas are forced at different time points.

From the point of view of the preceding section, this means that information states now have some internal structure — states are sequences of sub-states. From the outside, we can ignore this internal structure. Thus a formula is asserted by a state iff it comes to be forced at any point within that state. Without delving into the inner structure of states, formulas assumed by a state will obey the intuitionistic principles outlined previously. But they also possess further properties through the internal structure of states. It turns out that the internal and external properties of states and formulas give rise to an intuitionistic logic within an intuitionistic logic, just as information states are sequences of states within states.

5.2.1 Semantics for Verified and Unverified Assertions

Information Models

An information model M is now a quintuple

$$M = \langle S, \sqsubseteq_t, T, \leq, V \rangle$$

where S is a set of information states, s
$\sqsubseteq_t$ is a relation in $S \times S \times T$
and is transitive and reflexive over S for any t
T is a set of time instants, t
$\leq$ is a (linear) temporal order over T, and
V is a valuation function

The valuation function V is a function from states, times and atomic sentences in some language L onto the (verification) values 1 or 0. This specifies for a state, time and sentence whether the sentence is verified at that time in that state.

We impose two monotonicity conditions on the valuation function V:

Monotonicity of direct verification ('in-state' monotonicity):
For every state s and atomic sentence p
$t_1 \leq t_2$ implies if $V(s, t_1, p) = 1$ then $V(s, t_2, p) = 1$

Monotonicity of information growth ('out-of-state' monotonicity):
If $s_1 \sqsubseteq_t s_2$ then for atomic sentences p
(a) $\{p \mid V(s_1, t, p) = 1\} \subseteq \{p \mid V(s_2, t, p) = 1\}$
(b) $\{p \mid \exists t : V(s_1, t, p) = 1\} \subseteq \{p \mid \exists t : V(s_2, t, p) = 1\}$

Monotonicity of direct verification says that once an atomic sentence is verified within a state, if remains verified in that state. Monotonicity of direct verification will turn out to apply to all sentences, and not just to atomic ones.

Monotonicity of information growth says that moving from s_1 to a more fully informed state s_2 at some time t verifies the same atomic sentences in s_2 at t as were verified in s_1 at t (condition (a)). It also says (condition (b)) that the sentences asserted in s_1 but not yet verified in at t continue to be asserted in s_2, though the times at which they becomes verified in s_2 may differ form those in s_1. Indeed, some unverified assertions in s_1 may count as already verified in s_2. In addition, s_2 may include verified and unverified assertions that were not present at all in s_1. If s_2 does not contain any extra assertions (i.e. the $\sqsubseteq_t$ becomes $\approx_t$ in (a) and (b)), then the two states are *informationally equivalent* at t. Monotonicity of information growth does not hold for all non-atomic sentences.

In addition, a convergence of verification property is imposed:

Convergence of Verification:
If $s_1 \sqsubseteq_{t_1} s_2 \sqsubseteq_{t_2} s_3$,
then there is a time t_3 such that $t_3 \geq t_1$, $t_3 \geq t_2$ and $\forall t_4 \geq t_3 \; s_1 \sqsubseteq_{t_4} s_3$

as well as the constraint on absurdity:

No Absurdity:
For no s or t is it the case that $V(s, t, \bot) = 1$

The other constraints on information models noted in Chapter 3 are no longer of relevance once the temporal reference of atomic sentences is fixed in advance.

Semantic Definitions of Connectives

We now give semantic definitions for a propositional language involving the connectives $\wedge, \vee, \rightarrow, \neg$ and $\sim$. These specify what is required for a sentence to be verified as true at a time t in a state s.

1. $s, t \Vdash p$ iff $V(s, t, p) = 1$ if p is atomic

2. $s, t \Vdash \phi \wedge \psi$ iff $s, t \Vdash \phi$ and $s, t \Vdash \psi$

3. $s,t \Vdash \phi \vee \psi$ iff $s,t \Vdash \phi$ or $s,t \Vdash \psi$

4. $s,t \Vdash \phi \rightarrow \psi$ iff $\forall t_1 \geq t, s_1 \sqsupseteq_{t_1}^{\phi,t} s : \exists t_2 \geq t_1$ such that $s_1, t_2 \Vdash \psi$

5. $s,t \Vdash \neg\phi$ iff $\forall t_1 \geq t, s_1 \sqsupseteq_{t_1}^{\phi,t} s : \exists t_2 \geq t_1$ such that $s_1, t_2 \Vdash \perp$

6. $s,t \Vdash \sim \phi$ iff $\forall t_1 \geq t : s, t_1 \not\Vdash \phi$

Minimal Information Extensions The semantic definitions for $\rightarrow$ and $\neg$ rely on minimal information extensions. Parallel to Chapter 3, a minimal information extension can be defined as follows:

- $s_1 \sqsupseteq_{t_1}^{\phi,t} s$ iff
  a) $s_1 \sqsupseteq_{t_1} s$
  b) $s_1, t_1 \Vdash \phi$, and
  c) $\nexists t_2, s_2$ such that $t \leq t_2 < t_1$, $s \sqsupseteq_{t_2} s_2 \sqsupseteq_{t_2} s_1$, and $s_2, t_2 \Vdash \phi$

That is, if s_1 is a minimal extension of s with respect to ϕ at time t, then s_2 is the first state extending s that verifies ϕ at the earliest time t_1.

Negation Two types of negation are defined: 'out-of-state' negation, $\neg$, and 'in-state' negation $\sim$. Out-of-state negation says that a sentence will never be verified in any future state at any future time. In-state negation says that a sentence will never be verified in the current state at any future time.

If a sentence $\sim \phi$ holds in a state s at a time t, the sentence ϕ is *not* asserted by state s. (Note that the monotonicity of direct verification will ensure that if $\sim \phi$ holds in s at any time t, it holds at all times). If $\sim\sim \phi$ holds in s at t, then it follows that at some later time t_1, ϕ will hold in s. Thus $\sim\sim \phi$ gives a way of saying that ϕ is asserted by a state, though not necessarily as yet verified in that state.

As well as referring to the two negations as in-state and out-of-state negation, we can also say that $\sim$ amounts to a denial of assertion, while $\neg$ amounts to an assertion of denial.

Monotonicity

A significant feature of the forcing relation in intuitionistic logic is its monotonicity: once a sentence is forced in one state, it remains forced in all subsequent states. This holds for all sentences.

Within the information models currently being dealt with, we need to consider two distinct kinds of monotonicity: in-state monotonicity, and out-of-state monotonicity. In-state monotonicity holds for all sentences, but out-of-state monotonicity holds only for a restricted set of *stable* sentences.

Stable sentences are defined below, and are approximately those whose sub-formulas contain no odd-numbered sequences of in-state negations — thus $\sim\sim \phi$ is stable, but $\psi \wedge \sim \phi$ is not. Unstable sentences are ones that state that a particular sentence is not asserted in a given information state. That a sentence is not asserted in one state does not preclude it from coming to be asserted in a later state.

In-State Monotonicity First, we show that all sentences satisfy in-state monotonicity:

In-State Monotonicity:
If $s, t \Vdash \phi$ then $\forall t_1 \geq t : s, t_1 \Vdash \phi$

This is a direct consequence of the condition of monotonicity of direct verification (in-state monotonicity for atomic sentences) that was imposed on the valuation function in information models. For conjunction and disjunction, it is easy to see that monotonicity carries across from monotonic sentences.

For in-state negation, note that $\sim \phi$ is only forced in state s at t if ϕ fails to be forced in s at any time t_1 at or after t. Assuming that ϕ itself behaves monotonically within the state, this means that ϕ is forced at no time within the state, and so $\sim \phi$ is forced at all times within the state. Exactly similar arguments apply to $\neg\phi$ — just ignore the reference to future information states in the semantic definition, and concentrate on reference to future times within the current state.

For implication, $\phi \rightarrow \psi$, assume both ϕ and ψ behave monotonically. Restricting attention to what must go on within the state if the semantic conditions for the implication hold at time t, we find that if ϕ holds at any time t_1 after t, then ψ must hold at some time t_2 after t_1. Since ϕ and ψ are monotonic, if this holds at t, it will also hold at all times t' after t.

Stable Sentences Next we define what stable sentences are.

- If p is atomic, then p is stable.
- If ϕ and ψ are stable, then $\phi \wedge \phi$ and $\phi \vee \psi$ are stable.
- $\phi \rightarrow \psi$ is stable if ψ is stable. (Otherwise, it is semi-stable.)
- $\neg\phi$ is stable.
- If ϕ is stable, then $\sim\sim \phi$ is stable.

- Anything not classified as stable by the above is unstable.

A piece of notation that will be used below is that if Γ is a set of formulas or sentences, then $\mathit{Stable}(\Gamma)$ denotes the stable subset of Γ.

Out-of-State Monotonicity It can be shown that all stable sentences satisfy out-of-state monotonicity:

> Out-of-State Monotonicity:
> If $s, t \Vdash \phi$ then $\forall s_1 \sqsupseteq_t s, t_1 \geq t : s_1, t_1 \Vdash \phi$

(Proof: by induction on the complexity of ϕ — given in Appendix C).

The converse does not quite hold. There are some unstable formulas that satisfy out-of-state monotonicity: $\sim \phi \vee \sim\sim \phi$ is an example. The reason for this is that some unstable formulas, like $\sim \phi \vee \sim\sim \phi$ are valid in all models, and others like $\phi \wedge \sim \phi$ are invalid. Both valid and invalid formulas obey out-of-state monotonicity by dint of holding (or failing to hold) in all information states at all times in all models. Conjunctions and disjunctions of stable formulas with these valid or invalid unstable formulas can give rise to monotonic unstable formulas that are not themselves valid or invalid.

Thus, the instability of a formula is no guarantee of its non-monotonicity. One could perhaps try to rectify this by adding other cases to the definition of stability in an attempt to classify all the unstable validities and invalidities as stable, e.g.

- $\sim (\phi \wedge \psi)$ and $\sim (\phi \vee \psi)$ are stable if $\sim \phi$ and $\sim \psi$ are stable.
- $(\phi \rightarrow \psi) \vee \sim (\sim\sim)^*(\phi \rightarrow \psi)$ is stable (and valid).
- $\neg\phi \vee \sim (\sim\sim)^* \neg\phi$ is stable (and valid).
- $\sim \phi \vee \sim (\sim\sim)^* \sim \phi$ is stable (and valid).
- $\phi \wedge \sim (\sim\sim)^* \phi$ is stable (and invalid).
- $\phi \rightarrow \sim (\sim\sim)^* \phi$ is stable (and invalid).
- $\sim \phi \wedge \neg \sim \phi$ is stable (and invalid).

(where $(\sim\sim)^*$ means zero or more iterations of the double negation). Unfortunately, I do not know if it is possible to give a complete enumeration of valid and invalid formulas involving $\sim$.

It turns out that none of this is necessary and that the first definition of stability will do. The stable subset of a set of formula, $\mathit{Stable}(\Gamma)$, is used in two places below. The first is in limiting the set of premises, Γ, to be used by certain rules of inference. The other rules of inference that are not subject to this limitation ensure that monotonic formulas in Γ but excluded from $Stable(\Gamma)$ can be brought into any proof by another route. The second place is in constructing canonical models for proving completeness (Appendix C). Here it is shown that the definition of stability has no ill-effects.

5.2.2 Extending Intuitionistic Logic

The semantic definitions provided above can be given (informal) justification in terms of preserving as much as possible of intuitionism for the logic of asserted sentences. This also serves to illustrate how the connectives $\wedge, \vee, \neg$ and $\rightarrow$ defined above correspond to ordinary intuitionistic connectives.

To derive our connectives of $\wedge, \vee, \neg$ and $\rightarrow$ from the intuitionistic ones, start by temporarily ignoring the internal structure of information states. That is, we are only concerned with those assertions supported by a state, and not with the time at which they become verified in that state.

A sentence ϕ is asserted by state s iff it is verified in s at some time t:

$$s \Vdash \phi \text{ iff } \exists t.\, s, t \Vdash \phi$$

Recall that the intuitionistic forcing conditions are, with the necessary modification to the atomic case:

- $s \Vdash p$ iff $\exists t.\, V(s,t,p) = 1$ if p is atomic
- $s \Vdash \phi \wedge \psi$ iff $s \Vdash \phi$ and $s \Vdash \psi$
- $s \Vdash \phi \vee \psi$ iff $s \Vdash \phi$ or $s \Vdash \psi$
- $s \Vdash \neg\phi$ iff $\forall s' \sqsupseteq s$: $s' \nVdash \phi$
- $s \Vdash \phi \rightarrow \psi$ iff $\forall s' \sqsupseteq s$: if $s' \Vdash \phi$ then $s' \Vdash \psi$

We now reinstate the suppressed reference to time according to the definition relating $s \Vdash \phi$ and $s, t \Vdash \phi$. This is the first step back towards defining the relation $s, t \Vdash \phi$.

- $s \Vdash p$ iff $\exists t.\, V(s,t,p) = 1$ if p is atomic
- $s \Vdash \phi \wedge \psi$ iff $\exists t_1 : s, t_1 \Vdash \phi$ and $\exists t_2 : s, t_2 \Vdash \psi$
- $s \Vdash \phi \vee \psi$ iff $\exists t_1 : s, t_1 \Vdash \phi$ or $\exists t_2 : s, t_2 \Vdash \psi$

- $s \Vdash \neg\phi$ iff $\forall s' \sqsupseteq s, \forall t$: $s', t \not\Vdash \phi$
- $s \Vdash \phi \rightarrow \psi$ iff $\forall s' \sqsupseteq s$: if $\exists t_1 . s', t_1 \Vdash \phi$ then $\exists t_2 . s' t_2 \Vdash \psi$

It can be shown that the monotonicity of direct verification extends to non-atomic sentences:

For all s, ϕ, if $t_1 \leq t_2$, then $s, t_1 \Vdash \phi \Rightarrow s, t_2 \Vdash \phi$

Given this form of monotonicity, we can simplify some of the above clauses. For conjunction, if ϕ is forced at t_1 it is forced at all times after t_1, and if ψ is forced at t_2 it is forced at all times after t_2. Therefore, if ϕ is forced at t_1 and ψ at t_2, then there will be some time t_3 later than either at which both are forced. Similarly with disjunction. With implication, if ϕ being forced at t_1 means that ψ is also forced at some time, then ψ is bound to be forced at a time at or after t_1 (even if it is also forced before). We can therefore reformulate the above clauses without loss of generality as:

- $s \Vdash p$ iff $\exists t . V(s, t, p) = 1$ if p is atomic
- $s \Vdash \phi \wedge \psi$ iff $\exists t : s, t \Vdash \phi$ and $s, t \Vdash \psi$
- $s \Vdash \phi \vee \psi$ iff $\exists t : s, t \Vdash \phi$ or $s, t \Vdash \psi$
- $s \Vdash \neg\phi$ iff $\forall s' \sqsupseteq s, \forall t : s', t \not\Vdash \phi$
- $s \Vdash \phi \rightarrow \psi$ iff $\forall s' \sqsupseteq s, t_1$: if $s', t_1 \Vdash \phi$ then $\exists t_2 \geq t_1 . s' t_2 \Vdash \psi$

For atomic, conjunctive and disjunctive sentences, it obvious how forcing relative to a state and a time should be defined:

- $s, t \Vdash p$ iff $V(s, t, p) = 1$ if p is atomic
- $s, t \Vdash \phi \wedge \psi$ iff $s, t \Vdash \phi$ and $s, t \Vdash \psi$
- $s, t \Vdash \phi \vee \psi$ iff $s, t \Vdash \phi$ or $s, t \Vdash \psi$

However, for the out-of-state connectives of negation and implication there is more room for manouevre. The two possibilities open for implication correspond to the specific and habitual conditionals $\rightarrow$ and $\Rightarrow$. defined in Chapter 3.

$\Rightarrow$ version:
$s, t \Vdash \phi \rightarrow \psi$ iff $\forall s' \sqsupseteq s, t_1 \geq t$: if $s', t_1 \Vdash \phi$ then $\exists t_2 \geq t_1 . s' t_2 \Vdash \psi$

$\rightarrow$ version:
$s, t \Vdash \phi \rightarrow \psi$ iff $\forall s' \sqsupseteq s, t_1 \geq t$: if $s', t_1 \Vdash \phi$ then $\exists s_2 \approx_{t_1} s' . s_2, t_1 \Vdash \psi$

Both are equivalent, and we will use the simpler, $\Rightarrow$ version.

For negation, there is no absolute necessity to restrict attention to times $t_1 \geq t$. But monotonicity of direct verification means that nothing extra is gained through considering times before t as well. Therefore, the following may as well be adopted:

$$s, t \Vdash \neg\phi \text{ iff } \forall s_1 \sqsupseteq s, t_1 \geq t : s_1, t_1 \nVdash \phi$$

Finally, noting that all intuitionistic formulas are monotonic, it is possible to replace the universal quantification over states in negation and implication by universal quantification over minimally extending states. This is because of the following theorem

Minimal Extension Theorem:
If $s' \sqsupseteq_{t'} s$ and $s', t' \Vdash \phi$, then $\exists s_0, t_0$ such that $t_0 \leq t'$, $s_0 \sqsupseteq_{t0}^{\phi, t_0} s'$ and $s' \sqsupseteq_{t'} s_0$

(Proof: simple variant on proof given in Appendix B). That is, all states extending s supporting ϕ will extend a minimal extension of s with respect to ϕ. Provided ϕ and ψ are monotonic, if ψ is supported in all minimal extensions of a state s with respect to ϕ, then ψ is also supported in *all* extensions of s supporting ϕ. We therefore arrive at the final definitions for $\rightarrow$ and $\neg$:

- $s, t \Vdash \phi \rightarrow \psi$ iff $\forall t_1 \geq t, s_1 \sqsupseteq_{t_1}^{\phi, t} s : \exists t_2 \geq t_1$ such that $s_1, t_2 \Vdash \psi$
- $s, t \Vdash \neg\phi$ iff $\forall t_1 \geq t, s_1 \sqsupseteq_{t_1}^{\phi, t} s : \exists t_2 \geq t_1$ such that $s_1, t_2 \Vdash \bot$

By introducing extra internal structure to information states (through distinguishing between verified and unverified assertions), the stock of intuitionistic connectives can be extended. In-state negation, $\sim$, is one additional connective that results, and has no counterpart in ordinary intuitionistic logic.

5.2.3 Proof Theory

The semantic definitions presented above are sound and complete with respect to the following sequent calculus system (Proof: Appendix C).

$$\wedge I \frac{\Gamma \vdash \phi;\ \Gamma \vdash \psi}{\Gamma \vdash \phi \wedge \psi} \qquad \wedge E \frac{\Gamma \vdash \phi \wedge \psi}{\Gamma \vdash \phi}$$

$$\vee I \frac{\Gamma \vdash \phi}{\Gamma \vdash \phi \vee \psi} \qquad \vee E \frac{\Gamma \vdash \phi \vee \psi;\ \Gamma, \phi \vdash \chi;\ \Gamma, \psi \vdash \chi}{\Gamma \vdash \chi}$$

$$\rightarrow I \frac{\mathit{Stable}(\Gamma), \phi \vdash \sim\sim \psi}{\Gamma \vdash \phi \rightarrow \psi} \qquad \rightarrow E \frac{\Gamma \vdash \phi;\ \Gamma \vdash \phi \rightarrow \psi}{\Gamma \vdash \sim\sim \psi}$$

$$\neg I \frac{\mathit{Stable}(\Gamma), \phi \vdash \sim\sim \bot}{\Gamma \vdash \neg\phi} \qquad \bot \frac{\Gamma \vdash \bot}{\Gamma \vdash \phi}$$

$$\sim I \frac{\Gamma, \phi \vdash \bot}{\Gamma \vdash \sim \phi} \qquad \sim Ax \frac{}{\Gamma \vdash \sim \phi \vee \sim\sim \phi}$$

$$\sim\rightarrow \frac{\Gamma \vdash \sim\sim \phi;\ \Gamma \vdash \phi \rightarrow \psi}{\Gamma \vdash \sim\sim \psi} \qquad \sim\sim \bot \frac{\Gamma \vdash \sim\sim \bot}{\Gamma \vdash \bot}$$

(The conjunction and disjunction rules are symmetric with respect to conjuncts/disjuncts in the usual way). The axiom stating that $\vdash \sim \phi \vee \sim\sim \phi$ is a direct result of the fact that the ordering over temporal snapshots is linear, giving rise to the characteristic axiom of Gabbay's logic **KC**.

5.3 Modality, Conditionality and Negation

In-state negation has a dual character, depending on whether one is looking at information states from the inside or the outside. From the inside, it behaves like ordinary intuitionistic negation over snapshots of the state. From the outside it behaves much more like classical negation: any sentence ϕ is either asserted by a state — $s \Vdash \phi$ — or is not asserted — $s \Vdash \sim \phi$.

In-state and out-of-state negation provides a convenient way of defining epistemic modalities. What is more, they provide *three* modalities, corresponding loosely to the natural language modals *must, may* and *will.*

(a) 'May': $\mathbf{may}(\phi) =_{df} \sim \neg\phi$

$s, t \Vdash \sim \neg\phi$

iff

$\forall t_1 \geq t : s, t_1 \nVdash \neg\phi$

iff

$\forall t_1 \geq t : \exists s_1, t_2$ s.t. $s_1 \sqsupseteq_{t_2}^{\phi, t_1} s$

(b*) 'Strong Must': $\mathbf{must}_s(\phi) =_{df} \neg\neg\phi$

$s, t \Vdash \neg\neg\phi$

iff

$\forall s_1, t_1$ if $s_1 \sqsupseteq_{t_1}^{\neg\phi, t} s$, then $s_1, t_1 \Vdash \bot$,

iff

not($\exists s_1, t_1$ s.t. $s_1 \sqsupseteq_{t_1}^{\neg\phi, t} s$),

iff

not[$\exists s_1, t_1$ s.t. $s_1 \sqsupseteq_{t_1} s$, $s_1, t_1 \Vdash \neg\phi$ and not($\exists s_2, t_2$ s.t. $t \leq t_2 < t_1$, $s \sqsubseteq_{t_2} s_2 \sqsubseteq_{t_2} s_1$ and $s_2, t_2 \Vdash \neg\phi$)],

iff

not[$\exists s_1, t_1$ s.t. $s_1 \sqsupseteq_{t_1} s$, not($\exists s_2, t_2$ s.t. $t \leq t_2 < t_1$, $s \sqsubseteq_{t_2} s_2 \sqsubseteq_{t_2} s_1$ and $s_2, t_2 \Vdash \neg\phi$), and not($\exists s_2, t_2$ s.t. $s_2 \sqsupseteq_{t_2}^{\phi, t_1} s_1$)],

iff

$\forall s_1, t_1$ if $s_1 \sqsupseteq_{t_1} s$ and not($\exists s_2, t_2$ s.t. $t \leq t_2 < t_1$, $s \sqsubseteq_{t_2} s_2 \sqsubseteq_{t_2} s_1$ and $s_2, t_2 \Vdash \neg\phi$), then $\exists s_2, t_2$ s.t. $s_2 \sqsupseteq_{t_2}^{\phi, t_1} s_1$

(c) 'Will': $\mathbf{will}(\phi) =_{df} \sim\sim \phi$

$s, t \Vdash \sim\sim \phi$

iff

$\forall t_1 \geq t : s, t_1 \nVdash \neg\phi$

iff

$\forall t_1 \geq t : \exists t_2 \geq t_1 : s, t_2 \Vdash \phi$

(d) 'Strong Will': $\mathbf{will}_s(\phi) =_{df} \neg \sim \phi$

$s,t \mid\!\vdash \neg \sim \phi$

iff

not($\exists s_1, t_1$ s.t. $s_1 \sqsupseteq_{t_1}^{\sim\phi,t} s$),

iff

not[$\exists s_1, t_1$ s.t. $s_1 \sqsupseteq_{t_1} s$, $s_1, t_1 \mid\!\vdash \sim \phi$ and not($\exists s_2, t_2$ s.t. $t \leq t_2 < t_1$, $s \sqsubseteq_{t_2} s_2 \sqsubseteq_{t_2} s_1$ and $s_2, t_2 \mid\!\vdash \sim \phi$)],

iff

$\forall s_1, t_1$ if $s_1 \sqsupseteq_{t_1} s$ and not($\exists s_2, t_2$ s.t. $t \leq t_2 < t_1$, $s \sqsubseteq_{t_2} s_2 \sqsubseteq_{t_2} s_1$ and $s_2, t_2 \mid\!\vdash \sim \phi$), then $s_1, t_1 \mid\!\not\vdash \sim \phi$,

iff

$\forall s_1, t_1$ if $s_1 \sqsupseteq_{t_1} s$ and not($\exists s_2, t_2$ s.t. $t \leq t_2 < t_1$, $s \sqsubseteq_{t_2} s_2 \sqsubseteq_{t_2} s_1$ and $s_2, t_2 \mid\!\vdash \sim \phi$), then $s_1, t_1 \mid\!\vdash \sim\sim \phi$

Since $s \sqsupseteq_t s$ and not($\exists s_2, t_2$ s.t. $t \leq t_2 < t_1$, $s \sqsubseteq_{t_2} s_2 \sqsubseteq_{t_2} s_1$), it follows that $s, t \mid\!\vdash \sim\sim \phi$

It turns out that the definition of necessity in terms of double out-of-state negation is too strong for our purposes, since it can be shown that $\mathbf{must}_s(\mathbf{may}(\phi))$ entails $\mathbf{must}_s(\phi)$ (see Chapter 3).

At first sight, this might appear not to be the case. $\mathbf{must}_s(\phi)$ says that any state extending the current one that is not preceded by a state in which ϕ is false will have a minimal extension supporting ϕ. The problem with a formula $\phi = \mathbf{may}(\psi)$ is that $\mathbf{may}(\psi)$ can start off true, but on the accumulation of more information become false. It would seem that the definition of **must** takes account of this by only demanding that minimal extensions to support $\mathbf{may}(\psi)$ are possible before $\mathbf{may}(\psi)$ becomes false.

That this is not true can be seen from the example illustrated in Figure 5.1. In state

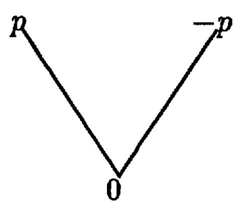

Figure 5.1: Counterexample to **must** as double negation

0, $\mathbf{may}(p)$ holds, and $\mathbf{must}_s(\mathbf{may}(p))$ ought to. But consider the state supporting $-p$. There is no state preceding it in which $\mathbf{may}(p)$ holds, and therefore, if $\mathbf{must}_s(\mathbf{may}(p))$ holds in state 0, the state $-p$ should have a minimal extension supporting p. It cannot. This means that $\mathbf{must}_s(\mathbf{may}(p))$ can only hold in state 0 if there are no states like $-p$ extending it. And this is only the case if $\mathbf{must}(p)$ holds in state 0.

This undesirable state of affairs can be rectified simply. We can define necessity as

(b) 'Must': $\mathbf{must}(\phi) =_{df} \phi \vee \neg\neg\phi$

When ϕ is monotonic, this is equivalent to the definition in terms of double negation, since ϕ entails $\neg\neg\phi$ if ϕ is monotonic. But if ϕ is non-monotonic, it allows for the possibility that although ϕ holds now, it may not always continue to hold.

There is a comparable difference between **will** and its stronger version $\mathbf{will}_s$. When ϕ is monotonic, $\mathbf{will}(\phi)$ and $\mathbf{will}_s(\phi)$ are equivalent, just as $\mathbf{must}(\phi)$ and $\mathbf{must}_s(\phi)$ are. But when ϕ is non-monotonic (e.g $\mathbf{may}(\psi)$), then $\mathbf{will}(\phi)$ is non-monotonic, whereas $\mathbf{will}_s(\phi)$ is monotonic.

The modal **may** is a denial of an assertion of denial: $\mathbf{may}(\phi)$ denies that it is asserted that ϕ is not the case.

5.3.1 Iterated Modalities and Conditionals

Given the three modalities defined above, the following equivalences and entailments hold:

$\mathbf{must}(\mathbf{must}(\phi))$	iff	$\mathbf{must}(\phi)$
$\mathbf{must}(\mathbf{may}(\phi))$	iff	$\mathbf{may}(\phi)$
$\mathbf{must}(\mathbf{will}(\phi))$	iff	$\mathbf{must}(\phi)$
$\mathbf{may}(\mathbf{must}(\phi))$	iff	$\mathbf{may}(\phi)$
$\mathbf{may}(\mathbf{may}(\phi))$	iff	$\mathbf{may}(\phi)$
$\mathbf{may}(\mathbf{will}(\phi))$	iff	$\mathbf{may}(\phi)$
$\mathbf{will}(\mathbf{must}(\phi))$	iff	$\mathbf{must}(\phi)$
$\mathbf{will}(\mathbf{may}(\phi))$	iff	$\mathbf{may}(\phi)$
$\mathbf{will}(\mathbf{will}(\phi))$	iff	$\mathbf{will}(\phi)$
$\mathbf{must}_s(\mathbf{may}(\phi))$	iff	$\mathbf{must}_s(\phi)$
$\mathbf{will}_s(\mathbf{may}(\phi))$	iff	$\mathbf{must}_s(\phi)$
If $\mathbf{will}(\phi)$	then	$\mathbf{must}(\phi)$
If $\mathbf{must}(\phi)$	then	$\mathbf{may}(\phi)$
If ϕ	then	$\mathbf{will}(\phi)$

Iterations of modals with negations have a few surprises in store. First, to discharge the disjunction in **must**, note that:

$\mathbf{must}(\neg\phi)$ iff $\neg\phi \vee \neg\neg\neg\phi$ iff $\neg\phi$
(since $\neg\neg\neg\phi$ entails $\neg\phi$)

$\mathbf{must}(\sim\phi)$ iff $\sim\phi \vee \neg\neg\sim\phi$ iff $\sim\phi$
(since $\neg\neg\sim\phi$ entails $\sim\phi$)

As a result, we get:

$$
\begin{array}{lcl}
\neg\mathbf{must}\neg\phi & = \neg\neg\phi & = \mathbf{must}_s(\phi) \\
\neg\mathbf{must}\sim\phi & = \neg\sim\phi & = \mathbf{will}_s(\phi) \\
\sim\mathbf{must}\neg\phi & = \sim\neg\phi & = \mathbf{may}(\phi) \\
\sim\mathbf{must}\sim\phi & = \sim\sim\phi & = \mathbf{will}(\phi) \\[4pt]
\neg\mathbf{may}\neg\phi & = \neg\sim\neg\neg\phi & = \mathbf{must}_s(\phi) \\
\neg\mathbf{may}\sim\phi & = \neg\sim\neg\sim\phi & = \mathbf{will}_s(\phi) \\
\sim\mathbf{may}\neg\phi & = \sim\sim\neg\neg\phi & = \mathbf{must}_s(\phi) \\
\sim\mathbf{may}\sim\phi & = \sim\sim\neg\sim\phi & = \mathbf{will}_s(\phi) \\[4pt]
\neg\mathbf{will}\neg\phi & = \neg\sim\sim\neg\phi & = \mathbf{must}_s(\phi) \\
\neg\mathbf{will}\sim\phi & = \neg\sim\sim\sim\phi & = \mathbf{will}_s(\phi) \\
\sim\mathbf{will}\neg\phi & = \sim\sim\sim\neg\phi & = \mathbf{may}(\phi) \\
\sim\mathbf{will}\sim\phi & = \sim\sim\sim\sim\phi & = \mathbf{will}(\phi) \\[4pt]
\neg\mathbf{must}_s\neg\phi & = \neg\neg\neg\neg\phi & = \mathbf{must}_s(\phi) \\
\neg\mathbf{must}_s\sim\phi & = \neg\neg\neg\sim\phi & = \mathbf{will}_s(\phi) \\
\sim\mathbf{must}_s\neg\phi & = \sim\neg\neg\neg\phi & = \mathbf{may}(\phi) \\[4pt]
\neg\mathbf{will}_s\neg\phi & = \neg\neg\sim\neg\phi & = \mathbf{must}_s(\phi) \\
\neg\mathbf{will}_s\sim\phi & = \neg\neg\sim\sim\phi & = \mathbf{must}_s(\phi) \\
\sim\mathbf{will}_s\neg\phi & = \sim\neg\sim\neg\phi & = \mathbf{may}(\phi) \\
\sim\mathbf{will}_s\sim\phi & = \sim\neg\sim\sim\phi & = \mathbf{may}(\phi)
\end{array}
$$

This is a rather bewildering array of dualities. Provided that the two different forms of negation are read in the appropriate way, they are all plausible, though it can sometimes be a bit of a mouthful to paraphrase what is meant. For example, $\neg\mathbf{must}\neg\phi$ says roughly that it will never be the case that it must never be the case that ϕ, which means that eventually ϕ has to turn out to be the case. Similarly, $\neg\mathbf{must}\sim\phi$ says roughly that will never be the case that ϕ must not be asserted, which means that ϕ has to be asserted, since either ϕ is asserted or it isn't.

The dualities that come closest to those familiar from modal logic are of the form $\sim$ *modal*$\neg\phi$. Thus $\sim\mathbf{must}\neg\phi$ says that it is not asserted that it must never be the case that ϕ, which means that it may be the case that ϕ. The dual $\sim\mathbf{may}\neg\phi$ is one of the ones that gives rise to a lengthy paraphrase. Recall that $\mathbf{may}(\phi)$ denies that it is asserted that ϕ is never the case. Thus $\mathbf{may}(\neg\phi)$ denies that it is asserted that $\neg\phi$ is never the case. And $\sim\mathbf{may}(\neg\phi)$ denies the denial: that is, it is asserted that $\neg\phi$ is never the case. Which means that ϕ must eventually turn out to be the case.

Where conditionals are concerned, the following facts are worthy of note:

If $\sim(\phi\rightarrow\psi)$ then $\mathbf{may}(\phi\wedge\neg\psi)$

If $\neg(\phi \rightarrow \psi)$ then $\phi \rightarrow \neg\psi$

From $\phi \rightarrow$ **may**(ψ) it does not follow that $\phi \rightarrow \psi$

That is, there are two ways of negating a conditional — denying that the conditional is asserted ($\sim$), and asserting the denial of the conditional ($\neg$). These give rise to rather different results. Also, the semi-stability of a conditional like $\phi \rightarrow$ **may**(ψ) ensures that it does not entail $\phi \rightarrow \psi$.

Comparison to Data Semantics

The logic of verified and unverified assertions presented here does not fall prey to two of the problems confronting basic Data Semantics (Chapter 3), namely that **must**(**may**(ϕ)) $\equiv$ **must**(ϕ) and $\phi \rightarrow$ (**may**(ϕ)) $\equiv \phi \rightarrow \psi$. Both Veltman (1985) and Landman (1985) propose solutions to these problems, and it is worth comparing them to the solution here.

Landman's solution runs as follows. For any information state s, there will be different branches of states growing out of s. Each branch represents one way in which the information in s might grow. One branch might first add the information that ϕ holds, and then later add the information that ψ holds. Another branch might reverse this order, so that first ψ is added and then ϕ. In this case, both branches lead to the same information state, but by different routes. Other branches may lead to entirely different states, e.g. one that adds $\neg\phi$ and then $\neg\psi$.

Under Landman's revised treatment, s supports **must**(ϕ) if either (a) ϕ already holds in s, or (b) ϕ is neither asserted nor denied in s, but in all branches growing out of s the first thing that happens to ϕ is that it becomes asserted. Should ϕ turn from being asserted to being denied after that point in the branch, as it might if $\phi =$ **may**(p), this does not affect the truth of **must**(ϕ).

Similarly, s supports **may**(ϕ) if either (a) ϕ already holds in s, or (b) ϕ does not hold in s, but there is at least one branch growing out of s where the first thing that happens to ϕ is that it becomes asserted. Once more, what happens to ϕ after this point does not matter.

With the conditional, s supports $\phi \rightarrow \psi$ if in all branches out of s as soon as ϕ holds, then in all branches growing from that point the first thing that happens to ψ is that it either already holds or comes to hold.

While this blocks the two troublesome equivalences, it is not entirely satisfactory. Consider the set of information states shown in Figure 5.2. The state 0 ought to support $p \rightarrow$ **may**(q). But suppose you follow the branch passing through the state $-q$ first. Once p is added to this branch, q neither holds nor will come to hold. Thus the state 0 does not support the required conditional.

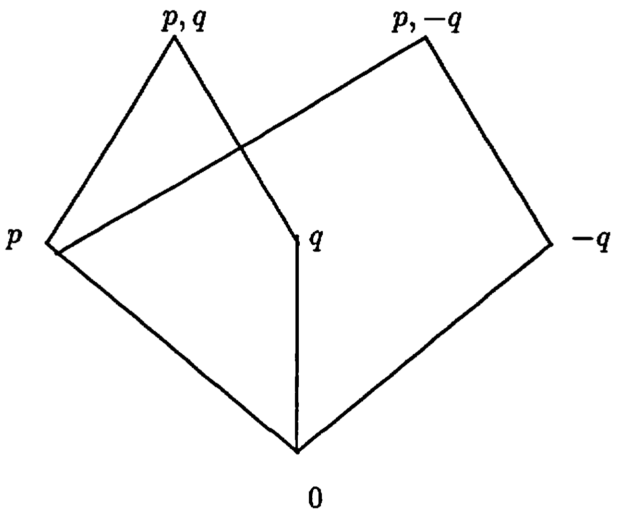

Figure 5.2: Landman's implication into possibility

Veltman's solution does not suffer from this drawback. He says that $\mathbf{must}(\phi)$ holds in a state s if for every state extending s it that denies ϕ, there is one between it and s that supports ϕ. A state s supports $\mathbf{may}(\phi)$ if there is some state extending s and supporting ϕ, with no intervening state denying ϕ. A state s supports $\phi \rightarrow \psi$ if for every state extending s and supporting ϕ while denying ψ, there is an intervening one supporting both ϕ and ψ.

However, as Veltman notes, this scheme of things licences the inference from $\mathbf{may}(\phi \rightarrow \psi)$ to $\phi \rightarrow \psi$. This inference does not affect the logic presented here, where $\mathbf{may}(\phi \rightarrow \psi)$ only allows one to conclude that $\mathbf{may}(\phi \wedge \mathbf{must}_s(\psi))$.

Veltman (1985:210) perhaps pinpoints the difference between his revised version of Data Semantics and the kind of approach pursued here, when he says

> [The] truth definitions assume that one's knowledge of the *changes* which one's partial knowledge *could yet undergo* is complete. By freely quantifying over *all* possible extensions of partial information they assume that one, in evaluating conditionals, is in a position to take all of these possibilities into account .... This is of course not very true to real life.

In the current treatment, modals and conditionals are defined in terms of minimal information extensions. This avoids free quantification over all possible information extensions, while still sometimes constraining what may happen in any possible information extension. This seems at least slightly truer to real life.

'Not' and Negation

In introducing two forms of negation, an obvious question arises: which corresponds to the word *not*? It would take us too far afield to explore this question in sufficient depth. But it is plausible to suppose the *not* can sometimes correspond to a denial of assertion, $\sim$, and sometimes to an assertion of denial, $\neg$.

Lyons (1977) makes a roughly similar point in distinguishing between negation of the *tropic* (roughly, propositional status) and *neustic* (speaker commitment). Denial of assertion, $\sim$ corresponds to Lyons' neustic negation — I do not say that it is the case. Assertion of denial corresponds to tropic negation — I say it is not the case.

In this connection, it is interesting to note that the normal modal duals correspond to a denial of assertion followed by an assertion of denial. Lyons assumes that the neustic always has scope over the tropic. This exactly reflects the order of $\sim$ and $\neg$ in this case.

5.3.2 Paradoxes of Implication

There are a number of traditional problems in the logic of conditionals that I have not so far addressed. The implication defined above validates the principles of strengthening the antecedent, transitivity and contraposition (in one direction). These have often been felt to be intuitively invalid principles, witness:

- Strengthening of the antecedent:
  If you put sugar in the coffee, it will taste nice.
  Therefore if you put sugar and diesel oil in the coffee it will taste nice.
  $\phi \rightarrow \chi \vdash (\phi \wedge \psi) \rightarrow \chi$
- Transitivity:
  If there is an election, the president will be voted back into office.
  If the president dies, there will be an election.
  Therefore, if the president dies, he will be voted back into office.
  $\phi \rightarrow \psi, \psi \rightarrow \chi \vdash \phi \rightarrow \chi$
- Contraposition:
  If it is after 5pm, it is not much after 5pm.
  Therefore, if it is much after 5pm, it is not after 5pm.
  $\phi \rightarrow \psi \vdash \neg\psi \rightarrow \neg\phi$

It is commonly argued that these inferences fail because the premise conditional and the conclusion conditional make different presuppositions. I think this is correct, and moreover that a very strong argument can be given in favour of this position.

Consider strengthening the antecedent, and let us construe the conditional as material implication for the time being. Truth tables reveal that if $\phi \supset \chi$ and $(\phi \wedge \psi) \supset \chi$, it follows that $\neg(\phi \wedge \psi)$. That is, if someone genuinely asserts both that if you put sugar in the coffee it will taste nice and that if you put sugar and oil in then it won't, then that person is also committed to the claim that you will not add both sugar and oil. Exactly the same applies to transitivity. If someone assents to both the conditionals about the election, they are tacitly assenting to the fact that the president won't die. And similarly with contraposition: someone who assents to the first conditional assumes that it is indeed not much after 5pm, if indeed it is after 5pm at all. This brings out in a very direct form the presuppositions lying behind the conditionals. The same arguments apply when the conditional is construed as $\rightarrow$, though truth tables cannot be used to bring this out.

In other words, $\rightarrow$ validates strengthening of the antecedent, transitivity and contraposition (in one direction) and is perfectly correct to do so.

5.4 Non-Epistemic Modalities

Up until now, I have been concentrating exclusively on epistemic modalities. These can be seen as special cases of the more general modalities defined below

- $s, t \Vdash$ **must**(ϕ) iff
  $s, t \Vdash \phi$ and sR_ts, or
  $not(\exists s', t'$ such that $s' \sqsupseteq_{t'}^{\neg\phi,t} s$ and $sR_{t'}s')$
- $s, t \Vdash$ **may**(ϕ) iff
  $\exists s', t'$ such that $s' \sqsupseteq_{t'}^{\phi,t} s$ and $sR_{t'}s'$
- $s, t \Vdash$ **will**(ϕ) iff
  $\exists t'$ such that $s \sqsupseteq_{t'}^{\phi,t} s$ and $sR_{t'}s$.

Here, R_t is a time-dependent binary accessibility relation on information states. For epistemic modals, $R_t = \sqsubseteq_t$, and the epistemic modals can be reduced to sequences of negations.

Non-epistemic modalities are defined in terms of minimal information extensions that are also modally accessible. For example, the deontic **may** of permission says that **may**(ϕ) holds in state s if s can be minimally extended to form a deontically acceptable state that supports ϕ. Deontic **must**(ϕ) says that there is no minimal extension of s making ϕ false that is also deontically acceptable. This is not to say that s has no minimal extensions falsifying ϕ, merely that none of these are deontically acceptable.

It is beyond the scope of this thesis to explore what modal properties result from different conditions imposed on the accessibility relation. I shall confine myself to a few general comments.

First, non-epistemic modals are constrained by what is epistemically possible. Permission to do something known to be epistemically impossible, e.g. travelling to Mars tomorrow, does not make sense. Obligations, abilities, volitions and so forth are likewise constrained to be realistic.

Some obligations and permissions are non-realistic, e.g. *I ought to go to the lecture, but I am not going to.* However, these are realistic relative to some counterfactual information state. The non-hypothetical modal *must* is anomalous if used to express an unrealistic obligation: *?I must go the lecture, but I am not going to.*

Second, while epistemic modals are sometimes non-monotonic, non-epistemic modals are liable to be considerably more non-monotonic. The following example of this is adapted from Thomason (1981a).

Suppose that I must go and get a tin of dog-food by four o'clock. Suppose that shortly before four, in a moment of weakness, I spend all my money on something else. Once this happens, I must not get a tin of dog-food, since to get it I would have to steal it, having no money left. This example illustrates that obligations can not only come, but also go, with varying circumstances. The original information state in which the obligation arose did not take account of the possibility of my spending all the money on something else. Once the information state is extended to include this fact, the obligation no longer holds.

Formally, this means from the fact that sR_ts_2 and $s \sqsubseteq_t s_1$ we are not allowed to conclude that $s_1R_ts_2$.

Third, it was noted in Chapter 2 that non-epistemic modals cannot be used to refer to past eventualities, even when the perfect *have* is employed. This restriction could be enforced by ensuring that for any state s' such that sR_ts', s' does not include any extra assertions that are verifiable before t.

Fourth, the treatment of non-epistemic modalities here appears to be compatible with that proposed by Kratzer (1977, 1981b). She argues that modals are dependent on a contextually given modal source. For epistemic modals, the source consists of what is known, for deontic modals it consists of what is required or permitted, and so forth. These sources consist of 'atomic' facts, plus conditional rules or regularities. Just as the informational ordering expresses epistemic regularities relative to a given state supporting atomic facts, the accessibility relation can be seen as expressing deontic regularities relative to a state. I conjecture that Kratzer's modal sources correspond to an information state plus accessibility relation. As with Kratzer's account, different modal auxiliaries are consistent with different types of modal source / accessibility relation. Thus, *may* can be used with an epistemic or deontic accessibility relation, while *can* is used with a deontic accessibility relation or one pertaining to abilities.

Clearly, much work still needs to be done in exploring this treatment of non-epistemic modality. But the basic idea holds promise.

Modal Intuitionistic Logic

There are other ways of treating modality within an intuitionistic framework, e.g. Bozic and Dosen 1984, Plotkin and Stirling 1986, Wijesekera 1990. However, these differ considerably from the treatment suggested here, and have entirely different aims.

All the work on intuitionistic modal logic of which I am aware strives to preserve the monotonicity property for all formulas, modal and non-modal. This is in sharp contrast to the logic of verified and unverified assertions, where modal formulas can sometimes be non-monotonic.

Non-monotonicity is inevitable for epistemic **may**: as information grows, possibilities diminish. This means that whatever modal notion is captured by the $\Diamond$ operator in modal intuitionistic logic, it is not epistemic possibility. It is no surprise then that most modal intuitionistic logics have been developed for the purposes of reasoning about computer programs. An example from Wijesekera (1990) will illustrate the point.

Let $M = \langle S, \sqsubseteq, R, V\rangle$ be a modal intuitionistic model where S is the set of information states, $\sqsubseteq$ the partial information order, R is the accessibility relation and is in any binary relation on S, and V is a valuation function. The semantic definitions for the modal operators $\Box$ and $\Diamond$ are

- $s \Vdash \Box\phi$ iff for all $s' \sqsupseteq s$ and for all s'' such that $s'Rs''$, $s'' \Vdash \phi$
- $s \Vdash \Diamond\phi$ iff for all $s' \sqsupseteq s$ there is a s'' such that $s'Rs''$, $s'' \Vdash \phi$

It can be shown that if $s \Vdash \Box\phi$, then all $s' \sqsupseteq s$ are such that $s' \Vdash \Box\phi$, and similarly for $\Diamond$.

This logic is useful for reasoning about the behaviour of programs, where information states represent partial information about the machines internal state. The ordering $\sqsubseteq$ corresponds to the way that this partial information may be extended, ultimately to give a complete description of a single machine state. The accessibility relation R describes the effect that program executions have on machine states. This means that $s \sqsupseteq s'$ concerns partial descriptions of a total machine state for a given time, where the underlying machine state remains constant. But sRs' concerns how one machine state can lead to another through the execution of a program.

Under this construal, $s \Vdash \Diamond\phi$ says that it is always going to be possible for ϕ to hold after a given program (R) has executed from any state compatible with the partial description furnished by s. Similarly $s \Vdash \Box\phi$ guarantees that executing R in any state compatible with s will lead to a state in which ϕ holds. When describing program behaviour, it is convenient to be able to focus on just one part of the machine state (the computer may be running other processes at the same time which affect other memory locations, but which are of no relevance). The modals provide a way of saying that the program will have a certain effect, independently of how other processes have configured other parts of the machine state. For this kind of information, one clearly needs the modals to behave monotonically.

Plotkin and Stirling (1986) suggest a somewhat different approach. The relevant modal definitions are:

- $s \Vdash \Diamond\phi$ iff $\exists s'.\ sRs'$ and $s' \Vdash \phi$
- $s \Vdash \Box\phi$ iff $\forall s'\ s' \sqsupseteq s$, $\forall s''$ if $s'Rs''$ then $s'' \Vdash \phi$

Two conditions are imposed on the informational ordering and the accessibility relation:

- If $s \sqsubseteq s'$ and sRr, then $\exists r'.\ s'Rr'$ and $r \sqsubseteq r'$
- If $r \sqsubseteq r'$ and sRr, then $\exists s'.\ s'Rr'$ and $s \sqsubseteq s'$

The first condition ensures the monotonicity of modal formulas, and also the implication $\Diamond\phi \rightarrow \neg\Box\neg\phi$. The second ensures that $\neg\Diamond\phi \rightarrow \Box\neg\phi$. Subject to these constraints, different accessibility relations will give rise to different logics, just as with Kripke semantics for classical modal logic.

Neither of these ways of defining the modal operators is of much help to us (nor is that proposed by Bozic and Dosen). Wijesekera's automatically imposes monotonicity, which is undesirable for epistemic modality. Plotkin and Stirling impose monotonicity via a condition on modal frames. But even if this were dropped, taking the accessibility relation R to be the same as the informational ordering $\sqsubseteq$ in the case of epistemic modalities would cause necessity to collapse.

5.5 Conclusions

The primary purpose of this chapter has been to show that the exclusively temporal focus on modals and conditionals in the previous chapters does not lead to something that is logically absurd. It remains to be seen how useful the logic of verified and unverified assertion is as a logic of (non-temporal) conditionals and modals.

One remaining question concerns whether there is a unique or overridingly obvious way of re-incorporating tense into this logic. I suspect that there is not, and that this is as it should be. It is well known that even quite closely related languages can differ widely in the way that their tense systems operate. Supposing that the logic in this chapter is the one that underlies all these languages, one would not wish to be forced into a just one legitimate way of incorporating tense. Some of the temporally based definitions of tenses, modals and conditionals may appear somewhat arbitrary from the point of view of the logic presented here. For example, the 'specific' and 'habitual' conditionals connectives of Chapter 3 both correspond to the same underlying implication as defined here. Why, then, have two different versions? The answer is: that's just the way that English happens to handle conditionals[3].

[3]Whenever someone says something like this, it is a sure sign that they are missing an important generalisation. In this connection, I have to admit that I do not feel that I have got right to the bottom of the difference between habitual and specific conditionals.

Chapter 6

Hypothetical Modals and Subjunctive Conditionals

We now turn to the hypothetical modal auxiliaries *would, could, should* and *might*, and to conditionals containing them. This encompasses what are often called subjunctive or counterfactual conditionals.

The hypothetical modals are sometimes described as past tense modals. By analogy with the present tense modals, one might hope to treat them in terms of a past tense operator applied to the modal operators previously defined. Unfortunately, this treatment is not viable (while *would* and *could* can sometimes be used to describe past futures and past abilities, *might* almost always has present modal time reference, and *should* invariably so).

Perkins (1983) has suggested that *would, could, should* and *might* are implicitly conditional. Following him in this, sentences like *John might go* are analysed as having (roughly) the following logical form

(1) John might go.
$\phi \mathrel{\Box\!\!\rightarrow} \mathbf{pres}(\mathbf{may}(go(j)))$

where ϕ gives the implicit antecedent of the implicit conditional, and $\mathrel{\Box\!\!\rightarrow}$ is the implicit (subjunctive) conditional connective. The implicit antecedent is filled in by context, so that (1) could be equivalent to *John might go if you asked him* or *John might go if you gave him enough money*, depending on context. It is the implicit conditionality of these modals that lends them their tentative air.

In subjunctive conditionals, the modal's implicit antecedent is filled in by the conditional's explicit antecedent, e.g

(2) John might go if you asked him.
$\mathbf{pres}(ask) \mathrel{\Box\!\!\rightarrow} \mathbf{pres}(\mathbf{may}(go))$

The apparently past tense on the conditional antecedent is in fact a subjunctive or conditional tense (as is the tense marking on the modal), and semantically corresponds to an ordinary present tense within the scope of the subjunctive conditional connective.

The subjunctive connective $\Box\!\!\rightarrow$ is a 'tensed' version of the indicative connective $\rightarrow$. The 'tense' on the subjunctive connective causes an underlying indicative conditional, e.g. $\mathbf{pres}(ask) \rightarrow \mathbf{pres}(\mathbf{may}(go))$ in the case of (2), to be evaluated relative to different information states and assertion times, picked out by a state selection function.

The chapter is organised as follows. Section 6.1 describes the temporal properties of conditionals with hypothetical modal consequents. Section 6.2 argues that the apparently past surface tenses in subjunctive conditionals reflect subjunctive (or conditional) tenses. Section 6.3 discusses the logico-syntactic structure of subjunctive conditionals, and gives semantic definitions for a past and a present subjunctive conditional connective. Section 6.4 makes some proposals about the behaviour of the information state selection function employed by the semantics of these two connectives. This is intended more to show that such a thing can be done than to give a detailed analysis.

6.1 Hypothetical Temporal Reference

Conditionals with hypothetical modal consequents divide into three classes

- Indicative (Non Back-Shifted): The tense form of the antecedent corresponds to its underlying semantic tense. Thus present tense antecedents refer to present or future times, and past tense antecedents refer to past times, e.g.

  (3) If John was at the meeting yesterday, he would be able to tell us what happened.

  (4) If John is here today, he could help us out.

- Semi-Indicative (Semi Back-Shifted): The antecedent has a past tense form and refers to a past time. However, the conditional behaves as though it were a conditional with a present tense antecedent and present tense modal consequent that was uttered at some point in the past, e.g.

  (5) If Dick Whittington turned again, he would become Lord Mayor of London.

- Subjunctive (Back-Shifted): The temporal reference of the antecedent is shifted forward from what might be expected from its surface form. The past tense is used to refer to present or future times, and the past perfect to past times, e.g.

  (6) If it rained tomorrow, we would call the match off. (Future reference)

  (7) If I knew the answer, I would tell you. (Present reference)

(8) If it had rained yesterday, the match would have been called off. (Past reference)

Antecedents with present tense surface forms do not occur in subjunctive conditionals.

6.1.1 Indicative Conditionals

Present tense antecedent

When the antecedent is in the present tense, the modal consequent is nearly always to be understood as expressing a tentative version of the corresponding present tense modal.

(9) If we arrive early, we could/would/might/should get good seats.
(If we arrive early, we can/will/may/ought to/shall get good seats.)

Under this interpretation, the conditionals display exactly the same temporal properties as the corresponding less tentative conditionals.

Very occasionally, though, *could* or *would* can be used to express past (non-epistemic) modalities when combined with with a present antecedent.

(10) If Thor Heyerdahl reaches America, then (maybe) the Ancient Egyptians could build papyrus boats capable of crossing the Atlantic.

Without the bracketed epistemic adverbial, this sentence is of marginal acceptability. What is dependent on the verification of the antecedent is not the boat building ability of the Egyptians, but our knowledge of it. The temporal reference of the consequent modal does not depend on the antecedent.

Past tense antecedent

When the antecedent is past tensed, tentative and past/epistemic conditionals are again possible:

(11) If you really did see the news last night, you would/should/might know what we are talking about.

(12) If Thor Heyerdahl reached America, then (maybe) the Ancient Egyptians could build papyrus boats capable of crossing the Atlantic.

6.1.2 Semi-Indicative Conditionals

Two forms of semi-indicative conditional may be distinguished: past narrative and past regularity. In both cases, the antecedent must have a past tense form.

Narrative In a past tense narrative,

(13) If Dick Whittington turned again, he would become Lord Mayor of London.

expresses the same thing as would

(14) If Dick Whittington turns again, he will become Lord Mayor of London.

uttered by a character in the narrative at that point. Past narrative conditionals behave exactly like conditionals with present antecedents and present modal consequents transposed to a past time of utterance. Thus the consequent eventuality tends to be simultaneous with or follow the antecedent eventuality, subject to the normal exceptions.

Regularity Conditionals about past regularities also behave as though they were present tense conditionals transposed to some past time of utterance.

(15) If the Emperor lowered his thumb, the Christian would be thrown to the lions.

In (15) the consequent event depends on the occurrence of the antecedent event. Thus, an antecedent precedes consequent reading arises. Because of this, a conditional like

(16) ?If the Christian was thrown to the lions, the Emperor would lower his thumb.

is pragmatically unacceptable.

6.1.3 Subjunctive Conditionals

The most common type of conditional with a past tense antecedent and hypothetical modal is subjunctive (or back-shifted). In back-shifted conditionals, simple past tense antecedents refer to non-past eventualities, and past perfect antecedents are needed to refer to past eventualities, e.g.

(6) If it rained tomorrow, we would call the match off. (Future reference)

(7) If I knew the answer, I would tell you. (Present reference)

(8) If it had rained yesterday, the match would have been called off. (Past reference)

The back-shift exhibited here is akin to a sequence of tense rule, where semantically underlying present tenses appear in surface form as past tenses and underlying past tenses appear as past perfects. However, a sequence of tense rule is not quite the right way to describe what is going on. For in some cases past perfect antecedents can have futurate reference, as in Dudman's (1983) example

(17) If the auditors had come tomorrow, they would have found the accounts in order.

This is not the behaviour one would expect of an underlying semantic past tense.

The explanation for a sentence like (17) is that although the antecedent event is a future one, it is incompatible with some event in the past (e.g. that the auditors came today and found the accounts in a mess, and they only visit once a year). What determines the choice of a past perfect tense form is potential incompatibility with past facts, and what determines the choice of a simple past form is potential incompatibility with present or future facts.

The tense form of the antecedent and consequent determines whether the conditional is evaluated relative to a 'past' or 'present' information state. Subject to this back-shifting over information states, the antecedent and consequent then behave as though they were semantically present tensed. In this respect, subjunctive conditionals are similar to semi-indicative conditionals.

It is useful to look at subjunctive conditionals under four sub-headings, which depend on whether the antecedent and consequent contain perfect auxiliaries.

1. Non-Perfective When neither antecedent nor consequent are perfective, the temporal properties of the back-shifted conditionals are exactly the same as those for the corresponding present tense conditionals. Thus

(18) a. If the bimetallic strip bent, the temperature would rise.

b. If the bimetallic strip bends, the temperature will rise.

are more or less synonymous, and express the same temporal relations.

2. Perfective Consequent Only When the antecedent is non-perfective but the consequent is perfective, the consequent modal has present or future time reference, and behaves like a future perfect. Thus the following two are more or less synonymous

(19) a. If the bimetallic strip bent, the temperature would have risen.

b. If the bimetallic strip bends, the temperature will have risen.

(The first of these has a preferred interpretation as a non back-shifted conditional, but this should be ignored). As with ordinary present tense modal conditionals, the consequent modal reference may either depend on the verification of the antecedent or on the time of utterance.

3. Perfective Antecedent Only When the antecedent is perfective but the consequent is not, the modal has present time reference. A conditional like

(20) ?If Hitler had not attacked the Soviet Union, he would win the war.

is pragmatically unacceptable since there is now no Hitler nor a war for him to win (though it would have been acceptable if uttered in 1944). A conditional like

(21) ?If the auditors had come tomorrow, they would find the books in order

also appears odd, since the present time reference of the consequent modal precedes that of the hypothetical antecedent eventuality.

4. Perfective Antecedent and Consequent When both antecedent and consequent are perfective, there are two possibilities for the temporal reference of the modal. The modal may either have present time reference, and behave like a future/modal perfect. Or it may depend on the time at which the antecedent event occurs, and so normally have past time reference. In the second case, the modal may behave either like a past modal perfect, or just a past modal. An example of the latter is

(22) If he had been found guilty, we could have imprisoned him.

which refers to a past ability to stop someone, and not to a present ability directed towards the past, or even a past ability directed to the further past. The *could have* in (22) marks a past, semantically non-perfective modal: the *have* acts as a marker of modal pastness, and not of perfectivity within the scope of the modal.

If the consequent modal is epistemic, it is hard to judge whether the modal has present (perfective) reference, past (non-perfective) reference, or past (perfective) reference.

(23) If the strip had bent, the temperature would have risen.

Sentence (23) can either mean that if the strip had bent, you would now be able to conclude that that the temperature had risen. Or it can mean that at the time of bending you could have concluded that the temperature would rise. Or it can mean that at the time of bending you could have concluded that the temperature had already risen. This ambiguity does not arise with non-epistemic modals (as (22) shows), since these are largely incompatible with perfectivity.

6.2 Back-Shift and Conditional Tenses

What distinguishes subjunctive conditionals from indicative conditionals is the fact that apparently past tense antecedents receive back-shifted interpretations. This section argues for two points. First, that back-shift is a more reliable criterion for distinguishing between subjunctive and indicative conditionals than the presumed counterfactuality of subjunctive antecedents. Second, that back-shifted tenses are subjunctive tenses: they are not deictically shifted past tenses, and they are not past and present tenses that have undergone some kind of sequencing of tense.

6.2.1 Counterfactuality

The distinction between indicative and subjunctive conditionals is usually drawn in completely non-temporal terms. Subjunctive conditionals are those where the antecedent is known or presumed to be false. Indicative conditionals have no presumption of falsity attached to the antecedent. The non-standard tense markings on subjunctive conditionals serve as a way of signalling the presumed counterfactuality, but do not have any significant temporal meaning[1].

It does not take much to point out the problems with this view, as Stalnaker (1975), Dudman (1991) and others, have already done.

First, some indicative conditionals have antecedents that are known to be false[2], e.g.

(24) If the Sun orbits the Earth, modern astronomy is fundamentally misconceived. (But of course, the Earth orbits the Sun).

[1]Kutschera (1974) argues that subjunctive conditionals should be subdivided into counterfactual and potential conditionals. In counterfactuals, the antecedent is presumed to be false. In potential conditionals, the antecedent is assumed to be probably false. In indicative conditionals, the truth value of the antecedent is assumed to be either completely unknown or true. He claims that Latin makes a syntactic distinction between counterfactual and potential conditionals.

In English, a similar distinction is perhaps marked by use of either a past perfect antecedent (counterfactual) or a simple past antecedent (potential). However, this is only because past antecedents refer to future events, and counterfactuality relative to future events is a good deal more open to doubt than counterfactuality relative to past events.

[2]Sentence (24) is also interesting in that it casts doubt on approaches to indicative conditionals that treat them as material implication plus 'assertability' conditions (e.g. Adams 1975, Grice 1976, Jackson 1979). A material implication with a false antecedent is trivially true (paradoxes of material implication). To prevent this paradox from carrying over to indicative conditionals, it is claimed that indicative conditionals with false antecedents are also trivially true, but that the falsity of the antecedent renders the conditional pragmatically unassertable: what is the point of asserting a conditional where it is already known to be inapplicable through the antecedent failing to be true? On this approach, a sentence like (24) should be pragmatically anomalous, yet it is not.

To be fair to Jackson, he would present the assertability conditional in terms of the probability of the implication given the probability of the antecedent being true. To save dividing by a zero probability, we would have to regard the antecedent as having a very low probability rather than simply false (zero probability). But this in turn raises the question of what happens when we make a declarative utterance: are we saying the expressed proposition is true, or only highly probable? If the latter, how can one read off the degree of probability from the surface form of the utterance?

Second, subjunctive conditionals can sometimes have antecedents known to be true, e.g.

(25) If the burglars had come in through the window, they would have picked up shards of glass on their clothing. (And since we know that they did get in that way, the forensic lab should check this coat for glass).

(26) If the butler had done it, we would have found exactly the clues we in fact found.

In (26) (due to Anderson, and reported in Stalnaker 1975), the truth of the consequent is used as evidence for the truth of the antecedent, which therefore cannot be viewed as being known or even presumed to be false.

Stalnaker and Presupposed Counterfactuality

Stalnaker (1975) suggests one way in which a presuppositional mechanism might apply to both subjunctive and indicative conditionals, while preserving the element of counterfactuality distinguishing subjunctive from indicative conditionals. He assumes that utterances are evaluated relative to some shared background of information between speaker and hearer. This background need not include everything that the speaker and hearer both know — it may be that the information is confined to what is germane to the topic of discourse. In Stalnaker's account, the background context is represented as the set of possible worlds compatible with the shared information.

Stalnaker's account of conditionals (Stalnaker 1968) makes use of a similarity ordering over possible worlds. To evaluate a conditional *If A then C* in a world w, one looks at the most similar world to w in which A is true, and sees whether C is true there — iff it is, the conditional is true in w[3]. Stalnaker proposes that worlds within the background context set (which if the shared information is correct includes w itself) are always regarded as more similar to w then worlds outside this set.

Indicative conditionals are ones where only similar worlds within the context set may be taken into consideration. The form of subjunctive conditionals acts as an explicit signal that worlds outside the context set can also be considered.

This allows for the possibility of non-counterfactual subjunctive conditionals, such as (26). Here, one can look at a similar world *within* the current context set. Sentences like (26) are the exception, however. Although subjunctive conditionals permit one to use similar worlds within the context sets as well as those outside it, a scalar implicature dictates that usually only worlds outside the context set should be considered. For, since the indicative form is set aside especially for dealing with similar worlds within the context

[3]Lewis (1973) allows for the possibility of there being more than one most similar world to w. This leads to some minor differences between his conditional logic and Stalnaker's, but they are both similar in spirit.

set, use of the subjunctive form is normally to be taken as implying that an indicative conditional would not suffice.

Stalnaker's treatment of context can also be extended to cover counterfactual indicative conditionals like (24). Recall that the shared context need not include everything known by both speaker and hearer. So while they both know that the Earth orbits the Sun, this fact does not have to be included within the background context.

As Stalnaker points out, while context can affect the interpretation of an utterance, the interpretation of an utterance can also affect the context. Often, the effect on context will be simply to update it with the interpreted contents of the utterance. But sometimes, the utterance may establish facts about the initial context. For example, the utterance of (24) may reveal that the known fact that the Earth orbits the Sun should *not* be included within the current context. This may force the hearer to revise their view of what the current shared context actually is, so as to exclude this information. Against this (revised) shared context, the conditional does not count as having an antecedent presupposed to be false, and so an indicative conditional is appropriate.

A presuppositional treatment like this fits in well with the indicative / subjunctive distinction. However, there are problems with taking it as defining the distinction. First, presuppositions are taken to be what is true / false in context. Second, the presuppositional context is not fixed, but may be revised to accommodate new utterances, as with (24). There must therefore be some way of recognising the presuppositional import of a sentence like (24) independently of what is already presupposed by the initial context. Put another way, if the conditional can change the presuppositional context as well as register presuppositions within it, we cannot rely simply on presuppositions to underpin the subjunctive / indicative distinction. This is why an alternate temporal, and non-presuppositional characterisation of the distinction is required.

6.2.2 Back-Shift and Deictic Shift

Dudman and Back-Shift

Dudman (1991) objects to an indicative / subjunctive distinction, even when drawn in temporal terms (Chapter 2.4). Dudman groups modalised indicative and subjunctive conditionals together in the class of 'real conditionals'. Common to real conditionals, he claims, is the fact that the time referred to by the antecedent tense is later than the time indicated by surface form. Thus antecedent present tenses are used to refer to the future, past tenses to the present or future, and past perfects to the past, present or future. In other words, all the tenses are back-shifted.

As shown in Chapter 2.4, the futurate reference of present tense antecedents is nothing out of the ordinary. It can be attributed to a form of deictic shift that occurs to present tenses in conditional antecedents. After revisions are made to Dudman's conditional classification to incorporate certain kinds of mixed tense conditional, modalised indicative conditionals should be included within Dudman's class of 'compound conditionals'.

Dudman's denial of a temporally based subjunctive / indicative distinction relies on being able to group subjunctive conditionals under the same heading as (modalised) indicative conditionals. That is, if modalised indicative conditionals are compound, subjunctive conditionals must be as well.

In compound conditionals, futurate readings for present tense antecedents arise from (secondary) deictic shift. One would therefore expect back-shift in subjunctive conditionals also to be the result of deictic shift.

This is initially plausible in view of the possibility of having deictically shifted past-in-the-future tenses, as in

(27) If I smile when I get out, the interview went well.

(28) By 1998, everybody will know someone who died from AIDS.

However, three arguments can be given against treating conditional back-shift as deictic shift. Consequently, the subjunctive / indicative distinction must be retained.

Back-Shift is not Deictic Shift

1. Deictic Shift of Past Tense Antecedents Futurate past tense antecedents only ever occur in conditionals with hypothetical modal consequents. In non-modalised conditionals or conditionals with present modal consequents, past tense antecedents invariably have past time reference (Chapter 2). That is, whatever mechanism applies in these cases to give futurate present tense antecedents, it does not have the same effect on past tense antecedents.

Therefore, the mechanism that produces back-shift in subjunctive conditionals must be different from the one that produces futurate antecedents in present tense indicative antecedents. This should be obvious anyway — futurate present tense antecedents undergo secondary deictic shift, a possibility not open to past tenses.

2. Futurate Adverbials Back-shifted past tenses in conditionals can readily co-occur with futurate temporal adverbials. Deictically shifted past tenses cannot. Thus (29), where the past tense is back-shifted, is acceptable

(29) If the interview went well tomorrow, I would be very surprised.

but (30), where the past tense is deictically shifted, is not

(30) *If I smile when I get out, the interview went well tomorrow.

Without the futurate adverbial, (30) functions as an ordinary deictically shifted conditional where reference is made to a future interview:

(31) If I smile when I get out, the interview went well.

The difference between (29) and (30) points to different mechanisms underlying the interpretation of the two tenses, and is evidence that conditional back-shift is not an instance of deictic shift.

Comrie (1986) notes, in connection with indirect speech, that opinions differ over whether a sentence like

(32) He will say that he came tomorrow.

is acceptable (for me it is not). It is therefore possible that some people will also find (30) acceptable. But even if this is so, the fact that opinions divide over (30) but not (29) is still evidence that different mechanisms are at work.

3. No Back-Shifted Past Tenses with Past Time Reference A deictically shifted past-in-the-future tense still permits past time reference. Therefore, if back-shift in conditionals were due to deictic shift, we would expect to find back-shifted conditionals where past tense antecedents had past time reference. This never occurs.

This argument may seem definitional. Back-shift occurs when a past tense form has non-past reference. So any past tense form that does have past reference cannot, by definition, be back-shifted. By starting out with this definition of back-shift, perhaps we have already begged the question with regard to deictic shift.

We can go beyond citing definitions. For presuppositions of falsity are symptomatic of subjunctive conditionals, even if we cannot safely use them to define what it is to be a subjunctive conditional. And conditionals with past antecedents referring to past time do not possess the normal counterfactual presuppositions.

Consider (33), which is unproblematically counterfactual and back-shifted.

(33) If the Soviet Union hadn't invaded Afghanistan, the S.U. would have survived beyond 1991.

If back-shift were deictic shift, the following conditional ought to be open to exactly the same counterfactual interpretation

(34) If the Soviet Union didn't invade Afghanistan, the S.U. would have survived beyond 1991.

Yet (33) and (34) differ in meaning: (34) could only be uttered in a context where one makes no presuppositions about whether the Soviet Union invaded Afghanistan.

6.2.3 Back-Shift and Sequence of Tense

If back-shift is not due to deictic shift, what does cause it? The traditional story told about back-shifted antecedents is that there is something like a sequence of tense rule at work (Quirk *et al* (1985), Comrie (1986), Palmer (1990)). That is, a back-shifted past tense corresponds to an underlying present tense where the sequence of tense rule converts the present tense morphology to a past tense form. Similarly, underlying past tenses are converted to past perfect surface forms. Underlying past perfects, if they occur, are left unaltered in surface form, since there is no past past perfect form for them to be back-shifted to. A present tense surface form cannot be the result of the application of a back-shifting sequence of tense rule, since there is no future tense which could be back-shifted to a present tense form.

A sequence of tense rule would predict the transformations shown in Table refseqtense-tab As can be seen, a back-shifted past perfect may correspond to a number of different

Underlying Tense	Surface Realisation
Simple Present	Simple Past
Simple Past	Past Perfect
Present Perfect	Past Perfect
Past Perfect	Past Perfect

Table 6.1: Effects of sequence of tense

underlying tenses. Indeed, a back-shifted past perfect may be very hard to distinguish from an ordinary, non back-shifted past perfect.

Evidence Against Sequence of Tense: Futurate Perfects If a back-shifted past perfect corresponds to an underlying past tense (or present or past perfect), it must have past temporal reference. An example given by Dudman (1983) shows that is not so:

(17) If the auditors had come tomorrow, they would have found the accounts in order.

Here, the past perfect antecedent has futurate reference. This is not the behaviour one would expect of an underlying past tense.

6.2.4 Back-Shift and Subjunctive Tenses

Rather than appealing to a sequence of tense rule, we should take the expression 'subjunctive conditional' at face value. That is, the antecedent and consequent clauses are in the subjunctive tense rather than the indicative tenses ('conditional tense' would perhaps be a happier choice of terminology than 'subjunctive tense').

What makes matters confusing is that subjunctive / conditional tenses have very similar forms to indicative tenses. Adapting Table 6.1 dealing with sequence of tense, we get Table 6.2. The present conditional tense looks a great deal like the indicative past tense.

Conditional Tense	Apparent Indicative Surface Realisation
Present	Simple Past
Past	Past Perfect
Present Perfect	Past Perfect
Past Perfect	Past Perfect

Table 6.2: Conditional tenses: apparent surface forms

The distribution of conditional tenses is limited. For example, in a sentence like *John slept* there is no ambiguity about whether *slept* is in the past indicative or the present subjunctive tense: it is in the past indicative. Conditional tenses only appear in conditionals with hypothetical modal consequents, and perhaps in a few other complement clauses (e.g. complements to *wish* — *I wish I were right about all this*).

Syntactic Evidence

The syntactic evidence in favour of this position is, admittedly, not very strong. The conditional tense forms are mostly identical to past and past perfect forms. Syntactically, there is little reason to parse subjunctive conditionals as having anything other than ordinary tense forms.

One possible exception to this is the use of *were* in some conditionals where *was* would ordinarily be used, e.g. *If I were/was you....* But even here, we could regard *were* as an idiomatic or stylistic variation on *was*, rather than a sign of a distinct subjunctive tense. For example, either *was* or *were* are acceptable in the following non-back-shifted conditional

(35) If Ivan IV was/were in a benign mood, he would kill his victims quickly.

(The antecedent tense refers to the past activities of Ivan the Terrible, and not to what he would be doing today were he alive).

Other evidence in favour of a separate subjunctive tense comes in the form of sentences like

(36) If I had've known, I would have helped you.

Here the perfect auxiliary occurs twice in the antecedent: once to express the subjunctive past, and then a second time to express perfectivity. The ambiguity in form between the simple past conditional, present perfect conditional and past perfect conditional is partly resolved by this more complicated form of construction. On its own, the antecedent

(37) *I had have known.

is unacceptable, suggesting that this construction is peculiar to conditional contexts. (No examples of this kind of construction occur in the LOB or Penn Treebank corpora; I have seen it in written (British) English elsewhere, but its acceptability is apparently dependent on dialect).[4]

Syntactically, there is limited motivation for parsing back-shifted conditionals as having anything other than ordinary tense forms in them. The same is not true for other languages. Portuguese has a separate conditional tense, though in Brazil it is tending to be replaced by the past imperfect tense. And both French and German have subjunctive tenses that are commonly used, and quite distinct in surface form from the non-subjunctive tenses. Cross-linguistic generalisations are a little suspect, especially when applied to tense systems, which can vary greatly even between quite closely related languages. The appeal to a separate conditional tense remains motivated on largely semantic grounds.

6.3 The Logical Structure of Subjunctive Conditionals

6.3.1 Compositional Semantics

Appendix A gives a unification based grammar showing how indicative and subjunctive conditionals are parsed and assigned logical forms. This section describes some of the main features about the logical forms of hypothetical modals and subjunctive conditionals that result.

[4]Sir John Lyons has pointed out that inverted conditionals

(38) Had I known, I would have helped you.

are further evidence for the existence of an English subjunctive tense, given the unacceptabilty of

(39) *Had I known, I will help you.

(40) *Have I known, I would help you.

(41) *Have I known, I will help you.

Hypothetical Modals

Following Perkins' (1983) informal suggestion, the hypothetical modals *would, could, should* and *might* are analysed as implicitly conditional present tense modals, e.g.

(42) would
$\lambda\phi.\ \langle Ante\rangle\langle\Box\!\!\rightarrow\rangle\mathbf{pres}(\mathbf{will}(\phi))$

The antecedent of the implicit conditional, $\langle Ante\rangle$ is specified by context, unless the modal occurs in the consequent of a subjunctive conditional. The implicit connective, $\Box\!\!\rightarrow$ is one of either (i) a present subjunctive conditional, $\Box\!\!\rightarrow_{pr}$, (ii) a past subjunctive conditional, $\Box\!\!\rightarrow_{pd}$, or (iii) a past narrative conditional, $\Box\!\!\rightarrow_{n}$. The abstracted propositional variable, ϕ, corresponds to the sentential argument to the modal operator.

Subjunctive Conditionals

In subjunctive and semi-indicative past narrative or regularity conditionals, the hypothetical modal's implicit antecedent is specified by the explicit conditional antecedent, e.g.

(43) *Present Subjunctive:*
If you asked him, John would go.
$\mathbf{pres}(ask(j))\Box\!\!\rightarrow_{pr}\mathbf{pres}(will(go(j)))$

(44) *Past Subjunctive:*
If you had asked him, John would have gone.
$\mathbf{pres}(ask(j))\Box\!\!\rightarrow_{pd}\mathbf{pres}(will(go(j)))$

(45) *Past Narrative / Regularity:*
If you asked him, John would (usually) go.
$\mathbf{pres}(ask(j))\Box\!\!\rightarrow_{n}\mathbf{pres}(will(go(j)))$

In all these cases, the antecedent and consequent are treated as having an underlying semantic present tense. The syntactic tense forms are determined by the conditional connective. If the present subjunctive conditional is used, both antecedent and consequent must be in the present subjunctive tense. If the past subjunctive conditional is used, both antecedent and consequent must be in the past subjunctive tense. If the past narrative conditional is used, both antecedent and consequent must be in the past (narrative) tense.

Indicative Conditionals

Some conditionals with hypothetical modal consequents are indicative. In such cases, the explicit conditional antecedent does not specify the modal's implicit antecedent, e.g.

(46) If John is here, he could help us.
$\mathbf{pres}(here(j)) \rightarrow [\langle Ante \rangle \square\!\!\rightarrow \mathbf{pres}(\mathbf{may}(help(j)))]$

The modal's conditional still needs to be specified by context. In these cases, the antecedent is in the normal indicative tense. The implicit conditional, $\square\!\!\rightarrow$, may be resolved to any one of $\square\!\!\rightarrow_{\overline{pr}}$, $\square\!\!\rightarrow_{\overline{pd}}$ or $\square\!\!\rightarrow_{\overline{nr}}$.

Adverbial Subjunctives

It is necessary to preserve the distinction between adverbial and conjunctive conditionals, even for subjunctive conditionals. This is to account for sentences such as

(47) If the letter arrived tomorrow, it would already be in the post.
$\mathbf{pres}[\mathbf{pres}(arrive) \square\!\!\rightarrow_{\overline{pr}} \mathbf{will}(inpost)]$

In the adverbial configuration, the present tense assigned to the modal consequent gets wide scope over the conditional. This ensures that the modal consequent has present time reference, despite the future time reference of the antecedent.

Tense Agreement

In subjunctive and semi-indicative conditionals, there must be tense agreement between antecedent and consequent. That is, a past subjunctive *had V-en* antecedent demands a past subjunctive *would have V-en* consequent, and may only be used in conjunction with a past subjunctive connective, $\square\!\!\rightarrow_{\overline{pd}}$. A present subjunctive *V-ed* antecedent requires a present subjunctive *would V* consequent, and may only be used with a present subjunctive connective. A past narrative antecedent *V-ed* requires a past narrative consequent *would V*, and may only be used with a past narrative connective.

At first sight, there appear to be violations of tense agreement, e.g.

(48) If Iraq had not invaded Kuwait, the Middle East would be a lot more stable.

However, the *had V-en* form is ambiguous between the past subjunctive and the present perfect subjunctive. In (48), the antecedent is a perfective present subjunctive, and not a past subjunctive.

The *would have V-en* form is also ambiguous between the past subjunctive and the perfective present subjunctive. In

(49) If the letter arrived tomorrow, it would have been posted today.

the consequent is a perfective present subjunctive, and not a past subjunctive.

Tense agreement is not necessary with indicative conditionals. Its presence in subjunctive conditionals is a direct result of the different kind of connective being used. The surface tenses of the antecedent and consequent do not reflect their underlying semantic tenses (which are uniformly present tense). Instead, they reflect the kind of subjunctive connective being used.

6.3.2 Semantic Definitions for Subjunctive Connectives

The connectives $\Box\!\!\rightarrow_{pr}$, $\Box\!\!\rightarrow_{pd}$ and $\Box\!\!\rightarrow_{nr}$ are versions of the indicative connective $\rightarrow$, except that the indicative connective is evaluated relative to altered information states and / or assertion and verification times. The subjunctive connectives employ a Stalnaker-like selection function, f, to select information states compatible with the antecedent and as similar as possible to the original information state. The selection function also returns a time, corresponding to the earliest verification time of any assertion removed from the information state to ensure compatibility with the conditional antecedent.

Def: $\Box\!\!\rightarrow_{pr}$
$[s, a, v, e] \vdash \phi \Box\!\!\rightarrow_{pr} \psi$ iff
for all s', t' such that $\langle s', t' \rangle \in f(s, \phi, a)$, $t' \geq a$ and $s', t', t', e \models \phi \rightarrow \psi$

Def: $\Box\!\!\rightarrow_{pd}$
$[s, a, v, e] \vdash \phi \Box\!\!\rightarrow_{pd} \psi$ iff
for all s', t' such that $\langle s', t' \rangle \in f(s, \phi, a)$, $t' < a$ and $s', t', t', e \models \phi \rightarrow \psi$

Def: $\Box\!\!\rightarrow_{nr}$
$[s, a, v, e] \vdash \phi \Box\!\!\rightarrow_{nr} \psi$ iff
there exists a $t' < a$ such that $s, t', t', e \models \phi \rightarrow \psi$

The selection function f is discussed in Section 6.4 below. It returns a set of state-time pairs, and the underlying indicative conditional must hold with respect to all the state-time pairs selected.

The main difference between the past and present subjunctive is that the past subjunctive requires that all the times returned by the selection function precede the assertion time for the conditional as a whole. The present subjunctive demands that the times returned are all either simultaneous with or succeed the assertion time.

Because past subjunctives lead to present tense indicative conditionals being evaluated relative to a past assertion time, the antecedent and consequent will typically refer to (hypothetical) past events. There is nothing forcing this to be the case: a futurate present tense evaluated relative to a past time can still refer to a time that is future with respect to the present. This is what happens in Dudman's conditional (17). Present subjunctives, on the other hand, do not refer to past eventualities, unless they are combined with perfective aspect.

The temporal ordering between antecedent and consequent in subjunctive conditionals is exactly the same as for any present tense indicative modalised conditional.

The past narrative conditional is like the past subjunctive in that it leads to a present tense indicative conditional being evaluated relative to a past time. Unlike the past subjunctive, the underlying indicative is evaluated relative to the current information state.

6.3.3 Resulting Temporal Properties

Present Subjunctive

Temporally, present subjunctive conditionals behave exactly like present tense indicative conditionals. A present subjunctive holds with respect to an assertion time a and a verification v iff the corresponding indicative holds relative to the same assertion and verification times. The only difference is that the underlying indicative may be evaluated relative to a slightly different information state.

This explains why conditionals like (18a) and (18b)

(18) a. If the bimetallic strip bent, the temperature would rise.

b. If the bimetallic strip bends, the temperature will rise.

are more or less synonymous, and exhibit the same patterns of temporal reference. As with the present tense indicative (18b), the temporal reference of the consequent modal in (18a) is the earliest possible verification time of the antecedent — usually the (end point of the) antecedent event.

Past Subjunctive

Temporally, past subjunctives behave like present tense indicatives uttered at some earlier time. The past subjunctive shifts the assertion and verification times for the underlying indicative back to some earlier time. This shifts the deictic centre for the present tense antecedents and consequents, allowing them to have past time reference.

In fact, the antecedent and consequent have futurate-in-the-past interpretations. This explains why sentences like Dudman's (17) are possible.

(17) If the auditors had come tomorrow, they would have found the accounts in order.

Again, the temporal reference of the consequent modal depends on the earliest possible verification time of the antecedent.

Past Narrative and Regularity

Like the past subjunctive, the past narrative / regularity conditional sets the assertion and verification times of the underlying indicative to some earlier time. Unlike the past subjunctive, the underlying indicative is evaluated relative to the original information state.

This means that past narrative and regularity conditionals like

(50) If Dick Whittington turned again, he would become Lord Mayor of London.

(51) If the Emperor lowered his thumb, the Christian would be thrown to the lions.

can be used to refer to (actual) historical events, where the antecedent event precedes or is simultaneous with the consequent event.

In fact, past narrative and regularity conditionals potentially can do more than this. As with past subjunctives, both antecedent and consequent have futurate-in-the-past interpretations. In principle, they could refer to future events, as well as historical ones. This seems not to occur, however.

One explanation for the absence of futurate interpretations in these conditionals is to exploit the contextual localisation of tenses mentioned in Chapter 3. Both the past and present tenses must pick event times falling within some contextually given localisation period (where the event time must additionally bear the appropriate relation to the assertion and verification times). With past narrative and regularity conditionals, it is safe to assume that the localisation period will refer to some specific stretch of history (e.g. the Middle Ages, or the late Roman Empire). This forces the antecedent and consequent to have past time reference.

It may seem surprising that past narrative and regularity conditionals are assigned the same logical form. However, as pointed out in Chapter 3, the difference between habitual and non-habitual conditionals is not simply a matter of logical form. Other factors, such as the anaphoric resolution of the tenses, may also determine the kind of interpretation.

Adverbial Subjunctives

It should be apparent from the analysis given for

(47) If the letter arrived tomorrow, it would already be in the post.
$\mathbf{pres}[\mathbf{pres}(arrive) \Box\!\!\rightarrow_{pr} \mathbf{will}(inpost)]$

that the consequent modal will have present time reference, and will not be dependent on the future verification of the antecedent. Adverbial subjunctive conditionals behave in a similar way to adverbial indicatives.

Indicatives

An example of an indicative conditional with a hypothetical consequent is

(46) If John is here, he could help us.
$\mathbf{pres}(here(j)) \rightarrow [\langle Ante \rangle \Box\!\!\rightarrow \mathbf{pres}(\mathbf{may}(help(j)))]$

On the assumption that the implicit conditional is not contextually resolved as a past narrative one, this predicts that John's ability to help is either simultaneous with the (present) reference of the antecedent, or follows it.

In a conditional like

(52) If John is here, he could have helped us.
$\mathbf{pres}(here(j)) \rightarrow [\langle Ante \rangle \Box\!\!\rightarrow \mathbf{pres}(\mathbf{may}(help(j)))]$

it is possible to construe the implicit conditional as a past subjunctive. Thus, John's ability to help may have preceded the present time reference of the antecedent. This interpretation is not open to (46), since there the hypothetical conditional cannot be taken to exhibit a past subjunctive tense.

In conditionals like

(53) If Thor Heyerdahl reaches America, then (maybe) the Ancient Egyptians could build papyrus boats capable of crossing the Atlantic.

it is plausible to construe the implicit conditional as a past narrative / regularity one. Therefore, the Egyptians' ability to cross the Atlantic precedes the futurate reference of the antecedent.

Is the Past Subjunctive Needed?

Could we dispense with the past subjunctive in favour of the present subjunctive plus perfect? That is, could we justify the following kind of analysis

(54) If you had asked him, John would have gone.
$\mathbf{pres}(perf((ask(j))) \Box\!\!\rightarrow_{pr} \mathbf{pres}(\mathbf{will}(perf((go(j)))))$

where *perf* is a perfect operator, selecting an event time that is past relative to the event time chosen by present tense. This would permit past time reference for both antecedent and consequent, as well as future time reference in Dudman's 'auditor' conditional.

This is not possible. In

(22) If he had been found guilty, we could have imprisoned him.

it is not possible to interpret the consequent as a non-epistemic modal plus perfect. The modal itself has past time reference, i.e. in the past it would have been possible for us to imprison him, and not that it would now be possible for us to have imprisoned him in the past.

6.3.4 Thomason & Gupta: Tensed Conditionals

Thomason and Gupta (1981) propose a slightly different formalisation of subjunctive conditionals (van Fraassen's (1981) is essentially the same). They suggest (tentatively) that a subjunctive conditional asserts that its corresponding indicative conditional *was* true at some contextually specified interval. To this end, they formalise (past) subjunctive conditionals as

(55) $Pa[Pr(A) \rightarrow Pr(C)]$

where Pa and Pr are their past and present tense operators[5]. This bears a similarity to the analysis above, which has semantically present tensed antecedents and consequents within the scope of subjunctive conditional connectives. Perhaps the subjunctive connectives could be decomposed into tense operators plus an indicative connective?

We cannot straightforwardly incorporate this treatment into our previous analysis of tense and indicative conditionals. Prefixing a **past** tense operator to a present tense indicative conditional makes no difference to the way in which the conditional is interpreted:

$$\mathbf{past}[\mathbf{pres}(A) \rightarrow \mathbf{pres}(C)] \equiv \mathbf{pres}(A) \rightarrow \mathbf{pres}(C)$$

The reason that this formalisation works for Thomason and Gupta is that they treat their past tense operator Pa in terms of a single indexed tensed logic, i.e.

$Pa(\phi)$ is true at t iff there is some time $t' < t$ such that ϕ is true at t'

Present tenses within the scope of past tenses express presentness relative to a past time. This explains how the present tenses in (55) can be used to refer to past times. Our **past** tense operator does not affect the deictic indices (the assertion and verification times) of any embedded tenses. It would therefore not have the required temporal effects.

Subjunctive Tense Operators?

Of course, there is nothing to prevent us from introducing subjunctive tense operators that do alter the assertion and verification times, for example

[5]In fact, Thomason and Gupta follow Prior in marking the present tense by the absence of any tense operator, so they do not actually define anything like Pr.

(56) Present subjunctive
If it rained tomorrow, the match would be cancelled
$SPres(\mathbf{pres}(rain) \rightarrow \mathbf{pres}(\mathbf{will}(cancel)))$

$s, a, v, e \models SPres(\phi)$ iff $\exists s'$ such that $f(s', s)$ and $s', a, v, e \models \phi$

(57) Past subjunctive
If it had rained, the match would have been cancelled
$SPast(\mathbf{pres}(rain) \rightarrow \mathbf{pres}(\mathbf{will}(cancel)))$

$s, a, v, e \models SPast(\phi)$ iff $\exists s'\ f(s', s)$ and $\exists t \leq a$ such that $s', t, t, e \models \phi$

(58) Past regularity / narrative
If the Emperor lowered his thumb, the Christian would be thrown to the lions.
$NPast(\mathbf{pres}(lower) \rightarrow \mathbf{pres}(\mathbf{will}(throw)))$

$s, a, v, e \models NPast(\phi)$ iff $\exists t \leq a$ such that $s, t, t, e \models \phi$

The present subjunctive tense, *SPres* picks out an information state s' 'relevantly similar' to the current state s. The selection function, f, is a variant on Stalknaker's selection function. The tense leaves the assertion and verification times unaltered. Present subjunctive conditionals consequently behave temporally exactly like present indicative sentences, the non-temporal difference being that the antecedent may in fact be false according to the state s (e.g. *If Churchill was alive today, he would intervene in Bosnia*).

The past subjunctive tense, *SPast*, also picks out a relevantly similar state s'. But by setting the assertion and verification indices to some earlier time, the embedded indicative is treated as though it were uttered in the past, relative to s'. This explains why the antecedent and consequent may have past time reference, but also why they may have future reference in examples like Dudman's (17).

The past narrative tense, *NPast*, does not select a new information state. However, it does reset the assertion and verification indices to some past time. So the embedded indicative is treated as though it were uttered in the past, but relative to the original information state.

Compositionality and the Selection Function

Decomposing the subjunctive connectives into subjunctive tenses plus indicative connectives makes the same temporal predictions as before. But it does not sit easily with the selection of 'relevantly similar' information states.

Stalnaker's selection function takes both a possible world and an (antecedent) proposition as arguments in order to return a similar possible world. But the subjunctive tenses in the formulas above — (56)–(58) — do not have access to the antecedent of the embedded indicative conditional. The only way of providing the required access would be to make the subjunctive tenses non-compositional, so that they could delve into their embedded formula to find the appropriate proposition.

How does Thomason and Gupta's formalisation avoid this problem? They assume that *indicative* conditionals employ a selection function. Consequently, it is not the job of the conditional tense to select a similar possible world — this is done by the conditional itself. And of course, the conditional does have compositional access to the antecedent proposition.

Unlike Stalnaker's selection function, Thomason and Gupta's takes a history h (i.e. a course of events / possible world), a proposition ϕ and a time t as arguments. It returns an alternative history, h', as close as possible to the original, in which the proposition is true.

Thomason and Gupta argue that the selection function should obey the property of *past predominance*. That is, differences between the two histories occurring before the time t count for more in terms of dissimilarity than differences occurring after t.

This means that past subjunctive conditionals, on their account, will select a possible history as much like the actual one up until some past time t, where the value of t is determined by the past tense operator Pa. After t, the actual history and the selected history may diverge. The past subjunctive in effect retreats to an earlier point in history, and considers how things would have been if matters had preceded differently from then on.

Nute (1984) objects to Thomason and Gupta's formalisation, on the grounds that the time selected by the past tense operator is arbitrary. Nute cites an example of using a computer that generates random numbers on the basis of certain locations in its memory state that are affected by other operations the machine is performing. At some point, he stops using the computer to roll a pair of dice, and gets a number 9. He then considers the conditional

(59) If I had used the computer instead of the dice, I would have got a 5.

He goes on to say that 'there is certainly some time in the past such that if I had used the computer at that time I would have got a five, so that a formula corresponding to (59) of the form $[Pa(Pr\phi \rightarrow Pr\psi)]$ is certainly true. Yet (59) [may not be] true. Depending on when I used the computer and what operations the computer had performed before I used it, I could have obtained any integer.'

Presumably, (59) should be interpreted as saying that if at the precise moment I had used the dice I had used the computer instead, I would have got a five. That is, the past time selected should be one at which either some antecedent event that didn't occur might have occurred, or did occur but might not have. But to select the appropriate time it is necessary once more to delve into the embedded conditional formula to retrieve the antecedent proposition. Thus problems with compositionality resurface.

It is possible that one could get round these difficulties by appeal to suitable contextual / anaphoric resolution of the tenses. However, there are worse difficulties for Thomason and Gupta's formalisation. It is unclear how present subjunctive conditionals differ from present indicative conditionals under their account.

Thomason and Gupta are not explicit about the formalisation of present subjunctive conditionals. But given that the antecedents and consequents in such conditionals do not have past time reference, it would have to be something like

(60) $Pr(Pr\phi \rightarrow Pr\psi)$

That is, the time fed into the selection function is the time of utterance. But this is precisely what happens with ordinary indicative conditionals as well. While it may not be entirely unreasonable to claim that conditionals like (18a) and (18b) are equivalent,

(18) a. If the bimetallic strip bent, the temperature would rise.
b. If the bimetallic strip bends, the temperature will rise.

it is inconceivable that conditionals like (61a) and (61b) are equivalent:

(61) a. If Caesar were alive today, he would bomb Serbia.
b. If Caesar is alive today, he will bomb Serbia.

This leads to the same conclusion as Nute's: subjunctive conditionals cannot be understood as a combination of tense operators plus indicative conditionals. Distinct conditional connectives are required.

6.4 State Selection Functions

This section sketches out a way in which selection functions could be defined for information states, but does not pursue the matter in full detail: the preceding section has already achieved the aim of accounting for the temporal properties of subjunctive conditionals. This section serves a rhetorical purpose — to show that the idea of a state selection function is a plausible one — and does not attempt to go into all the details about what these functions are like, and what logical properties emerge as a result.

Most counterfactual logics assume a possible worlds semantics. In this context, selection functions, $f(w, \phi)$, pick out the world or worlds most similar to w in which ϕ is true. Possible worlds are not the same things as information states. Work therefore needs to be done in finding appropriate information state analogues of world selection functions, so that $f(s, \phi)$ picks out the information states most similar to s in which ϕ holds or which is compatible with ϕ.

There are a number of ways in which state selection functions could be defined. Morreau (1992) for example identifies information states with sets of possible worlds. One can then construct a state selection function by imaging on the world selection function:

$$f(s, \phi) = \bigcup_{w \in s} f(w, \phi)$$

(that is, collect together the results of applying the world selection function to all the worlds w contained in the state s). Provided that the world selection function satisfies certain properties, Morreau shows that the state selection function defined gives rise to a well defined conditional logic.

This section approaches matters in a different way, and adapts the basic ideas behind Kratzer's premise semantics (Kratzer 1979, 1981a; Veltman 1976, 1985; Ginsberg 1985 also makes a similar proposal). It shows how, in a simple class of cases, premise semantics can provide the basis for an information state selection function without appeal to objects like possible worlds.

6.4.1 Premise Semantics for Possible Worlds

In the possible worlds version of premise semantics, each world w has a premise set $H(w)$ assigned to it, where a premise set is a set of propositions (Kratzer treats propositions as sets of possible worlds).

The premise set must uniquely characterise the world, w, as ensured by the following centering condition.

- $\bigcap H(w) = \{w\}$

This means that the conjunction of all the propositions contained in the premise set is true of exactly world w, and of no other world.

Usually, there will be more than one set of propositions that uniquely characterise a world. Suppose we have a world w in which exactly the propositions p, q and r are true. Then the following sets of propositions all characterise w:

$\{p, q, r\}$, $\{p \wedge q, r\}$, $\{p, q \wedge r\}$, $\{p \wedge r, q\}$, $\{p \wedge q \wedge r\}$
(As do: $\{p \wedge q, p, q, r\}$, $\{p \wedge q, p, r\}$, and so on)

These different sets correspond to different ways of 'lumping together' the facts in w. Which of these sets counts as the premise set for w depends on context — how one chooses to carve up the world for a particular purpose. The contextually determined choice of a premise set for a world w affects which counterfactual conditionals count as true in w and which do not.

Kratzer proposes the following semantics for counterfactuals of the form "If it were the case that ϕ then it would be the case that ψ" ($\phi \rightarrow_{\Box} \psi$):

- Let $\mathcal{P}$ be the set of all consistent subsets of $H(w) \cup \phi$ that contain ϕ.
  Then $\phi \rightarrow_{\Box} \psi$ is true in w iff for every set $P \in \mathcal{P}$ that is maximal entails ψ.
  (A set P in $\mathcal{P}$ is maximal if there is no proper superset of P in $\mathcal{P}$)[6]

For counterfactuals of the form "If it were the case that ϕ then it might be the case that ψ" ($\phi \rightarrow_{\Diamond} \psi$), we have

- $\phi \rightarrow_{\Diamond} \psi$ is true in w iff there exists a maximal set in $\mathcal{P}$ that is compatible with ψ.

It will help to give an example of how these definitions work.

Let us return to the world w in which p, q and r are true. Consider first the conditional

(A) If it were the case that $\neg p$ it would be the case that r.

And consider (A) relative to two different premise sets:

(H1) $\{p, q, r\}$

(H2) $\{p \wedge r, q\}$

The consistent subsets of $\mathrm{H1}(w) \cup \neg p$, $\mathcal{P}_1$, and $\mathrm{H2}(w) \cup \neg p$, $\mathcal{P}_2$, are

($\mathcal{P}_1$) $\{\neg p\}, \{\neg p, q\}, \{\neg p, r\}, \{\neg p, q, r\}$

($\mathcal{P}_2$) $\{\neg p\}, \{\neg p, q\}$

It should be apparent that every maximal set in $\mathcal{P}_1$, namely $\{q, r\}$, entails r. So according to the premise set H1, the conditional (A) is true in w. But according to premise set H2, (A) is false, since there is no member of $\mathcal{P}_2$ that entails r.

This example illustrates how the choice of premise set can affect the truth of counterfactuals. In H2, the propositions p and r are lumped together. When we remove p from the premise set to make it compatible with $\neg p$, we have to remove r along with it. But in H1, p and r can be removed independently of one another.

Note, however, that relative to both H1 and H2 the counterfactual

[6]In fact, this is not precisely the definition that Kratzer gives. She allows for the fact that compactness might fail for infinite premise sets (compactness: if all finite subsets of a set of formulas are satisfiable, so is the set). For this reason, Kratzer gives a slightly more complicated definition:

– $\phi \Box\!\!\rightarrow \psi$ is true in w iff for every set $P \in \mathcal{P}$ there is a superset of P contained in $\mathcal{P}$ which entails ψ

(B) If it were the case that $\neg p$ it might be the case that r.

counts as true.

Premise Sets and Selection Functions

Premise sets define a form of selection function. According to H1, the most similar $\neg p$-worlds to w are all ones in which r is true. But according to H2, the most similar $\neg p$-worlds to w split into those where r is true and those where r is false.

Lewis (1981) establishes the equivalence between premise semantics and a version of his system of spheres semantics. To derive a similarity ordering from premise sets, let $W = \bigcup H(w)$. For any $u, v \in W$ let $u \succeq_w v$ iff all propositions in $H(w)$ that hold in v also hold in u, but some propositions in $H(w)$ hold in u but not v. Here, $u \succeq_w v$ means that u is more similar to w than v.

The system of spheres semantics does not, strictly speaking, employ world selection functions. Lewis's denial of the limit assumption means that he does not hold that there are necessarily such things as the *closest* worlds to w in which some proposition ϕ is true. Instead, a counterfactual like $\phi \rightarrow_{\Box} \psi$ holds in w if there is a set of worlds surrounding w where ϕ is true in at least some of these worlds, and for all worlds as close or closer to w in which ϕ holds, ψ also holds. But if one is dealing with a finite set of worlds in total, or one is happy to accept the limit assumption, this turns out to be the same as selecting a set of closest ϕ-worlds.

In (Lewis 1973) it is assumed that $\succeq_w$ is a total order, i.e. that the relative similarity of two worlds to w can always be compared. This is not so with the similarity ordering defined from premise semantics — incomparabilities are possible. However, Lewis (1981) argues that this incomparability can be explained by means of a supervaluation on comparable orderings: where there is an incomparable ordering, this corresponds to a number of distinct possible total orders, and the conditional must hold when the comparison is made in all possible ways.

Lewis (1981) also shows that the equivalence between premise semantics and a version the system of spheres semantics holds even when the centering condition on premise sets is dropped. When a premise set contains propositions that are all true of the world, but which together do not uniquely characterise the world, it is *weakly centred.* Lewis (1973) provides an axiomatisation for a conditional logic with a weakly centred system of spheres, **VW**:

A0. Axioms and rules of inference of classical propositional logic
A1. $\phi \rightarrow \phi$
A2. $(\phi \rightarrow \psi) \supset (\phi \supset \psi)$
A3. $(\neg\phi \rightarrow \phi) \supset (\psi \rightarrow \phi)$

A4. $((\phi \rightarrow \psi) \wedge (\psi \rightarrow \phi)) \supset ((\phi \rightarrow \chi) \equiv (\psi \rightarrow \chi))$
A5. $((\phi \rightarrow \psi) \wedge \neg(\phi \rightarrow \neg\phi)) \supset ((\phi \wedge \chi) \rightarrow \psi)$

$$RCEC\frac{\phi \equiv \psi}{(\chi \rightarrow \phi) \equiv (\chi \rightarrow \psi)}$$

$$RCK\frac{(\phi_1 \wedge \ldots \phi_n) \supset \psi}{((\chi \rightarrow \phi_1) \wedge \ldots \wedge (\chi \rightarrow \phi_n)) \supset (\chi \rightarrow \psi)}$$

To account for the partial similarity ordering resulting from premise sets, the axiom A5 should be replaced by

A5'. $[(\phi \rightarrow \psi) \wedge (\chi \rightarrow \psi)] \supset [(\phi \vee \chi) \rightarrow \psi)]$

(Lewis does not discuss this logic, which I shall call **VW'**)

6.4.2 Premise Semantics for Information States

In premise semantics, the truth of a counterfactual does not depend on the identity of an individual world so much as on the premise set associated with it. Distinct worlds with identical premise sets will support the same conditionals.

Of course, while centering of premise sets remains in force, the possibility of different worlds having the same premise set will not arise. But once this requirement is dropped (so that, e.g., the premise set represents what is known in a world rather than everything that is true in the world), then different worlds can have the same premise set.

Without centering of premise sets, it becomes permissible to assign premise sets to information states. With centering, a premise set identifies a possible world; assigning a premise set to an information state would erroneously tie it to a specific possible world. But without centering, no such link between information states and possible worlds is imposed.

One might therefore conjecture that if premise sets are assigned to and centred on information states instead of possible worlds, the result would be the weakly centred logic **VW'**. I leave this as a conjecture. Note that it still presupposes an underlying possible worlds semantics. Below, a rather different approach is followed.

Motivation

The basic idea is to take the set of atomic assertions supported by an information state, and use these to form the state's premise set. Subsets of the premise set of a state s will correspond to premise sets of states informationally preceding s. By taking subsets of the premise set consistent with some conditional antecedent, one is in effect specifying preceding information states compatible with the antecedent. Thus, taking subsets of the premise set gives rise to a form of selection function.

The temporally specific nature of the selection function arises from the temporal reference of the assertions removed from the premise set to ensure compatibility. The time selected by the function corresponds to the earliest verification time of any assertion removed.

Assertion Sets

Definition: Basic Assertions The basic assertions of an information state are those that are supported by the state without reference to preceding or succeeding information states.

- $s \vdash_b p(t)$ if $\exists t_v.\ V(s, t_v, p, t)$
- $s \vdash_b \phi \wedge \psi$ if $s \vdash \phi$ and $s \vdash \psi$

The basic assertions of a state are temporally specific eternal propositions.

Definition: Assertion Set The assertion set of an information state is the set of basic assertions supported by that state

- $Ass(s) = \{\phi : s \vdash_b \phi\}$

To prevent assertion sets from containing infinite numbers of equivalent conjunctions, we treat assertion sets as being sets of propositions rather than formulas.

Definition: Conditional Assertions A (first degree) conditional assertion is one that depends on basic assertions in other information states.

- $s \vdash_c \phi \rightarrow \psi$ iff for all $s' \sqsupseteq s$ if $s' \vdash_b \phi$ then $s' \vdash_b \psi$

I am simplifying matters here by ignoring non-monotonic conditionals (e.g. $\phi \rightarrow \mathbf{may}(\psi)$).

We can go beyond first degree conditional assertions if we allow conditional assertions as antecedents and/or consequents to conditionals. Defining assertions supported by a state (basic and conditional) as

- $s \vdash \phi$ iff $s \vdash_b \phi$ or $s \vdash_c \phi$

we can then define conditional assertions of any degree as

- $s \vdash_c \phi \rightarrow \psi$ iff for all $s' \sqsupseteq s$ if $s' \vdash \phi$ then $s' \vdash \psi$

However, I will not dwell on conditional assertions of greater than degree one.

Definition: Non-Trivial Conditional Assertions Given the monotonicity of information extension, if ϕ and ψ are basic assertions supported by state s, then the conditional assertion $\phi \rightarrow \psi$ is also trivially supported by s. This is not to say that there is actually some kind of conditional connection between ϕ and ψ, merely that they both happen to be asserted and will continue to be so.

It is useful to identify non-trivial conditional assertions. To do this, we must retreat to a preceding state where neither antecedent nor consequent hold, and see if this state supports the conditional assertion.

- $s \vdash_{ntc} \phi \rightarrow \psi$ iff $\exists s'.\ s' \sqsubseteq s$ such that
  (a) $s \not\vdash \phi$,
  (b) $s \not\vdash \psi$, and
  (c) $s \vdash_c \phi \rightarrow \psi$

It should be borne in mind that the assertions occurring in the antecedent and consequent of a conditional assertion will be temporally specific, eternal propositions. Conditional assertions thus connect an instance of one proposition being true at one specific time with an instance of another proposition being true at another time. In this respect, they are unlike many conditional sentences in English, which express a more general connection between propositions holding at non-specific times.

Assumption: State Existence We need to allow ourselves plenty of information states in what follows. We therefore assume that if you take any consistent set of basic propositions, it will be logically equivalent to the assertion set of some information state:

- Let Γ be any consistent set of basic propositions.
  Then $\exists s$ such that $\Gamma \equiv Ass(s)$

Two sets of propositions are equivalent, $\Gamma \equiv \Delta$, if one set classically entails all the propositions contained in the other, and vice versa.

Lemma: State Ordering The informational ordering of states must respect the conditional assertions holding in states: if a conditional assertion holds in an earlier state, it must also hold in all later states. This follows trivially from the definition of a conditional assertion. This licences the following trivial lemma, which will be useful later on.

State Ordering Lemma
For all s_1, s_2 such that $s_2 \sqsubset s_1$:
for all ϕ, ψsuch that
$Ass(s_1) \vdash \phi$, $Ass(s_1) \vdash \psi$ and $Ass(s_2) \vdash \phi$
if there exists an s' such that $s' \sqsubset s_1$ and $s' \sqsubset s_2$
and $s' \vdash_{ntc} \phi \rightarrow \psi$
then $Ass(s_2) \vdash \psi$

(Note: I use the turnstile, $\vdash$, both to express which propositions are asserted by a state, and to express which propositions follow classically from a set of propositions.)

Premise Sets

Premise sets consist of two sets of assertions: the set of basic assertions supported by a state, and the set of non-trivial conditionals supported by the state. These will be referred to as the basic premise set and the conditional premise set respectively. Some conditions need to be imposed on basic premise sets.

Definition: Centering A basic premise set, $P(s)$ is centred (on s) as follows:

- Centering: $P(s) \equiv Ass(s)$

In other words, a centred basic premise set incorporates all and only the basic assertions made by the information state. Centering here is an obvious analogue of centering in the possible worlds case, although centred premise sets do not uniquely characterise information states. The same premise set may be centred on more information state, provided that the two states make the same basic assertions (the states may differ in their conditional assertions).

Definition: Weak Conditional Lumping Weak conditional lumping ensures that basic premise sets convey a certain amount of information about the non-trivial conditional assertions holding in a state.

- Weak Conditional Lumping:
  If $s \vdash_{ntc} \phi \rightarrow \psi$ and $s \vdash \phi$ (and so $s \vdash \psi$), then $\phi \notin P(s)$

A simple example will illustrate the effect of this condition. Suppose state s asserts p and q, and it holds non-trivially that $p \rightarrow q$. The assertion set for s will be $\{p, q, p \wedge q\}$. Weak conditional lumping demands that p not be a member of $P(s)$. If $P(s)$ is also centred, this means that the options for $P(s)$ are either $\{q, p \wedge q\}$ or $\{p \wedge q\}$. In both cases, p has been absorbed into a conjunction with its antecedent. This means that if you were to find a subset of $P(s)$ consistent with $\neg q$, the subset would not entail p.

Weak conditional lumping ensures a kind of behaviour similar to that found in belief revision systems (Gärdenfors 1988). Suppose a belief state holds that p and *if p then q*, and so consequently also q. If you want to retract q, it is necessary to retract either p or *if p then q*. By lumping p together with its consequent q, we ensure that it is p that is retracted, rather than the conditional.

Definition: Strong Conditional Lumping Strong conditional lumping is an extension of weak conditional lumping:

- Strong Conditional Lumping:
  If $s \vdash_{ntc} \phi \rightarrow \psi$ and $s \vdash \phi$ (and so $s \vdash \psi$), then $\phi \not\in P(s)$ and $\psi \not\in P(s)$

In the example we gave above, this would mean that the only permissible premise set $P(s)$ would be $\{p \wedge q\}$. Strong conditional lumping ensures that should you retract the antecedent of a conditional, the consequent gets retracted with it. It reflects the view that the only reason that the antecedent held was because the consequent held. This is not always a correct view to take, but Tichy's (1976) 'man with a hat' example requires it (see p. 193).

Main Theorem: Subsets of Basic Premise Sets are Basic Premise Sets

We now come to stating and proving the theorem that allows us to forge a link between maximally consistent subsets of basic premise sets and 'similar' information states.

- Theorem:
  Let $P(s)$ be a basic premise set for s that satisfies centering and weak conditional lumping.

  For any subset, $P' \subset P(s)$, there exists an s' such that $s' \sqsubset s$ and P' is a centred, weakly lumped basic premise set for s'

To prove this, we divide the theorem into two parts

(A) $\forall s'$ such that $s' \sqsubset s$ and $P' \equiv Ass(s')$:
P' is centred on s', and
P' is weakly lumped wrt s'

(B) $\exists s'$ such that $s' \sqsubset s$ and $P' \equiv Ass(s')$

To prove part (A), that P' is centred on s' is a direct consequence of $P' \equiv Ass(s')$. To show that P' is weakly lumped: Suppose $\phi, \psi \in Ass(s')$ and $\exists s''$ such that $s'' \sqsubset s'$, $s'' \not\vdash \phi$, $s'' \not\vdash \psi$ and $s'' \vdash \phi \rightarrow \psi$. This state also precedes s. Since s is lumped, it follows that $\phi \notin P(s)$, and as P' is a subset of $P(s)$, $\phi \notin \mathcal{P}'$. Hence, P' is weakly lumped.

Proving part (B) is slightly harder. Suppose there is no s' such that $s' \sqsubset s$ and $P' \equiv Ass(s')$ and derive a contradiction. That is, $\forall s'$ if $s' \sqsubset s$ then $P' \not\equiv Ass(s')$. From the state ordering lemma, this entails

$$
\begin{array}{l}
\forall s', \text{ if} \\
\quad (\forall \phi, \psi \text{ such that } Ass(s) \vdash \phi,\ Ass(s) \vdash \psi\ Ass(s') \vdash \phi \\
\qquad\qquad \text{if } \exists s'' \text{ s.t. } s'' \sqsubset s,\ s'' \sqsubset s',\ s'' \not\vdash \phi\ s'' \not\vdash \psi \text{ and } s'' \vdash \phi \rightarrow \psi \\
\qquad\qquad\qquad \text{then } Ass(s') \vdash \psi) \\
\text{then} \\
P' \not\equiv Ass(s')
\end{array}
$$

Putting this in the contrapositive, exploiting the assumed existence of states for any assertion set so that $\exists s'.\ P' \equiv Ass(s')$, plus the fact that $Ass(s) \equiv P(s)$ we get

$$
\begin{array}{l}
\exists \phi, \psi \text{ such that } P(s) \vdash \phi,\ P(s) \vdash \psi\ P' \vdash \phi,\ P' \not\vdash \psi \\
\qquad \text{and } \exists s'' \text{ s.t. } s'' \sqsubset s,\ s'' \sqsubset s',\ s'' \not\vdash \phi\ s'' \not\vdash \psi \text{ and } s'' \vdash \phi \rightarrow \psi
\end{array}
$$

From this, it follows that $P(s)$ is lumped by $\phi \rightarrow \psi$, and so if $P(s) \vdash \phi$ then $P(s) \vdash \phi \wedge \psi$. It can be shown that

> If $P(s)$ is weakly lumped and $P(s) \vdash \phi$ implies $P(s) \vdash \phi \wedge \psi$,
> then for any subset P' of $P(s)$, $P' \vdash \phi$ implies $P' \vdash \phi \wedge \psi$

From this we get a contradiction, because $P(s) \vdash \phi$ implies $P(s) \vdash \phi \wedge \psi$ but $P' \vdash \phi$ does not imply that $P' \vdash \phi \wedge \psi$.

Deriving a Selection Function

How does one evaluate a counterfactual conditional given a premise set consisting of a basic premise set and a conditional premise set? Take maximal subsets of the basic premise set which together with the conditional premise set are consistent with the conditional antecedent. Then evaluate the corresponding present tense indicative conditional relative to the states corresponding to these reduced basic premise sets. The time at which the conditional is evaluated in such a state is the earliest time at which any of the assertions removed from the original premise set could have been verified (typically, the end point of the assertion's event time).

Example

As an example of the way this all works, consider Dudman's conditional

(17) If the auditors had come tomorrow, they would have found the accounts in order.

We will consider what happens if this conditional is uttered at time $t = 2$, where I am busy doing the accounts, but where the auditors have already visited at time $t = 1$. Tomorrow will be represented as time $t = 3$. Let us give some background domain rules:

$A(t)$ means 'the auditors come at time t'
$D(t)$ means 'the accounts are done at time t'
$F(t)$ means 'the auditors find the accounts in a mess at t

Conditionals: 1. $\forall t_1, t_2.\ A(t_1) \wedge A(t_2) \rightarrow t_1 = t_2$
(The auditors only come once)

2. $\forall t_1, t_2.\ D(t_1) \wedge D(t_2) \rightarrow t_1 = t_2$
(The accounts only get done once)

3. $\forall t_1, t_2.\ F(t_1) \wedge F(t_2) \rightarrow t_1 = t_2$
(The accounts are only found to be a mess once)

4. $\forall t_1, t_2.\ A(t_1) \wedge D(t_2) \wedge t_1 < t_2 \rightarrow F(t_1)$
(The accounts are found to be a mess if the auditors arrive before they are done)

5. $\forall t_1, t_2.\ A(t_1) \wedge D(t_2) \wedge t_1 < t_2 \rightarrow \neg F(t_1)$
(The accounts are not found to be a mess if the auditors arrive after they are done)

6. $\forall t_1.\ F(t_1) \rightarrow A(t_1)$
(If the accounts are found to be a mess, it happens when the auditors visit)

Now consider the assertion set for a state s in which $A(1)$ and $D(2)$ holds and where by implication $F(1)$ holds. This will be

$$\{A(1), D(2), F(1), A(1) \wedge D(2), A(1) \wedge F(1), D(2) \wedge F(1), A(1) \wedge D(2) \wedge F(1)\}$$

Weak conditional lumping means that this set will be cut down to the basic premise set P

$$P = \{A(1), D(2), A(1) \wedge F(1), A(1) \wedge D(2) \wedge F(1)\}$$

Given that the conditional premise set will contain $A(3) \rightarrow \neg A(1)$, it follows that the maximal subset of P consistent with $A(3)$ is $\{D(2)\}$. This means that to evaluate the counterfactual conditional, we go back to a state where just $D(2)$ holds as a basic assertion, and see if the indicative conditional $A(3) \rightarrow \neg F(3)$ holds there relative ime $t = 1$, which it does.

Although this example seems quite trivial, it is significant that while the assertions $A(1)$ and $F(1)$ have both been removed, the assertion $D(2)$ pertaining to a later time remains. Thus, the suggested treatment of counterfactuals is not the same as that proposed by Dudman, where one goes back to the state of the world at an earlier time ($t = 1$) and then lets a 'projective fantasy' take over. This fantasy would not guarantee that $D(2)$ still holds.

The role that strong conditional lumping plays can be seen by reference to Tichy's 'man with a hat' example. Here, we have a man who always wears his hat when it rains, but wears it only 50% of the time when it is not raining. Currently it is raining and the man is wearing his hat. Is it true that *if it was not raining, the man would have his hat on*?

With only weak conditional lumping, one could remove the assertion that it is raining while leaving the assertion that the man has his hat on intact — the basic premise set is $\{rain \wedge hat, hat\}$. This would predict that the conditional is true, whereas in fact it is not: the man might or might not be wearing his hat. But if the premise set is strongly lumped, the premise set is $\{rain \wedge hat\}$, and so the assertion that he is wearing his hat would also be removed.

Generalisation

The treatment of counterfactuals above assumes that non-trivial conditional assertion are sacrosanct, and that the counterfactual antecedent is only incompatible with basic assertions. This is an unreasonable limitation, and the treatment needs to be generalised.

One (conjectured) way of doing this is as follows:

$f(s, \phi, a) =$
$\langle s_\phi, t_\phi \rangle : t_\phi < a$ and
$\exists t', s' \leq_{t_\phi} s$ s.t. $s', t_\phi, t, t \not\models \neg\phi$ and $not(\exists t''_\phi, s''$ s.t. $a \geq t''_\phi > t_\phi,\ s \sqsubseteq_{t''_\phi} s'' \sqsubseteq_{t''_\phi} s'$
and $s'', t''_\phi, t, t \not\models \neg\phi)$
and $s_\phi \geq_{t_\phi} s'$, $P(s_\phi) \subseteq P(s)$ and
$not(\exists s'_\phi$ s.t. $s'_\phi \geq_{t_\phi}$ and $P(s_\phi) \subset P(s'_\phi \subseteq P(s))\}$

In this definition, s' is a state preceding s that is compatible with the antecedent, such that no other state s'' between s' and s is also compatible with the antecedent. The selected state s_ϕ extends s', and supports as much of state s's premise set as possible. As before, the way in which the premise set is lumped affects the behaviour of the selection function.

If the antecedent formula ϕ is compatible with the non-trivial conditionals holding in s, then $s_\phi \sqsubseteq s$. In such cases, the states selected by $f(s, \phi, a)$ will be just those that would be selected by the procedure above. In other cases s and s_ϕ may not be directly related to each other. This matter is not pursued here.

6.4.3 The Gärdenfors Triviality Result

The operations carried out on premise sets when evaluating counterfactuals are reminiscent of analyses of counterfactuals in terms of belief revision and the Ramsey Test (Gärdenfors 1988). One makes minimal revisions to the belief / information state to make it compatible with the antecedent, and then sees whether the consequent follows once the antecedent is added to the revised state. Gärdenfors has shown that this kind of belief revision based treatment of counterfactuals faces a serious problem. It is impossible to use the Ramsey test in conjunction with certain kinds of non-trivial belief revision systems.

Gärdenfors lays down a number of rationality postulates for the revision, expansion and contraction of belief states (while not essential, it is convenient to view belief states as sets of sentences in a language L closed under some form of logical consequence, $\vdash$). The revision, expansion and contraction of a belief state K relative to a sentence ϕ are symbolised as K^*_ϕ, K^+_ϕ and K^-_ϕ respectively. K^*_ϕ makes the minimal changes necessary to K for it to consistently contain ϕ, K^-_ϕ makes the minimal changes necessary to ensure that the belief state no longer includes ϕ, and K^+_ϕ simply adds ϕ to the belief state, even if the result is inconsistent.

The Ramsey Test (RT) amounts to

RT: $\phi \Box\!\!\rightarrow \psi \in K$ iff $\psi \in K^*_\phi$

That is, $\phi \Box\!\!\rightarrow \psi$ holds in K if minimally revising K to support ϕ leads to ψ also holding. The Ramsey Test as originally conceived is intended to apply to indicative as well as counterfactual / subjunctive conditionals. In the indicative case, the revision of K to support ϕ is simply an extension of K to support ϕ, K^+_ϕ. In terms of information states, the revision K^*_ϕ can be cast as (i) selecting an information state compatible with ϕ and then minimally extending it to support ϕ.

A non-trivial belief revision system is one where there are at least two contingent sentences in L. Letting p and q be two such sentences, in a non-trivial system there will be states compatible with p and q, p and $\neg(p \wedge q)$, and with q and $\neg(p \wedge q)$, and none except the absurd state belief state $K_\perp$ are compatible with all three of p, q and $\neg(p \wedge q)$.

Some general principles required for the impossibility proof are:

K*2: $\phi \in K^*_\phi$

K*P: If $\psi \in K$ and $\neg\phi \notin K$, then $\psi \in K^*_\phi$

K*5': If $K^*_\phi = K_\perp$, $\vdash \neg\phi$

K*M: If $K \subseteq K'$, then $K^*_\phi \subseteq K'^*_\phi$
(Entailed by RT: suppose $\psi \in K^*_\phi$. By RT, $\phi \Box\!\!\rightarrow \psi \in K$, and hence $\phi \Box\!\!\rightarrow \psi \in K'$. Thus, by RT again, $\psi \in K'^*_\phi$)

K$^+$: $K^+_\phi = \{\psi : K \cup \{\phi\} \vdash \psi\}$

The proof proceeds as follows.

1. Let p, q, r,be three sentences establishing non-triviality for a state K such that $p \wedge q$, $p \wedge r$ and $q \wedge r$ are all compatible with K (i.e. $\neg(p \wedge q) \notin K$ etc.), but $(p \wedge q \wedge r) \vdash \perp$
2. Since $\neg(p \wedge q) \notin K$, by the deduction theorem for $\vdash$ $\neg p \notin K^+_q$.
3. By K*P, $q \in (K^+_q)^*_p$
4. By RT, $p \Box\!\!\rightarrow q \in K^+_q$
5. By K*M, $p \Box\!\!\rightarrow q \in K^+_{q \wedge r}$
6. Repeating (2)–(5), $p \Box\!\!\rightarrow r \in K^+_{r \wedge q}$
7. $K^+_{r \wedge q} = K^+_{q \wedge r}$
8. By RT, $r, q \in (K^+_{q \wedge r})^*_p$
9. By K*2, $p \in (K^+_{q \wedge r})^*_p$
10. Thus $(p \wedge q \wedge r) \in (K^+_{q \wedge r})^*_p$, hence $(K^+_{q \wedge r})^*_p = K_\perp$
11. By K*5', $\vdash \neg p$. Contradiction.

Short of outright rejection, it is hard to modify the Ramsey Test in such a way that the impossibility result does not follow. One plausible modification takes the hint from the fact that the Ramsey Test suffers from a positive version of the paradoxes of implication: if ϕ and ψ both hold in a state K, then $\phi \Box\!\!\rightarrow \psi$ (trivially) holds in K. One might therefore weaken the Ramsey Test to

WRT: $\phi \Box\!\!\rightarrow \psi \in K$ iff $\psi \in (K^-_\psi)^*_\phi$

That is, first retract the consequent, and then revise to add the antecedent. Unfortunately, the test still entails a weaker version of the monotonicity property,

WK*M: If $K \subseteq K'$ and $\psi \notin K'$, then $\psi \in K'^*_\phi$ if $\psi \in K^*_\phi$

and this is enough to establish an impossibility result (Gärdenfors 1988).

This result does not afflict the treatment of subjunctive conditionals given previously and indicates that, contrary to appearances, the Ramsey Test is not employed. The point at which the proof fails is step (5). At step (4), K_q^+ supports the conditional $p \Box\!\!\rightarrow q$ only trivially, meaning that it does not have to be preserved when selecting states compatible with some further formula. The same holds with the extension of K_q^+, $K_{q \wedge r}^+$. And precisely because p is incompatible with $q \wedge r$, there is no reason to suppose that $p \Box\!\!\rightarrow q$ (and consequently q) holds in $(K_{q \wedge r}^+)_p^*$.

6.5 Conclusions

This chapter has shown how the temporal properties of subjunctive conditionals can be accounted for by assuming that the apparently past tense forms in antecedent and consequent are manifestations of subjunctive tenses. Semantically, subjunctive conditionals are present tense indicative conditionals, evaluated to information states and assertion times picked out by an information state selection function, analogous to the selection functions employed in possible world analyses of counterfactual conditionals. (The discussion of state selection functions is only intended to indicate how a fully worked out treatment might proceed).

This illustrates the point that one cannot come to grips with subjunctive conditionals until one has an adequate account of indicative conditionals: indicative conditionals are both temporally and logically primary.

Chapter 7

Conclusions and Further Directions

Our ordinary language shows a tiresome bias in its treatment of time. Relations of date are exalted grammatically as relations of position, weight and color are not. This bias is of itself an inelegance, or breach of theoretical simplicity. Moreover, the form that it takes — that of requiring every verb form to show a tense — is peculiarly productive of needless complications, since it demands lip service to be paid to time even when time is farthest from our thoughts. (Quine 1960:170).

7.1 Summary

This thesis has been concerned with the behaviour of the past and present tenses in modal and conditional sentences. The opinion is often voiced that the interactions between time and modality are crucial to understanding both, and that time, especially future time, has an irreducibly modal dimension, while modality has an irreducibly temporal dimension. Despite this, the interactions between tense, modals and conditionals in natural language have received surprisingly little attention, even at a purely descriptive level.

The first goal of the thesis was to describe what the interactions are. This goal is attained in two stages. The first part of Chapter 2 sets out the patterns of temporal reference exhibited by simple and modalised indicative conditional sentences, as well as non-conditional sentences containing present tense modal auxiliaries. The first part of Chapter 6 does the same for subjunctive conditionals, and shows how their temporal properties parallel those of present tense indicative conditionals. In both cases, the regularities reported are the result of a systematic survey performed by varying the tense and aspect of the antecedent and consequent clauses and the assumed pragmatic connection between antecedent and consequent. These results have been checked against roughly 2500 conditional sentences occurring in the LOB corpus of British English and the Wall St Journal extracts contained in the Penn Treebank.

It is generally acknowledged that subjunctive conditionals are temporally problematic: they contain apparently simple past tenses that can only be understood as referring to the future. This acknowledgement has diverted attention away from the more fundamental problems confronting the temporal interpretation of indicative conditionals.

The data presented in Chapter 2 raises three questions. First, why is it that in certain modal and conditional contexts, past and present tenses can be deictically shifted so that they refer to future times? Second, why do the past and present tenses behave asymmetrically in this respect? There are contexts in which present tenses can refer to the future while past tenses can only refer to the past. Third, there are some strong semantic constraints on the temporal ordering between eventualities described by the antecedent and consequent clauses, but these depend on the tenses of the antecedent and consequent. What gives rise to these constraints? No previous accounts of tenses, modals and conditionals offer satisfactory answers to all these questions.

The asymmetric behaviour of deictically shifted past and present tenses has not been noticed before, and provides the key to answering the other two questions. As long as one takes the tenses to have a single deictic centre, their asymmetry will remain elusive. Two deictic centres are required: a primary centre, known as the assertion time, and a secondary centre, known as the verification time.

The past tense expresses pastness with respect to the primary deictic centre and (typically) pastness or simultaneity with respect to the secondary deictic centre. The present tense expresses non-pastness with respect to the primary deictic centre and (typically) pastness or simultaneity with respect to the secondary centre. In normal contexts, both the primary and secondary deictic centres correspond to the time at which a sentence is uttered, the speech time. In these contexts, the past tense expresses pastness relative to the speech time, and the present tense expresses presentness relative to the speech time.

Modals and conditionals give rise to a mixture of primary and secondary deictic shift. Tenses in conditional antecedents undergo secondary deictic shift, so that the verification time may succeed the speech time. This leaves the interpretation of past tenses unaffected, since they must still refer to times preceding the unshifted assertion time. But it allows present tenses to refer to times after the speech time.

Tenses in conditional consequents can undergo primary and/or secondary deictic shift. Primary deictic shift moves the assertion time forward, and allows both past and present tenses to refer to future times. Secondary deictic shift, when it occurs, allows present tenses to refer to times that are future relative even to the shifted assertion time. There are in fact two kinds of conditional, differing in that one permits only primary deictic shift in the consequent ($\rightarrow$), and the other permits both primary and secondary deictic in the consequent ($\Rightarrow$). Modals give rise to primary and secondary deictic shift, subject to the constraint that the secondary deictic centre may not precede the primary deictic centre.

Returning to the first question, why do modals and conditionals give rise to deictic shift, and why are there two deictic centres? The presence of two deictic centres results from a view of the tenses as not only describing the way that the world changes over time, but also describing the way that information about the world changes over time.

Information comes in three categories: verified, unverified and indirect. With verified information, there is some direct verifying link, often causal, between the information and the state of affairs it is about. With unverified information, this link has yet to be established, but the information is still there in unverified form. Information about the future is typically unverified, since the future events have to take place before the information can finally be verified. A given information state will usually support both verified and unverified information, with more and more of the unverified information becoming verified as time goes by.

Indirect information concerns what information will be supported when further information comes in, but which is currently neither verified nor unverified. To say that some fact may or must be the case is to say that information will develop in such a way that the fact may or must eventually come to be verified in information states extending the current one. Beyond this, the current information says nothing about the fact.

The primary deictic centre, the assertion time, corresponds to the time at which an information state is extended to support an unverified assertion conveyed by the utterance of a sentence. The secondary deictic centre, the verification time, corresponds to the time at which this unverified assertion becomes verified.

Simple, non-modal, non-conditional utterances are used to convey verified information; this is why the assertion and verification times are both initially taken to be the speech time. A modal like *will* is used to convey unverified information: it asserts the truth of some fact, but qualifies this assertion by saying that it will not be until later that this assertion can be verified. Other modals and conditionals convey indirect information. *Must* and *may* describe what will be asserted and verified in subsequent information states. The conditional says that if you were to extend the current information state by asserting the antecedent, then as soon as this assertion becomes verified, a perhaps unverified assertion of the consequent will also be supported.

The semantic constraints on temporal ordering in conditionals follow from this last point. The primary deictic centre of the consequent is furnished by the secondary deictic centre of the antecedent — when the antecedent is verified, you can assert the consequent. When both antecedent and consequent are past tensed, they must both refer to times preceding the primary and secondary deictic centres, resulting in a lack of any semantically imposed ordering. But when the antecedent and consequent are present tensed, the antecedent time must precede its secondary deictic centre, while the consequent must succeed its primary deictic centre (the antecedent's secondary centre). As a result, an antecedent precedes consequent ordering is semantically enforced.

Chapters 3 and 4 set out the consequences of this approach to tense and information change in detail, and demonstrated that it successfully accounts for the patterns of temporal reference and deictic shift exhibited by indicative conditionals and sentences containing present tense modal auxiliaries. This fulfills the first part of the thesis' second goal: a semantic account of tense, modals and conditionals that makes the correct temporal predictions. It also goes substantially beyond this goal. It was shown how the same treatment could handle futurate uses of the present tense referring to planned or pre-determined events, habitual uses of the past and present tenses, the similarities and differences between the word *if* and the temporal connective *when*, and other non-modal, non-conditional subordinate tense constructions.

The second part of the second goal requires accounting for the temporal properties of subjunctive conditionals. Chapter 6 showed how this can be done by building on the semantics of indicative conditionals. The apparent temporal anomalies of subjunctive conditionals are shown to result from the occurrence of subjunctive tenses in these conditionals. Once this is realised, the temporal properties fall out quite smoothly.

It has not been the goal of this thesis to develop a logic for modals and conditionals: the aim has been temporal adequacy, not logical adequacy. However, it would be foolhardy to ignore logical questions altogether. Chapter 5 therefore demonstrated that the temporal perspective on tense and conditionals leads to an interesting and plausible logic of conditionals, which moreover naturally gives rise to three forms of epistemic modality corresponding to *must*, *may* and *will*. This confirms the status of *will* as a full modal auxiliary, and not a mere marker of future tense as some have supposed (e.g. Wekker 1976).

In a similar vein, the final part of Chapter 6 showed that the temporal perspective on subjunctive conditionals leads to a logical treatment that is not implausible, though no attempt is made to develop any logic.

If nothing else, this thesis has brought to light a new range of temporal data that ought to be considered in any semantic treatment of tense, especially when this is to be combined with modals and conditionals. But more than this, the thesis has presented a satisfying way of accounting for the data.

We can also offer a rejoinder to Quine's exasperated observation quoted at the start of this chapter. The bias towards time in natural language occurs precisely because time never is far from our thoughts: it is the dimension in which thinking takes place, and this is reflected in the tense system. It is Quine's insistence that tense should only describe the way the world changes over time, and not the way that information changes, that makes natural language appear inelegant in his eyes.

7.2 Further Work

Ramifications for Tense and Aspect

What consequences does the doubly centred treatment of the past and present tenses have for other areas of temporal reference in natural language? For the most part, other areas of temporal reference are non-indexical — temporal adverbials, temporal anaphora, aspect and aktionsart, and so forth. One might therefore expect the revision to the indexical treatment of tense to have little impact on non-indexical mechanisms for temporal reference.

This conclusion is premature. It has already been seen how the temporal connective *when* exploits the verification time. There are other areas where having both an assertion and verification time available can make a difference.

For example, one obvious question to ask is: what is the relation between Reichenbach's notion of reference time and the verification time? Reichenbach claimed that the reference time offered a perspective point from which to view events being described, and that it was required to distinguish between the simple past and present perfect tenses. But it has never been made entirely clear why natural language should require an extra perspective point.

The verification time acts as some form of perspective point from which to view events being described — the point at which the descriptions are verified. Moreover, the verification time is a temporal index that has independent motivation in terms of the way that information changes over time. Perhaps it is Reichenbach's reference time.

The correspondence between reference and verification time cannot be as direct as this would suggest. Reichenbach treats the simple past tense as having the following configuration: `E,R---S`, where `E`, `R` and `S` represent the event, reference and speech times, and the line indicates temporal precedence. We treat the simple past as `E---V,A`, which corresponds more closely to Reichenbach's account of the present perfect once `V` is identified with `R` and `A` with `S`. However, something of the spirit of Reichenbach's analysis can be retained if the reference time is identified with the time at which an event is first verified or verifiable, rather than with the time that an assertion just happens to be verified. This would give `E,V'---A`, as required.

If this is so, then perhaps the following account can be given of the perfect tense: it makes reference to the time at which an assertion first becomes verified or verifiable in a given context. The present perfect says that there is something that makes a past event all of a sudden presently verifiable, while the past perfect says that at some point in the past an event became verifiable. The stated verifiability of the event might well be connected to the implication of current relevance that the perfect carries.

Perhaps the verification time can also be used to account for the progressive, something that Reichenbach never handled. Simply put, the progressive makes an assertion that is not yet verifiable. As it stands this makes the progressive behave a little like *will*, which explains why the present progressive is more naturally used to describe future events than the simple present. However, this falls foul of the imperfective paradox — that John is writing a thesis does not necessarily mean that John will write a thesis; he may never finish it. In this case, perhaps the progressive conveys unverified indirect information.

All of this is highly speculative, but is worth further consideration.

Quantification

One obvious omission from this thesis has been any treatment of quantification. We need to consider both quantified noun phrases and quantificational adverbials like *usually, always, sometimes, never.*

With regard to quantified noun phrases, it is interesting to note that first order intuitionistic logic has in-state and out-of-state quantifiers, i.e. $\exists$ and $\forall$ respectively. This potentially introduces an interesting temporal dimension to certain kinds of quantified noun phrase. A sentence like *all men are beasts* expresses a generalisation holding across past, present and future information states: it legislates for all men you are ever likely to meet. By contrast, *each man is a beast* seems more naturally to refer to a certain set of men already encountered within a given information state. Just as the intuitionistic negation, and out-of-state connective, spawned an in-state version ($\sim$), perhaps the same holds for universal quantifiers. This might go some of the way towards explaining the differences between the determiners *each*, *every* and *all.*

Quantificational adverbs can usefully be seen as specifying the quantification over minimal information extensions employed by conditionals. Up to now, I have assumed that the conditional quantifies over *all* minimal extensions supporting the conditional antecedent. But when a habitual conditional is combined with a quantificational adverbial like *sometimes, often* or *usually*, it is more natural to construe it as describing what happens in some or most minimal extensions, and not necessarily all. Quantificational adverbs should be regarded as implicitly conditional in the same way that hypothetical modals are. When they occur in non-conditional sentences context furnishes as implicit antecedent, but when they occur in conditional sentences the conditional itself furnishes the antecedent.

There is nothing especially novel about this. Lewis's (1975) account of adverbs of quantification is roughly similar, van Benthem (1984) shows how conditionals can be treated as generalised quantifiers, and Slaney (1991) exploits the trick of deriving implicational connectives from quantifiers (e.g. $p \rightarrow q =_{df} (\forall x \mid p)q$, where x does not occur free in p or q). What is problematic, however, is the fact that this variation of the conditional's quantification over information states should apply only to habitual conditionals. Specific conditionals must always quantify over all possible information extensions.

This points to further differences between specific and habitual conditionals that have not been covered in this thesis. Perhaps one needs to say that specific conditionals are invariably adverbial, and that adverbial conditionals only permit a universal quantifier over information states.

7.3 Computational Application: A Planner Interface

Although this thesis is concerned with the formal semantics of English, it grew out of some very real problems that arose when trying to implement a natural language front end to a planning system (Crabtree *et al* 1990; Crouch and Pulman, 1993). An interface to a planning system provides a rich domain for asking temporal, modal and conditional questions, such as

> When does Smith go to Ipswich?
> Could Jones go to London?
> If Smith repaired fault3, who could repair fault4?
> Can Smith repair digital equipment?

The system developed consisted of three components: a natural language front end, which parsed and semantically interpreted English sentences, and then translated the result into an expression in a specially designed plan query language (PQL); a plan query language evaluator, an inference engine that answered queries by consulting a plan or trying to generate further plans; and the planner itself.

There was tolerable success in dealing with modals and conditionals in the plan query language evaluator (Moffat and Ritchie 1990) — essentially, the space of plans generated was used to provide a set of accessible 'possible worlds' for analysing modality. However, this relied on the temporal relations expressed by modal and conditional sentences being fixed. That is, a query like *Could Jones go to London?* had to lead a PQL expression of the form $\exists t.\, t \geq now \wedge \mathrm{may}(go(smith, london, t))$ where the temporal references have to be stated explicitly. Fixing temporal references was the job of the front end. It rapidly became apparent that finding a fully general way of fixing the temporal references was a non-trivial task. Hence the subject matter of this thesis.

In the context of an interface to a planner, it was possible to take some short-cuts. As noted previously, present tense conditionals about plans allow for a much freer order between antecedent and consequent than do formally similar conditionals that are not about plans. But one can envisage other, non-planning applications where time, modality and conditionality feature prominently: Isard's (1974) noughts-and-crosses program is an example much ahead of its time; a system for describing and reasoning about mechanical devices like thermostats would be another possibility. In these cases, it is necessary to have a general account of the temporal relations expressed by modals and conditionals, which would include an account of the difference between the past and present tense bimetallic strip conditionals.

The treatment of modals and conditionals outlined in this thesis offers a satisfying answer to a problem inherent in Moffat and Ritchie's treatment of modality. They formalise the domain in which a planner operates in terms of a stratified set of theories: a theory about the way the world works in general (axioms about the linearity of time, etc), a theory about the actions available in the planning domain (what their pre- and post-conditions are), a theory describing the task to be carried out (initial conditions, goal conditions), a theory describing a non-linear plan for the task, and a theory describing linearisation of the non-linear plan. Modality is treated in terms of what is consistent with or entailed by by the theories up to some level. For example, the performance of a certain action is necessary in a given plan if the occurrence of that action is entailed by world knowledge, action definitions and the task specification. The performance of a certain action is possible if it is consistent with these theories taken jointly.

Automatic planning systems only implement a heuristically driven subset of the entailment and consistency relations implicit in this formalisation. This means that there is bound to be a mismatch between what a given planner does and what the formalisation says it does. In practice this does not matter; in a planner interface the planner itself will be used to calculate entailment and consistency relations. But the formal mismatch in unwelcome.

This can be rectified if the hierarchy of theories is seen as an ordered set of information states, with the planner acting as a device for creating minimal extensions of task specifications. It just so happens that these minimal extensions are complete plans, but the account of modals and conditionals makes no assumptions about how small these minimal extensions have to be. Once this move is made, the treatment of modals and conditionals in this thesis can be used to give a faithful formalisation of the modalities and conditionalities engendered by the planning system. Moreover, it can be cleanly integrated with a linguistically satisfactory account of modals and conditionals.

References

Adams, E., (1975), *The Logic of Conditionals.* D. Reidel, Dordrecht.

Alshawi, H. (ed), (1992), *The Core Language Engine.* MIT Press, MA.

Alshawi, H., D. Carter, R. Crouch, S. Pulman, M. Rayner and A. Smith, (1992) *CLARE: A Contextual Reasoning and Cooperative Response Framework for the Core Language Engine (Final Report)*, SRI International, Cambridge Computer Science Research Centre, Cambridge, UK.

Alshawi, H. and R. Crouch, (1992), 'Monotonic Semantic Interpretation', *Proceedings of the 30th Annual Meeting of the Association for Computational Linguistics*, 33–39.

Bennett, J., (1988) 'A Farewell to the Phlogiston Theory of Conditionals', *Mind* **xcvii**, 509–527.

van Benthem, J. F. A. K., (1984), 'Foundations of Conditional Logic', *Journal of Philosophical Logic* **13**, 303–349.

van Benthem, J. F. A. K., (1991), *The Logic of Time* (2nd edition). D. Reidel, Dordrecht.

Bertinetto, P. M., (1985), 'Intrinsic and Extrinsic Temporal Reference: On Restricting the Notion of "Reference Time"', in V. Lo Cascio and C. Vet (eds), *Temporal Structure in Sentence and Discourse*, Foris, Dordrecht.

Božić, M. and K. Došen, (1984), 'Modals for Normal Intuitionistic Modal Logics', *Studia Logica*, XLIII, 3, 216–245.

Chellas, B., (1980) *Modal Logic: an Introduction*, Cambridge University Press, Cambridge.

Coates, J., (1983), *The Semantics of Modal Auxiliaries*, Croom Helm, London.

Comrie, B., (1986), *Tense*, Cambridge University Press, Cambridge.

Comrie, B., (1986a), 'Conditionals: a Typology', in E. C. Traugott (ed), *On Conditionals*, Cambridge University Press, Cambridge.

Cooper, R. (1983), *Quantification and Syntactic Theory.* D. Reidel, Dordrecht.

Crabtree, B., R. Crouch, D. Moffat, N. Pirie, S. Pulman, G. Ritchie and A. Tate, (1990), 'Interacting with an Intelligent Planning System Using English Sentences', *First International Conference on Expert Planning Systems*, IEE Conference Publication Number 322, 169–174.

Crouch, R. and S. Pulman, (1993), 'Time and Modality in a Natural Language Interface to a Planning System', *Artificial Intelligence* **63**, 265–304.

van Dalen, D., (1986), 'Intuitionistic Logic', in D. Gabbay and F. Guenthner (eds), *Handbook of Philosophical Logics, Vol III*, D. Reidel, Dordrecht.

Dowty, D. R., (1979), *Word Meaning and Montague Grammar*, D. Reidel, Dordrecht.

Dowty, D. R., (1982), 'Tenses, Time Adverbs and Compositional Semantic Theory', *Linguistics and Philosophy* **5**, 23–55.

Dowty, D. R., (1986), 'The Effects of Aspectual Class on the Temporal Structure of Discourse: Semantics or Pragmatics?', *Linguistics and Philosophy* **9**, 37–61.

Dowty, D. R., R. E. Wall and S. Peters, (1981), *Introduction to Montague Semantics*, D. Reidel, Dordrecht.

Dudman, V. H., (1983), 'Tense and Time in English Verb Clusters of the Primary Pattern', *Australian Journal of Linguistics* **3**, 25–44.

Dudman, V. H., (1984a) 'Conditional Interpretations of "If"-sentences', *Australian Journal of Linguistics* **4**, 143–204.

Dudman, V. H., (1984b), 'Parsing 'If'-Sentences', *Analysis* **44**, 145–153.

Dudman, V. H., (1985), 'Towards a Theory of Predication for English', *Australian Journal of Linguistics* **5**, 143–196.

Dudman, V. H., (1991), 'Interpretations of 'If'-Sentences', in F. Jackson (ed), *Conditionals*, Oxford Readings in Philosophy, Oxford University Press, Oxford.

Dummett, M. A. E., (1977), *Elements of Intuitionism*, Oxford University Press, Oxford.

Enç, M., (1987), 'Anchoring Conditions for Tense', *Linguistic Inquiry* **18**, 633–657.

Evans, G., (1982), *The Varieties of Reference*, Oxford University Press, Oxford.

Fenstad, J. E., P.-K. Halvorsen, T. Langholm and J. F. A. K. van Benthem, (1987), *Situations, Language and Logic*, D. Reidel, Dordrecht.

Fitting, M., (1983), *Proof Methods for Modal and Intuitionistic Logics*, D. Reidel, Dordrecht.

van Fraassen, B. C., (1981), 'A Temporal Framework for Conditionals and Chance', in W. L. Harper, R. Stalnaker and G. Pearce (eds) *Ifs*, D. Reidel, Dordrecht.

Gabbay, D. M., (1981), *Semantical Investigations in Heyting's Intuitionistic Logic*, D. Reidel, Dordrecht.

Galton, A., (1984), *The Logic of Aspect*, Clarendon Press, Oxford.

Galton, A., (1987), 'Temporal Logic and Computer Science: an Overview', in A. Galton (ed), *Temporal Logics and Their Applications*, Academic Press, New York.

Gärdenfors, P., (1988), *Knowledge in Flux*, MIT Press, MA.

Ginsberg, M. L., (1985), 'Counterfactuals', *Proceedings of the Ninth International Joint Conference on Artificial Intelligence*, 80–86.

Grice, H. P., (1976), 'Logic and Conversation' in P. Cole and J. L. Morgan (eds) *Syntax and Semantics, vol. 3, Speech Acts*, Academic Press. Reprinted in F. Jackson (ed) *Conditionals*, 155–175, Oxford Readings in Philosophy, Oxford University Press, Oxford.

Groenendijk, J. and M. Stokhof, (1991), 'Dynamic Predicate Logic', *Linguistics and Philosophy* **14**, 39–100.

Haegeman, L. and H. Wekker, (1984), 'The Syntax and Interpretation of Futurate Conditionals in English', *Journal of Linguistics* **20**, 45–55.

Harper, M. P. and E. Charniak, (1986), 'Time and Tense in English', *Proceedings 24th Annual Meeting of the Association for Computational Linguists*, 3–9.

Higginbotham, J., (1985), 'On Semantics', *Linguistic Inquiry* **16**:4. Reprinted in E. LePore (ed), *New Directions in Semantics*, 1–54, Academic Press, 1987.

Hinrichs, E. W., (1986), 'Temporal Anaphora in Discourses of English', *Linguistics and Philosophy* **9**, 63–82.

Hinrichs, E. W., (1988), 'Tense, Quantifiers and Contexts', *Computational Linguistics* **14:2**, 3–14.

Hornstein, N., (1991) *As Time Goes By: Tense and Universal Grammar*, MIT Press, MA.

Isard, S. D., (1974), 'What Would You Have Done If...?', *Theoretical Linguistics* **1**, 233–255.

Jackson, F., (1979), 'On Assertion and Indicative Conditionals', *Philosophical Review* **88**, 565-589.

Kamp, H. and C. and Rohrer, (1983), 'Tense in Texts', in R. Bäuerle, C. Schwarze and A. von Stechow (eds), *Meaning, Use and Interpretation of Language*, de Gruyter, Berlin.

Krifka, M., (1991), 'Thematic Relations as Links between Nominal Reference and Temporal Constitution', in I. Sag and A. Sabolsci (eds), *Lexical Matters*, Chicago University Press, Chicago.

Kratzer, A., (1977), 'What 'Must' and 'Can' Must and Can Mean', *Linguistics and Philosophy*, **1**, 337–355.

Kratzer, A., (1979), 'Conditional Necessity and Possibility', in R. Bäuerle, U. Egli and A. von Stechow (eds), *Semantics from Different Points of View*, Springer, Berlin.

Kratzer, A., (1981a), 'Partition and Revision: The Semantics of Counterfactuals', *Journal of Philosophical Logic* **10**, 201–216.

Kratzer, A., (1981b), 'The Notional Category of Modality', in H.-J. Eikmayer and H. Rieser (eds), *Words, Worlds and Contexts: New Approaches in Word Semantics*, de Gruyter, Berlin.

Kratzer, A., (1989) 'An Investigation of the Lumps of Thought', *Linguistics and Philosophy*, **12**, 607–653.

von Kutschera, F., (1974), 'Indicative Conditionals', *Theoretical Linguistics* **1**, 257–288.

Landman, F., (1984), 'Data Semantics for Attitude Reports', *Logique et Analyse* **106**. Reprinted in Landman (1986), 157–181.

Landman, F., (1985), 'Data Semantics: an epistemic theory of partial objects', in Landman (1986), 1–96.

Landman, F., (1986), *Towards a Theory of Information*, Foris, Dordrecht.

Langholm, T., (1987), *Truth, Partiality and Persistence*, CSLI Lecture Notes 15, Stanford, CA.

Lascarides, A., (1988), *A Formal Semantic Analysis of the Progressive*, PhD Dissertation, University of Edinburgh.

Lascarides, A. and N. Asher, (1991), 'Discourse Relations and Common Sense Entailment', Research Paper HCRC/RP-16, Human Communication Research Centre, University of Edinburgh.

Lascarides, A., N. Asher, J. Oberlander, (1992), 'Inferring Discourse Relations in Context', *Proceedings of the 30th Annual Meeting of the Association for Computational Linguistics*, 1–8.

Lewin, I., (1990), 'A Quantifier Scoping Algorithm Without a Free Variable Constraint', *Proceedings 13th International Conference on Computational Linguistics, Vol. 3*, 190–194.

Lewis, D., (1973), *Counterfactuals*, Blackwell, Oxford.

Lewis, D., (1975), 'Adverbs of Quantification', in E. L. Keenan (ed), *Formal Semantics of Natural Language*, Cambridge University Press, Cambridge.

Lewis, D., (1981), 'Ordering Semantics and Premise Semantics for Counterfactual', *Journal of Philosophical Logic* **10**, 217–234.

Lyons, J., (1977), *Semantics, vol II*, Cambridge University Press, Cambridge.

MacArthur, R. P., (1976), *Tense Logic*, D. Reidel, Dordrecht.

McCoard, R. W., (1978), *The English Perfect: Tense-Choice and Pragmatic Inferences*, North-Holland Publishing Company, Amsterdam.

Moens M. and M. Steedman, (1988), 'Temporal Ontology and Temporal Reference', *Computational Linguistics* **14:2**, 15–28.

Morreau, M., (1992), 'Epistemic Semantics for Counterfactuals', *Journal of Philosophical Logic*, **21**, 33–62.

Moffat, D. C. and G. D. Ritchie, (1990) 'Modal Queries about Partially Ordered Plans', *Journal of Experimental and Theoretical Artificial Intelligence* **2**, 341–368.

Nerbonne, J., (1985), *German Temporal Semantics: Three Dimensional Tense Logic*, Garland Press, New York.

Nieuwint, P., (1986), 'Present and Future in Conditional Protases' *Linguistics* **24**, 371–392.

Nute, D., (1984), 'Conditional Logic', in D. Gabbay and F. Guenthner (eds), *Handbook of Philosophical Logic, Vol. II*, D. Reidel, Dordrecht.

Palmer, F. R., (1986), *The English Verb* (2nd edition), Longman, London.

Palmer, F. R., (1990), *Modality and the English Modals* (2nd edition), Longman, London.

Partee, B. H., (1973), 'Some Structural Analogies between Tenses and Pronouns in English', *Journal of Philosophy* **70**, 601–609.

Partee, B. H., (1984), 'Nominal and Temporal Anaphora', *Linguistics and Philosophy* **7**, 243–286.

Pereira, F. C. N., (1990), 'Categorial Semantics and Scoping', *Computational Linguistics* **16**, 1–9.

Perkins, M. R., (1983), *Modal Expressions in English*, Frances Pinter, London.

Pitts, A. M., (1993), 'Semantics of Programming Languages', Lecture Notes, University of Cambridge, Computer Laboratory.

Plotkin, G. and C. Stirling, (1986), 'A Framework for Intuitionistic Modal Logic', in J. Y. Halpern (ed), *Theoretical Aspects of Reasoning about Knowledge*, Morgan Kaufmann, CA.

Pollard, C. and I. Sag, (1987), *Information-Based Syntax and Semantics*, CSLI Lecture Notes 13, Stanford, CA.

Poznanski, V., (1990), *A Relevance-Based Utterance Processing System*, PhD Thesis, University of Cambridge Computer Laboratory.

Prior, A. N., (1967), *Past, Present and Future*, Oxford University Press, Oxford.

Quine, W. V.O., (1960), *Word and Object*, MIT Press, MA.

Quirk, R., S. Greenbaum, G. Leech and J. Svartvik, (1985), *A Comprehensive Grammar of the English Language*, Longman, London.

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

Reichgelt, H., (1986), 'Semantics for Reified Temporal Logic', Technical Report DAI 125, Dept. Artificial Intelligence, Edinburgh University.

Richards, B., (1987), 'Subordinate Tense' in J. Oberlander (ed), *Temporal Reference and Quantification: an IQ Perspective*, Edinburgh Working Papers in Cognitive Science, Volume 2.

Rigter, B., (1985), 'Focus Matters', in V. Lo Cascio and C. Vet (eds), *Temporal Structure in Sentence and Discourse*, Foris, Dordrecht.

Ritchie, G. D., (1979), 'Temporal Clauses in English', *Theoretical Linguistics* **6**, 87–115.

Roberts, C., (1989), 'Modal Subordination and Pronominal Anaphora in Discourse', *Linguistics and Philosophy* **12**, 683–721.

Schubert, L. K. and F. J. Pelletier, (1989), 'Generically Speaking, or, Using Discourse Representation Theory to Interpret Generics', in G. Chierchia, B. H. Partee and R. Turner (eds) *Properties, Types and Meaning, II*, D. Reidel, Dordrecht.

Shoham, Y., (1986), 'Reified Temporal Logics: Semantical and Ontological Considerations', *Proceedings of the 7th European Conference on Artificial Intelligence*, 390-397.

Slaney, J., (1991), 'The Implications of Paraconsistency', *Proceedings of the 12th International Joint Conference on Artificial Intelligence*, 1052–1057.

Stalnaker, R., (1968), 'A Theory of Conditionals', *Studies in Logical Theory, American Philosophical Quarterly* Monograph 2, 98–112, Blackwell, Oxford. Reprinted in E. Sosa (ed) *Causation and Conditionals*, Oxford University Press, 1975; and in F. Jackson (ed), *Conditionals* Oxford University Press, 1991.

Stalnaker, R., (1975), 'Indicative Conditionals', *Philosophia* **5**, 269–286.

Stalnaker, R., (1981), 'A Defense of Conditional Excluded Middle', in W. L. Harper, R. Stalnaker and G. Pearce (eds) *Ifs*, D. Reidel, Dordrecht.

Stump, G. T., (1985), *The Semantic Variability of Absolute Constructions*, D. Reidel, Dordrecht.

Taylor, B., (1977), 'Tense and Continuity', *Linguistics and Philosophy* **1**, 199-220.

Thomason, R. H., (1981), 'Deontic Logic as Founded on Tense Logic', in R. Hilpinen (ed), *New Studies in Deontic Logic*, D. Reidel, Dordrecht.

Thomason, R. H., (1984), 'Combinations of Tense and Modality', in D. Gabbay and F. Guenthner (eds), *Handbook of Philosophical Logic, Vol. II*, D. Reidel, Dordrecht.

Thomason, R. H. and A. Gupta, (1981), 'A Theory of Conditionals in the Context of Branching Time', in W. L. Harper, R. Stalnaker and G. Pearce (eds) *Ifs*, D. Reidel, Dordrecht.

Tichý, P., (1976), 'A Counterexample to the Stalnaker-Lewis Analysis of Counterfactuals', *Philosophical Studies* **29**.

Tichý, P., (1985), 'Do We Need Interval Semantics', *Linguistics and Philosophy* **8**, 263–282.

Veltman, F., (1976), 'Prejudices, Presuppositions and the Theory of Counterfactuals' in J. Groenendijk and M. Stokhof (eds) *Amsterdam Papers in Formal Grammar, vol 1.*

Veltman, F., (1984), 'Data Semantics', in J. Groenendijk, T. M. V. Janssen and M. Stokhof (eds), *Truth, Interpretation and Information*, Foris, Dordrecht.

Veltman, F., (1985), *Logics for Conditionals*, PhD Dissertation, University of Amsterdam.

Veltman, F., (1986), 'Data Semantics and the Pragmatics of Indicative Conditionals' in E. C. Traugott (ed), *On Conditionals*, Cambridge University Press, Cambridge.

Veltman, F., (1990), 'Defaults in Update Semantics I', in H. Kamp (ed) *Conditionals, Defaults and Belief Revision*, DYANA Deliverable R2.5.A.

Vendler, Z., (1967), *Linguistics in Philosophy*, Cornell University Press, New York.

Verkuyl, H. J., (1989), 'Aspectual Class and Aspectual Composition' *Linguistics and Philosophy* **12** 39–94.

Wekker, H. Chr., (1976), *The Expression of Future Time in Contemporary British English*, North Holland, Amsterdam.

Wijesekera, D., (1990), 'Constructive Modal Logics I', *Annals of Pure and Applied Logic*, **50**, 271–301.

Appendix A

A Grammar for Conditionals

Background

This appendix presents a unification based grammar for conditional sentences. The formalism is a variant of the CLE formalism (Alshawi 1992, Alshawi *at al* 1992). Rules consist of a rule identifier, the mother category, an arrow and a list of daughter categories. Each category has the form ⟨*SynCat*⟩/(⟨*SemCat*⟩+⟨*LF*⟩), where ⟨*SynCat*⟩ is a syntactic category consisting of a main category followed by a list of syntactic feature-value equations. ⟨*SemCat*⟩ is a similar semantic category, and ⟨*LF*⟩ is the logical form associated with the constituent.

Examples of the logical forms built up by this grammar are (uppercase letters are variable, `X^[..X..]` represents lambda abstraction):

1. John ate:

   ```
   [past,eat(j)]
   ```

2. 2. If John ate, Mary drank:

   ```
   form(if(indic,past),
        A^[A,[past,eat(j)],
             [past,drink(m)]],
        _)
   ```

3. If John eats, Mary drinks:

   Conjunctive

   ```
   form(if(indic,past),
        A^[A,[pres,eat(j)],
             [pres,drink(m)]],_)
   ```

Adverbial

```
[pres,form(if(indic,past),
          A^[A,[pres,eat(j)],
          drink(m)],_)
```

4. If John ate, Mary would drink:

   Present Subjunctive

```
form(if(subj,pres),
     G^[G,[pres,eat(j)],
          [pres,[will,drink(m)]]],_)
```

   Narrative

```
form(if(narr,past),
     G^[G,[pres,eat(j)],
          [pres,[will,drink(m)]]],_)
```

   Semi-Indicative

```
form(if(indic,past),
     A^[A,[past,eat(j)],
          form(if(subj,pres),
               C^[C,form(impl_ante,C^C,_),
                    [pres,[will,drink(m)]]],_)],
     _)
```

   Adverbial Present Subjunctive

```
[pres,
 form(if(subj,pres),
      G^[G,[pres,eat(j)],
           [will,drink(m)]],_)
```

5. If John ate, Mary would have drunk:

   Present Subjunctive

```
form(if(subj,pres),
     M^[M,[pres,eat(j)],
          [pres,[will,[perf,drink(m)]]]],_)
```

Semi-Indicative (Present Subjunctive)

```
form(if(indic,past),
     A^
     [A,[past,eat(j)],
      form(if(subj,past),
           C^[C,form(impl_ante,C^C,_),
                [pres,[will,drink(m)]]],
           _)],
     _),
```

Semi-Indicative (Past Subjunctive)

```
form(if(indic,past),
     G^
     [G,[past,eat(j)],
        form(if(subj,pres),
             I^[I,form(impl_ante,I^I,_),
                  [pres,[will,[perf,drink(m)]]]],
             _)],
     _)
```

6. If John had eaten, Mary would have drunk:

Past Subjunctive

```
form(if(subj,past),
     G^[G,[pres,eat(j)],
          [pres,[will,drink(m)]]],
     _)
```

Present Subjunctive (Perfective)

```
form(if(subj,pres),
     O^[O,[pres,[perf,eat(j)]],
          [pres,[will,[perf,drink(m)]]]],
     _)
```

Semi-Indicative (Past Subjunctive)

```
form(if(indic,past),
     A^[A,[past,[perf,eat(j)]],
          form(if(subj,past),
```

```
                C^[C,form(impl_ante,C^C,D),
                     [pres,[will,drink(m)]]],
                _)],
      _)
```

Some explanation of the logical forms is required. The `form` construction, used in the representation of conditional formulas and implicit antecedents to semi-indicative conditionals, takes three arguments: a form-category, a restriction, and an (uninstantiated) resolvent. As in Alshawi and Crouch 1992, the form-category expresses a relation between context, the restriction and the resolvent.

For conditional forms, the form-categories express the following relations (independent of context, in fact)

`if(subj,past) ===> P^Q^[P`$\Box\!\!\rightarrow_{pd}$`Q]`

`if(subj,pres) ===> P^Q^[P`$\Box\!\!\rightarrow_{pf}$`Q]`

`if(narr,past) ===> P^Q^[P`$\Box\!\!\rightarrow_{nf}$`Q]`

`if(indic,past) ===> P^Q^[P`$\rightarrow$`Q]`

`if(indic,pres) ===> P^Q^[P`$\rightarrow$`Q]` or `P^Q^[P`$\Rightarrow$`Q]`

That is, if the category is `if(subj,past)`, then the resolvent is `P^Q^[P`$\Box\!\!\rightarrow_{pd}$`Q]`. The meaning of the form as a whole is obtained by applying the restriction to the resolvent. So for instance

```
form(if(subj,past),
     G^[G,[pres,eat(j)],
          [pres,[will,drink(m)]]],
P^Q^[P□→pd Q])
```

is equivalent to

`[pres,eat(j)]` $\Box\!\!\rightarrow_{pd}$ `[pres,[will,drink(m)]]`

Note that for (present) indicative conditionals, two different resolutions are possible: $\rightarrow$ or $\Rightarrow$. Which constitutes the appropriate resolution depends on the tense and aspectual class of the consequent.

Implicit antecedent forms are resolved by picking up some proposition salient in context.

Grammar and Lexicon

```
% Features on categories --- syntactic+semantic:
%
%  cond(mood=_,tense=_,gaps=_)+
%  scond(consq=_,sgap=_)
%
%  s(mood=_,vform=_,tense=_,gaps=_)+
%  ss(sgap=_,implicit_cond=_)
%
%  vp(mood=_,vform=_,tense=_,gaps=_)+
%  svp(subjval,sgap,tense_op,implicit_cond)
%
%  aux(mood=_,vform=_,tense=_,takes_vform=_)+
%  saux(tense_op=_,implicit_cond=_,subcat=_)
%
%
% Semantic Features:
%
% subjval: place holder in verb phrase meaning for subject NP
% consq:   place holder in conditional ('if' + antecedent) for
%          consequent clause
% implicit_cond: place holder for implicit conditionality of
%          hypothetical modals.  Either gets filled in by subjunctive
%          conditionals, or by an expression requiring contextual
%          resolution
% tense_op: value of tense operator to be applied at sentential level,
%          passed up from verb phrase.
% sgap:    semantic gaps
% subcat:  subcategorisation frame



rule(syn0, % Sentences: must be indicative mood, with no gaps
           % and any implicit conditionals filled in

     sigma/(ssigma+S),

 -->
     [s(mood=indic,tense=_,gaps=[])
      /(ss(sgap=[],implicit_cond=n)+S)]).
```

```
rule(syn1,  % Ordinary conditional --- 'if A, ___'
            % mood and tense instantiated on form category for if

     cond(mood=M,tense=T,gaps=[])
     /(scond(consq=Conseq,sgap=[])
     + form(if(M,T),P^[P,Ante,Conseq],_)),

 -->
    [if/(sif+_),

     s(mood=M,tense=T,gaps=[])
     /(ss(sgap=[],implicit_cond=n) + Ante)
  ]).


rule(syn2, % Gapped conditional
           % gap set up to be filled in by
           % preposed adverbial conditionals

     cond(mood=M,tense=T,gaps=[cond(mood=M,tense=T,gaps=[])])
     /(scond(consq=Conseq,
             sgap=[(scond(consq=Conseq,sgap=[]) + Cond)]) + Cond),

 -->
     []).


rule(syn3,  % Conjunctive indicative conditional sentence

    s(mood=indic,tense=_,gaps=[])
    /(ss(sgap=[],implicit_cond=n) + Cond),

 -->
```

```
    [cond(mood=indic,tense=_,gaps=[])
     /(scond(consq=Conseq,sgap=[]) + Cond),

     s(mood=_,tense=_,gaps=[])
     /(ss(sgap=[],implicit_cond=n) + Conseq)
 ]).




rule(syn3a,  % Conjunctive non-indicative conditional sentence
             % 1. forces mood and tense agreement
             % 2. instantiates implicit conditional meaning of sentence

     s(mood=_indic,tense=_,gaps=[])
     /(ss(sgap=[],implicit_cond=n) + Cond),

 -->
     [cond(mood=M,\+(mood=indic),tense=T,gaps=[])
      /(scond(consq=_Conseq,sgap=[]) + Cond),

      s(mood=M,tense=T,gaps=[])
      /(ss(sgap=[],implicit_cond=y) + Cond)
  ]).



rule(syn4,  % Preposed indicative adverbial conditional sentence
            % Conditional fills in gap features

    s(mood=indic,tense=_,gaps=[])
    /(ss(sgap=[],implicit_cond=n) + Cond),

 -->
    [cond(mood=indic,tense=T,gaps=[])
     /(scond(consq=Conseq,sgap=[]) + Ante),

    s(mood=indic,tense=_,
                 gaps=[cond(mood=indic,tense=T,gaps=[])])
     /(ss(sgap=[(scond(consq=Conseq,sgap=[]) + Ante)],
```

```
        implicit_cond=n)
     + Cond)
  ]).


rule(syn4a,  % Preposed non-indicative adverbial conditional sentence

     s(mood=indic,tense=_,gaps=[])
     /(ss(sgap=[],implicit_cond=n) + Cond),

 -->
    [cond(mood=M,\+(mood=indic),tense=T,gaps=[])
     /(scond(consq=Conseq,sgap=[]) + Ante),

    s(mood=M,tense=pres,
           gaps=[cond(mood=M,tense=T,gaps=[])])
     /(ss(sgap=[(scond(consq=Conseq,sgap=[]) + Ante)],
          implicit_cond=y)
     + Cond)
  ]).


rule(syn5,  %Adverbial indicative conditional modification

     vp(mood=indic,vform=finite,tense=T,gaps=G)
     /(svp(subjval=NP,sgap=SG,tense_op=Op,implicit_cond=n) + Cond),

 -->
      [vp(mood=indic,vform=finite,tense=T,gaps=[])
       /(svp(subjval=NP,sgap=[],tense_op=Op,implicit_cond=n)
       + Conseq),

      cond(mood=indic,tense=_,gaps=G)
      /(scond(consq=Conseq,sgap=SG) + Cond)
  ]).


rule(syn5a,  %Adverbial non-indicative conditional modification
```

```
    vp(mood=_indic,vform=finite,tense=pres,gaps=G)
    /(svp(subjval=NP,sgap=SG,tense_op=Op,implicit_cond=n)
    + Cond),

 -->
   [vp(mood=subj,vform=finite,tense=pres,gaps=[])
    /(svp(subjval=NP,sgap=[],tense_op=Op,
          implicit_cond=Conseq^Cond)
    + Conseq),

    cond(mood=subj,tense=past,gaps=G)
    /(scond(consq=Conseq,sgap=SG) + Cond)
  ]).



rule(syn6,  % S -> NP VP, subjunctive, no implicit conditional
    s(mood=M,tense=T,gaps=G)
    /(ss(sgap=SG,implicit_cond=n) + [Op,S]),

 -->
    [np/(snp+NP),

     vp(mood=M,vform=finite,tense=T,gaps=G)
     /(svp(subjval=NP,sgap=SG,tense_op=Op,implicit_cond=n)
     + S)
  ]).


rule(syn6a,  % S -> NP VP, implicitly conditional
    s(mood=M,tense=T,gaps=G)
    /(ss(sgap=SG,implicit_cond=y) + Cond),

 -->
    [np/(snp+NP),

    vp(mood=M,vform=finite,tense=T,gaps=G)
     /(svp(subjval=NP,sgap=SG,tense_op=Op,
         implicit_cond=[Op,S]^Cond)
     + S)
  ]).
```

```
rule(syn6b,  % S -> NP VP, indicative
     s(mood=indic,tense=T,gaps=G)
     /(ss(sgap=SG,implicit_cond=n) + Cond),

 -->
     [np/(snp+NP),

     vp(mood=subj,vform=finite,tense=T,gaps=G)
     /(svp(subjval=NP,sgap=SG,tense_op=Op,
            implicit_cond=C^Cond)
     + S)
  ])
:-
 Cond = form(if(_,_),P^[P,form(impl_ante,P^P,_),C],_),
 C = [Op,S].



rule(syn7,  % VP -> Aux VP
     vp(mood=M,vform=VOut,tense=T,gaps=[])
     /(svp(subjval=NP,sgap=SG,tense_op=Op,implicit_cond=Cond)
     + Aux),

 -->
    [aux(mood=M,vform=VOut,tense=T,takes_vform=VIn)
     /(saux(tense_op=Op,implicit_cond=Cond,subcat=S) + Aux),

     vp(mood=M,vform=VIn,tense=_,gaps=[])
     /(svp(subjval=NP,sgap=SG,tense_op=_,implicit_cond=n) + S)
  ]).


% Lexicon:

lex(john,
    np/(snp+j)).
lex(mary,
    np/(snp+m)).
```

```
lex(if,
    if/(sif+if).
lex(ate,
    vp(mood=indic,vform=finite,tense=past,gaps=[])
     /(svp(subjval=NP,sgap=[],tense_op=past,implicit_cond=n) +
    eat(NP)).
lex(ate,
    vp(mood=subj,vform=finite,tense=pres,gaps=[])
     /(svp(subjval=NP,sgap=[],tense_op=pres,implicit_cond=n) +
    eat(NP)).
lex(ate,vp(mood=narr,vform=finite,tense=narr,gaps=[])
     /(svp(subjval=NP,sgap=[],tense_op=pres,implicit_cond=n) +
    eat(NP)).
lex(drank,
    vp(mood=indic,vform=finite,tense=past,gaps=[])
     /(svp(subjval=NP,sgap=[],tense_op=past,implicit_cond=n) +
    drink(NP)).
lex(drank,
    vp(mood=subj,vform=finite,tense=pres,gaps=[])
     /(svp(subjval=NP,sgap=[],tense_op=pres,implicit_cond=n) +
    drink(NP)).
lex(drank,vp(mood=narr,vform=finite,tense=narr,gaps=[])
     /(svp(subjval=NP,sgap=[],tense_op=pres,implicit_cond=n) +
    drink(NP)).
lex(eats,
    vp(mood=indic,vform=finite,tense=pres,gaps=[])
     /(svp(subjval=NP,sgap=[],tense_op=pres,implicit_cond=n) +
    eat(NP)).
lex(drinks,
    vp(mood=indic,vform=finite,tense=pres,gaps=[])
     /(svp(subjval=NP,sgap=[],tense_op=pres,implicit_cond=n) +
    drink(NP)).
lex(eat,
    vp(mood=_,vform=inf,tense=_,gaps=[])
     /(svp(subjval=NP,sgap=[],tense_op=_,implicit_cond=n) +
    eat(NP)).
lex(drink,
    vp(mood=_,vform=inf,tense=_,gaps=[])
     /(svp(subjval=NP,sgap=[],tense_op=_,implicit_cond=n) +
    drink(NP)).
lex(eaten,
```

```
    vp(mood=_,vform=en,tense=_,gaps=[])
     /(svp(subjval=NP,sgap=[],tense_op=_,implicit_cond=n) +
    eat(NP)).
lex(drunk,
    vp(mood=_,vform=en,tense=_,gaps=[])
     /(svp(subjval=NP,sgap=[],tense_op=_,implicit_cond=n) +
    drink(NP)).
lex(has,
    aux(mood=indic,vform=finite,tense=pres,takes_vform=en)
     /(saux(tense_op=pres,implicit_cond=n,subcat=S) +
    [perf,S]).
lex(had,
    aux(mood=indic,vform=finite,tense=past,takes_vform=en)
     /(saux(tense_op=past,implicit_cond=n,subcat=S) +
    [perf,S]).
lex(had,
    aux(mood=subj,vform=finite,tense=past,takes_vform=en)
     /(saux(tense_op=pres,implicit_cond=n,subcat=S) +
    S).
lex(had,
    aux(mood=subj,vform=finite,tense=pres,takes_vform=en)
     /(saux(tense_op=pres,implicit_cond=n,subcat=S) +
    [perf,S]).
lex(have,aux(mood=_,vform=inf,tense=_,takes_vform=en)
     /(saux(tense_op=_,implicit_cond=n,subcat=S) +
    [perf,S]).
lex(will,
    aux(mood=indic,vform=finite,tense=pres,takes_vform=inf)
     /(saux(tense_op=pres,implicit_cond=n,subcat=S) +
    [will,S]).
lex(would,
    aux(mood=subj,vform=finite,tense=pres,takes_vform=inf)
     /(saux(tense_op=pres,
           implicit_cond=Cns^form(if(subj,_T),P^[P,_A,Cns],_R),
           subcat=C) +
    [will,C]).
lex(would,
    aux(mood=narr,vform=finite,tense=narr,takes_vform=inf)
     /(saux(tense_op=pres,
           implicit_cond=Cns^form(if(narr,_T),P^[P,_A,Cns],_R),
           subcat=C) +
```

```
    [will,C]).
lex(would,
    aux(mood=subj,vform=finite,tense=past,takes_vform=have)
    /(saux(tense_op=pres,
           implicit_cond=Cns^form(if(subj,_T),P^[P,_A,Cns],_R),
           subcat=C) +
    [will,C]).
lex(have,aux(mood=subj,vform=have,tense=_,takes_vform=en)
     /(saux(tense_op=_,implicit_cond=n,subcat=S) +
    S).
```

Semantic Definitions

For the sake of reference, we collect together the main semantic definitions used in this thesis.

Atomic sentences:
$s, a, v, e \models p$ iff $V(s, v, p, e) = 1$

Def: $\wedge$
$s, a, v, e \models \phi \wedge \psi$ iff $s, a, v, e \models \phi$ and $s, a, v, e \models \psi$

Def: $\vee$
$s, a, v, e \models \phi \vee \psi$ iff $s, a, v, e \models \phi$ or $s, a, v, e \models \psi$

Def: past
$s, a, v, e \models \mathbf{past}(\phi)$ iff there is some $e' < a$ such that $s, a, v, e' \models \phi$

Def: pres
$s, a, v, e \models \mathbf{pres}(\phi)$ iff $\exists e', a'$ such that $e' \geq a$, $a \leq a' \leq e'$, and $s, a', v, e' \models \phi$

Def: $s \sqsubseteq_v^{\phi,a,e} s_1$
$s \sqsubseteq_v^{\phi,a,e} s_1$ iff
(a) $s_1, a, v, e \models \phi$, and
(b) There is no s' or v' such that
$s \sqsubseteq_{v'} s' \sqsubseteq_{v'} s_1$, and
$a \leq v' < v$, and
$s', a, v', e \models \phi$

Def: $\Rightarrow$
$s, a, v, e \models \phi \Rightarrow \psi$ iff
$\forall s_1, v_1$ such that $v_1 > a$ and $s \sqsubseteq_{v_1}^{\phi,a,e} s_1$,
there exists a $v_2 \geq v_1$ such that $s_1, v_1, v_2, e \models \psi$

Def: $\rightarrow$
$s, a, v, e \models \phi \rightarrow \psi$ iff
$\forall s_1, v_1$ such that $v_1 \geq a$ and $s \sqsubseteq_{v_1}^{\phi,a,e} s_1$,
there exists an s_2 such that $s_2 \approx s_1$ and $s_2, v_1, v_1, e \models \psi$

Def: $\neg$
$s, a, v, e \models \neg(\phi)$ iff $\forall s_1$ such that $s \sqsubseteq_{a}^{\phi,a,e} s_1$, $s_1, a, a, e \models \bot$

Def: $\sim$
$s, a, v, e \models \sim \phi$ iff $\forall v_1$ such that $v_1 \geq v$, $s, a, v_1, e \not\models \phi$

Def: H
$s, a, v, e \models H(\phi)$ iff $\forall s', v'$ such that $s' \sqsupseteq_{v'}^{\psi,e,e} s$, $\exists v'', e'$ such that $v'' \geq v'$, $v' \leq e' \leq v''$ and $s', v', v'', e' \models \phi$

Def: may
$s, a, v, e \models \mathbf{may}(\phi)$ iff
$\exists s', a', v', e'$ such that $e' \geq e, v' \geq e$ and $s' \sqsupseteq_{v'}^{\phi,a',e'} s$.

Def: must
$s, a, v, e \models \mathbf{must}(\phi)$ iff
$\forall s'$ if $s' \sqsupseteq_v s$, then $\exists s_0, a', v', e'$ such that $e' \geq e, v' \geq e$ and $s_0 \sqsupseteq_{v'}^{\phi,a',e'} s$ and $\exists s''$ such that $s'' \sqsupseteq_v s_0$ and $s'' \sqsupseteq_v s'$.

Def: will
$s, a, v, e \models \mathbf{will}(\phi)$ iff
$\exists s', a', v', e'$ such that $e' \geq e, v' \geq e$, $s' \sqsupseteq_{v'}^{\phi,a',e'} s$ and $s' \approx s$.

Def: $\Box\!\!\rightarrow_{pr}$
$[s, a, v, e] \vdash \phi \Box\!\!\rightarrow_{pr} \psi$ iff
for all s', t' such that $\langle s', t' \rangle \in f(s, \phi, a)$, $t' \geq a$ and $s', t', t', e \models \phi \rightarrow \psi$
(where f is a state selection function).

Def: $\Box\!\!\rightarrow_{pd}$
$[s, a, v, e] \vdash \phi \Box\!\!\rightarrow_{pd} \psi$ iff
for all s', t' such that $\langle s', t' \rangle \in f(s, \phi, a)$, $t' < a$ and $s', t', t', e \models \phi \rightarrow \psi$
(where f is a state selection function).

Def: $\Box\!\!\rightarrow_{nr}$
$[s, a, v, e] \vdash \phi \Box\!\!\rightarrow_{nr} \psi$ iff
there exists a $t' < a$ such that $s, t', t', e \models \phi \rightarrow \psi$

Appendix B

Minimal Extensions and Monotonicity

B.1 Properties of Minimal Information Extensions

The notation

$$s \sqsubseteq_v^{\phi,a,e} s_1$$

says that s_1 is a minimal extension of s_0 relative to a formula ϕ, an assertion time a and an event time e. The verification time v is the first time after a at which the assertion made by ϕ at a is verified in s_1. This notation is subject to the following definition:

- $s \sqsubseteq_v^{\phi,a} s_1$ iff
  (a) $s_1, a, v, e \models \phi$, and
  (b) There is no s_2 or v_2 such that
  $s \sqsubseteq_v s_2 \sqsubseteq_v s$, and
  $a \leq v_2 < v$, and
  $s_2, a, v_2, e \models \phi$

In this appendix, a number of properties of minimal extensions are discussed. First, it is observed that for minimal extensions to exist, both the ordering over information states and the ordering over verification times must be well founded. Then it is shown that if state s_2 extends a state s_0 and verifies some assertion ϕ, then s_2 extends some minimal extension of s_0 with respect to ϕ. As a direct result of this, minimal extensions can be shown to exist, just so long as some extension verifying the assertion exists.

Well Foundedness

For minimal information extensions to be exist, it is important that the temporal ordering $<$ is well founded on the indices in I, and that the informational ordering $\sqsubseteq_v$ is well founded on the states in S. A binary relation R is well founded just in case that for any two objects x and y such that xRy, there is no infinite descending chain of objects such that

$$x \ldots x_i R x_j \ldots x_k R y$$

The ordering of 'greater than' on the real numbers is an example of a non well founded relation, since for any two numbers, there is always an infinite chain of numbers occurring between them.

Without well foundedness on states and verification times, there will not be such things as minimal extensions of states. For any extension of a state, there will always be another one preceding it. In the same way, for a given real number, i, there is no other number, j, that is minimally greater than i: for any j greater than i, there will always be another number, k, that is less than j but still greater than i.

Existence of Minimal Extensions

A Simplified Case

Before establishing the existence of minimal extensions as defined above, it is helpful to do the same thing for a simpler case. Let us therefore forget about verification times, and the dependence of the information ordering on verification times. Instead, for two states we can just say that $s_1 \sqsubseteq s_2$ without further temporal qualification. Given a (temporally fixed) assertion p, we can then define minimal extension for the simpler case as follows

$s \sqsubseteq^p s_1$ iff
(a) $s \sqsubseteq s_1$ and $s \models p$, and
(b) $\neg(\exists s_2.\, s \sqsubseteq s_2 \sqsubset s_1 \wedge s_2 \models p)$

To prove the existence of minimal extensions, we need to show something like the following:

If $\exists s_2.\, s_2 \sqsubseteq s \wedge s_2 \models p$, then $\exists s_1.\, s \sqsubseteq^p s_1$

That is, provided that a state s has at least one extension supporting p, then it has a minimal extension supporting p.

To prove this, we will in fact establish a slightly stronger result:

Theorem:

$$\begin{array}{l} \forall s_2 \text{ if } s_2 \sqsupseteq s \ \wedge s_2 \models p, \\ \quad \text{then } \forall s_3 \text{ if } s_2 \sqsupseteq s_3 \sqsupseteq s \ \wedge s_3 \models p, \\ \qquad \text{then } \exists s_1 .\, s \sqsubseteq^p s_1 \ \wedge s_3 \sqsupseteq s_1 \end{array}$$

The proof proceeds by a well founded induction (Pitts 1993). A well founded induction is a bit like a normal inductive proof thrown into reverse. In a normal induction, one shows that if a property P holds of some object n, $P(n)$, then it also holds of $n+1$, $P(n+1)$. The base case is normally to show that $P(0)$. With a well founded induction one establishes that if $P(n)$, then $P(n-1)$, with the base case being $P(0)$ for some $0 < n$. Provided that the chain $n, n-1, \ldots, 0$ is well founded, we can be sure that starting at n we will eventually work our way back to the base case.

Taking the property P to be

$$\begin{array}{l} P(s_2) = \text{if } s_2 \sqsupseteq s, \\ \quad \text{then } \forall s_3 \text{ if } s_2 \sqsupseteq s_3 \sqsupseteq s \ \wedge s_3 \models p, \\ \qquad \text{then } \exists s_1 .\, s \sqsubseteq^p s_1 \ \wedge s_3 \sqsupseteq s_1 \end{array}$$

it is easy to show that if P holds of some state s', then it also holds of $s'' \sqsubseteq s'$. For, if $s' \not\sqsupseteq s$ then $s'' \not\sqsupseteq s$, and P is vacuously satisfied. If $s' \sqsupseteq s$, then $P(s')$ legislates for what happens to all states s'' such that $s' \sqsupseteq s'' \sqsupseteq s$. If s'' is such that $s'' \not\sqsupseteq s$, then $P(s'')$ again holds vacuously.

It only remains to establish the base case. Here we simply show that $P(s)$. If $\models p$, then the minimal extension, s_1, is s itself. If not, $P(s)$ is vacuously satisfied.

This proof establishes that

$$\begin{array}{l} \forall s_2 \text{ if } s_2 \sqsupseteq s, \\ \quad \text{then } \forall s_3 \text{ if } s_2 \sqsupseteq s_3 \sqsupseteq s \ \wedge s_3 \models p, \\ \qquad \text{then } \exists s_1 .\, s \sqsubseteq^p s_1 \ \wedge s_3 \sqsupseteq s_1 \end{array}$$

from which, by strengthening of the antecedent we get the theorem we want.

$$\begin{array}{l} \forall s_2 \text{ if } s_2 \sqsupseteq s \ \wedge s_2 \models p, \\ \quad \text{then } \forall s_3 \text{ if } s_2 \sqsupseteq s_3 \sqsupseteq s \ \wedge s_3 \models p, \\ \qquad \text{then } \exists s_1 .\, s \sqsubseteq^p s_1 \ \wedge s_3 \sqsupseteq s_1 \end{array}$$

Existence Proof From this result, it is easy to prove the existence of minimal extensions. Suppose that there exists an s_2 extending s and supporting p. Then there exists a minimal extension of s supporting p. It is quite possible, of course, that s is its own minimal extension — i.e. when s already supports p.

Extensions of Minimal Extensions The same result also establishes the following. Any state s_2 extending s and supporting p is in turn of a minimal extension to s supporting p.

Intuitive Reconstruction The well founded induction proof given above might at first sight seem like a sleight of hand. However, the intuitive pattern of the proof can easily be described.

Suppose you start off with an s_2 extending s such that s_2 supports p. Three cases are possible:
(a) There is no state intervening between s_2 and s. Then s_2 is the minimal extension of s (in fact, s_2 is s)
(b) There are intervening states, but none of them support p. Again s_2 is the minimal extension of s.
(c) There are intervening states supporting p. Pick one of these, and treat it as though it were the original s_2.
Cases (a) and (b) correspond to terminations of the induction. Case (c) causes us to descend a chain from s_2 to s. Since the chain is not infinitely descending, sooner or later we will run into a state with no states preceding it that support p (case b), or eventually we will reach state s itself.

The Real Thing

It is now time to carry out the same kind of proof for temporally varying information states. The proof is similar to the one above, but we have to deal with both states and times. We are aiming to prove the following theorem:

> *Theorem: All extensions extend minimal extensions:*
> If $s_2 \sqsupseteq_{v_2} s_0$, $v_2 \geq a$ and $s_2, v_2, a, e \models \phi$,
> then $\exists s_1, v_1$ such that $a \leq v_1 \leq v_2$, $s_2 \sqsupseteq_{v_2} s_1$, and $s_0 \sqsubseteq^{\phi,a,e}_{v_1} s_1$

Well Founded Induction We will choose as the inductive property, P, the following:

$$\begin{aligned} P(s_2, v_2) = {} & \text{if } s_2 \sqsupseteq_{v_2} s_0 \text{ and } v_2 \geq a, \\ & \quad \text{then } \forall s_3, v_3.\ \text{if } v_2 \leq v_3 \leq a,\ s_2 \sqsupseteq_{v_3} s_3 \sqsupseteq_{v_3} s_0 \text{ and } s_3, v_3, a, e \models \phi, \\ & \quad\quad \text{then } \exists s_1, v_1.\ v_3 \geq v_1,\ s_3 \sqsupseteq_{v_3} s_1 \text{ and } s_1 \sqsupseteq^{\phi,a,e}_{v_1} s_0 \end{aligned}$$

Once more, it is easy to show that if P applies to s_2 and v_2, it will apply to any s_3, v_3 preceding them. This is because the property P universally quantifies over preceding states and times. Should s_3 and v_3 precede s_2 and v_2 but not succeed s_0 and a, property P will be vacuously true of them, as before. For the base case, choose $v_2 = a$ and $s_2 \sqsubseteq_a s_0$, $s_0 \sqsubseteq_a s_2$.

The induction proves that

$$\begin{aligned}
&\text{if } s_2 \sqsupseteq_{v_2} s_0 \text{ and } v_2 \geq a,\\
&\text{then } \forall s_3, v_3.\ \text{if } v_2 \leq v_3 \leq a,\ s_2 \sqsupseteq_{v_3} s_3 \sqsupseteq_{v_3} s_0 \text{ and } s_3, v_3, a, e \models \phi,\\
&\qquad\qquad\text{then } \exists s_1, v_1.\ v_3 \geq v_1,\ s_3 \sqsupseteq_{v_3} s_1 \text{ and } s_1 \sqsupseteq_{v_1}^{\phi,a,e} s_0
\end{aligned}$$

Minimal Extension Theorem The minimal extension theorem follows directly from this inductively derived proof. Select some s_2 and v_2 such that $s_2 \sqsupseteq_{v_2} s_0$, $v_2 \geq a$ *and* $s_2, v_2, a, e \models \phi$. Strengthening the antecedent does not invalidate the result above, and so we have

$$\begin{aligned}
&\text{if } s_2 \sqsupseteq_{v_2} s_0,\ v_2 \geq a \text{ and } s_2, a, v_2, e \models \phi,\\
&\text{then } \forall s_3, v_3.\ \text{if } v_2 \leq v_3 \leq a,\ s_2 \sqsupseteq_{v_3} s_3 \sqsupseteq_{v_3} s_0 \text{ and } s_3, a, v_3, e \models \phi,\\
&\qquad\qquad\text{then } \exists s_1, v_1.\ v_3 \geq v_1,\ s_3 \sqsupseteq_{v_3} s_1 \text{ and } s_1 \sqsupseteq_{v_1}^{\phi,a,e} s_0
\end{aligned}$$

Since $v_2 \leq v_2 \leq a$ and $s_2 \sqsupseteq_{v_2} s_2 \sqsupseteq_{v_2} s_0$, by existential instantiation, we have

$$\begin{aligned}
&\text{If } s_2 \sqsupseteq_{v_2} s_0,\ v_2 \geq a \text{ and } s_2, a, v_2, e \models \phi,\\
&\text{then } \exists s_1, v_1 \text{ such that } a \leq v_1 \leq v_2,\ s_2 \sqsupseteq_{v_2} s_1, \text{ and } s_0 \sqsubseteq_{v_1}^{\phi,a,e} s_1
\end{aligned}$$

Existence of Minimal Extensions That minimal extensions exist follows as a trivial consequence of the above theorem.

B.2 Monotonicity of Conditionals

It can be shown that all conditional formulas satisfy monotonicity of verification. It can also be shown that, *provided the consequent satisfies monotonicity of information growth*, conditionals also satisfy monotonicity of information growth.

Monotonicity of Verification

Taking $\rightarrow$ first, note that the value of the initial verification time v does not play any role in the definition. It therefore follows that if a $\rightarrow$ conditional is verified at any time in an information state, it is also verified at all times afterwards in that state. Exactly the same argument goes through for $\Rightarrow$, since it is also independent of the initial value of the verification time.

This result should not be confused with what might be called 'monotonicity of utterance'. This would say that once an utterance of a sentence is supported in a given state at a certain time, utterances of the same sentence made at all times afterwards are also supported. Monotonicity of utterance fails, as it should, for sentences containing tenses. Monotonicity of verification merely says that once the assertional content of a given utterance is verified, it will continue to be verified. Later utterances of the same sentence may well have different assertional contents, however.

Monotonicity of Information Growth

Monotonicity of $\Rightarrow$

To show that the $\Rightarrow$ conditional satisfies monotonicity of information growth, we must establish the following

> *Monotonicity of $\Rightarrow$:*
> If $s, a, v, e \models \phi \Rightarrow \psi$ then for all $s_2 \sqsupseteq_v s$, $s_2, a, v, e \models \phi \Rightarrow \psi$
> (provided ψ is monotonic)

We will take this in a number of steps.

1. What does it mean for $s, a, v, e \models \phi \Rightarrow \psi$? According to the semantic definition of $\Rightarrow$ it is as follows

> $s, a, v, e \models \phi \Rightarrow \psi$ iff
> $\forall s_1, v_1$ if $s \sqsubseteq_{v_1}^{\phi,a,e} s_1$ then $\exists v_2$ such that $v_2 \geq v_1$ and $s_1, v_1, v_2, e \models \psi$

From this we also have that (a) $s_1 \sqsupseteq_{v_1} s$ and (b) $s_1, a, v_1, e \models \phi$.

It is also possible to show that v_1 must either be identical to a, or to the earliest verification time permitted for ϕ, depending on which is latter. Suppose that the assertion of ϕ is already verified in s. Then the earliest verification time, v_1 not less than a is a itself. Suppose that ϕ is not verified in s, but is ideally verifiable at some time before a. Then the earliest verification time v_1 not less than a is again a. Suppose that ϕ is not verified, and its earliest verification time, t_v, is after a. Then $v_1 = t_v$.

2. Now consider some arbitrary s_2 such that $s_2 \sqsupseteq_v s$. We need to show that $s_2, a, v, e \models \phi \Rightarrow \psi$, where

> $s_2, a, v, e \models \phi \Rightarrow \psi$ iff
> $\forall s_3, v_3$ if $s_2 \sqsubseteq_{v_3}^{\phi,a,e} s_3$ then $\exists v_4$ such that $v_4 \geq v_3$ and $s_3, v_3, v_4, e \models \psi$

Let us assume that there is such a minimal extension of s_2 (if not, the conditional will be vacuously true, and monotonicity follows immediately). From this we also have that (a) $s_3 \sqsupseteq_{v_3} s_2$ and (b) $s_3, a, v_3, e \models \phi$.

As before, we can show that v_3 must be either a or t_v. That is, $t_3 = t_1$.

3. The states s, s_2 and s_3 are related as follows

$$s \sqsubseteq_v s_2 \sqsubseteq_{v_3} s_3$$

By the constraint on the convergence of verification, we can therefore conclude that

$$\exists v_5 \text{ such that } v_5 \geq v,\ v_5 \geq v_3 \text{ and } s \sqsubseteq_{v_5} s_3$$

4. From (3) we can show, using the theorem stating that all extensions are extensions of minimal extensions, that

$$\exists s_1, v_1 \text{ such that } s \sqsubseteq_{v_1}^{\phi,a,e} s_1,\ s_3 \sqsupseteq_{v_5} s_1 \text{ and } v_1 \leq v_5$$

5. From (1) we have that $s_1, v_1, v_2, e \models \psi$. Assuming monotonicity of verification for ψ, we also have that $s_1, v_1, v_4, e \models \psi$ for some verification time $t_4 \geq t_5$. If we assume that ψ is monotonic, then given that $s_3 \sqsupseteq_{v_5} s_1$ and $v_1 = v_3$ we have

$$s_3, v_3, v_4, e \models \psi\ (v_4 \geq v_3)$$

This is what is required of the consequent in (2). So by assuming the conditional in (1) and the antecedent of the conditional in (2), we can show that the consequent of the conditional in (2) holds. This establishes monotonicity of information growth.

Monotonicity of $\rightarrow$

Showing that $\rightarrow$ is monotonic follows essentially the same pattern as for $\Rightarrow$. Steps (1)–(4) proceed essentially as before. At the end of step (4), we can draw the following conclusions:

a. $s_2 \sqsubseteq_a^{\phi,a,e} s_3$
b. $s_3 \sqsupseteq_{v_5} s_1$
c. $s \sqsubseteq_{v_1}^{\phi,a,e} s_1$
d. $v_3 = v_1$ e. $\exists s_1' \approx s_1$ such that $s_1', v_3, v_3, e \models \psi$

We need to show

e. $\exists s_3' \approx s_3$ such that $s_3', a, a, e \models \psi$

Given that s_3 extends s_1, and there is a state equivalent to s_1 that verifies ψ at a, it follows that there will be a similar state equivalent to s_3. Thus the consequent of the conditional at s_2 is verified as required, and $\rightarrow$ is monotonic.

B.3 Worked Example

This section spells out the support conditions for the conditional formula $\mathbf{past}(A) \rightarrow \mathbf{past}(C)$. Evaluating this formula relative to a state s and a time n (n for 'now') amounts to the following:

$$s, n \models \mathbf{past}(A) \rightarrow \mathbf{past}(C) \text{ iff } [s, n, n, n] \models \mathbf{past}(A) \rightarrow \mathbf{past}(C)$$
$$\text{iff}$$
$$\forall v_1 \geq n,\ \forall s_1 \sqsupseteq_{v_1}^{\mathbf{past}(A),n,n} s :\ \exists s_2 \approx s_1 :\ [s_2, v_1, v_1, n] \models \mathbf{past}(C)))$$

For s_1 to be a minimal extension of s with respect to $\mathbf{past}(A)$, the following must hold

$$s_1 \sqsupseteq_{v_1}^{\mathbf{past}(A),n,n} s \text{ iff}$$
$$s_1 \sqsupseteq_{v_1} s,\ [s_1, n, v_1, n] \models \mathbf{past}(A),\ v_1 \geq n \text{ and}$$
$$\neg(\exists v_0, s_0 :\ n \leq v_0 < v_1 \& s \sqsubseteq_{v_0} s_0 \sqsubseteq_{v_0} s_1 \& [s_0, n, v_0, n] \models \mathbf{past}(A))$$

For a past tense formula to hold, we have:

$$[s_1, n, v_1, n] \models \mathbf{past}(A) \text{ iff } \exists t_A < n :\ [s_1, n, v_1, t_A] \models A$$

Because of the existence of ideally verifying states, there will be an s_1 where $\mathbf{past}(A)$ is verified at the end-point of t_A. Therefore, the earliest v_1 not preceding n at which any minimally extending state s_1 verifies $\mathbf{past}(A)$ is n itself. So,

$$v_1 = n,$$
$$\text{therefore}$$
$$s_1 \sqsupseteq_{v_1}^{\mathbf{past}(A),n,n} s \text{ implies}$$
$$s_1 \sqsupseteq_n s \text{ and } [s_1, n, n, n] \models \mathbf{past}(A)$$

For the consequent, we therefore have

$$\exists s_2 \approx s_1 :\ [s_2, n, n, n] \models \mathbf{past}(C)$$
$$\text{which implies}$$
$$[s_2, n, n, t_C] \models C, \text{ where } t_C < n$$

In summary, (i) $t_A < n$, (ii) the antecedent is verifiable at n, and so (iii) $t_C < n$. Thus both antecedent and consequent events precede the time of utterance, but are not ordered relative to each other.

Appendix C

Verified and Unverified Assertion: Soundness and Completeness

This appendix shows that the proof theory for a logic of verified and unverified assertions presented in Chapter 5 is sound and complete with respect to the intended semantics.

C.1 Syntax and Semantics

We start by recalling the proof theory and semantic definitions.

Proof Theory

The proof theory is shown in Figure C.1. Stability of a formula is a syntactic notion defined as follows:

- If p is atomic, then p is stable.
- If ϕ and ψ are stable, then $\phi \wedge \phi$ and $\phi \vee \psi$ are stable.
- $\phi \rightarrow \psi$ is stable if ψ is stable. (Otherwise, it is semi-stable.)
- $\neg\phi$ is stable.
- If ϕ is stable, then $\sim\sim \phi$ is stable.
- Anything not classified as stable by the above is unstable.

Stable(Γ) denotes the set of stable sentences contained in a set of sentences Γ.

$$\wedge I\,\frac{\Gamma \vdash \phi;\ \ \Gamma \vdash \psi}{\Gamma \vdash \phi \wedge \psi} \qquad \wedge E\,\frac{\Gamma \vdash \phi \wedge \psi}{\Gamma \vdash \phi}$$

$$\vee I\,\frac{\Gamma \vdash \phi}{\Gamma \vdash \phi \vee \psi} \qquad \vee E\,\frac{\Gamma \vdash \phi \vee \psi;\ \ \Gamma, \phi \vdash \chi;\ \ \Gamma, \psi \vdash \chi}{\Gamma \vdash \chi}$$

$$\rightarrow I\,\frac{\mathit{Stable}(\Gamma), \phi \vdash \sim\sim \psi}{\Gamma \vdash \phi \rightarrow \psi} \qquad \rightarrow E\,\frac{\Gamma \vdash \phi;\ \ \Gamma \vdash \phi \rightarrow \psi}{\Gamma \vdash \sim\sim \psi}$$

$$\neg I\,\frac{\mathit{Stable}(\Gamma), \phi \vdash \sim\sim \bot}{\Gamma \vdash \neg\phi} \qquad \bot\,\frac{\Gamma \vdash \bot}{\Gamma \vdash \phi}$$

$$\sim I\,\frac{\Gamma, \phi \vdash \bot}{\Gamma \vdash \sim \phi} \qquad \sim Ax\,\frac{}{\Gamma \vdash \sim \phi \vee \sim\sim \phi}$$

$$\sim\rightarrow\,\frac{\Gamma \vdash \sim\sim \phi;\ \ \Gamma \vdash \phi \rightarrow \psi}{\Gamma \vdash \sim\sim \psi} \qquad \sim\sim\bot\,\frac{\Gamma \vdash \sim\sim \bot}{\Gamma \vdash \bot}$$

Figure C.1: Proof theory for logic of verified and unverified assertions

Information Models

Information models M are quintuples:

$$M = \langle S, \sqsubseteq_t, T, \leq, V \rangle$$

where S is a set of information states, s
$\sqsubseteq_t$ is a relation in $S \times S \times T$
and is transitive and reflexive over S for any t
T is a set of time instants, t
$\leq$ is a (linear) temporal order over T, and
V is a valuation function

Three constraints are imposed on information models:

Monotonicity of direct verification ('in-state' monotonicity):
For every state s and atomic sentence p
if $t_1 \leq t_2$ then $V(s, t_1, p) = 1 \Rightarrow V(s, t_2, p) = 1$

Monotonicity of information growth ('out-of-state' monotonicity):
If $s_1 \sqsubseteq_t s_2$ then for atomic sentences p
(a) $\{p \mid V(s_1, t, p) = 1\} \subseteq \{p \mid V(s_2, t, p) = 1\}$
(b) $\{p \mid \exists t : V(s_1, t, p) = 1\} \subseteq \{p \mid \exists t : V(s_2, t, p) = 1\}$

Convergence of Verification:
If $s_1 \sqsubseteq_{t_1} s_2 \sqsubseteq_{t_2} s_3$,
then there is a time t_3 such that $t_3 \geq t_1$, $t_3 \geq t_2$ and $\forall t_4 \geq t_3$, $s_1 \sqsubseteq_{t_4} s_3$

No Absurdity:
For no s or t is it the case that $V(s, t, \perp) = 1$

Semantics

The semantic definitions for the connectives are:

1. $s, t \Vdash p$ iff $V(s, t, p) = 1$ if p is atomic
2. $s, t \Vdash \phi \wedge \psi$ iff $s, t \Vdash \phi$ and $s, t \Vdash \psi$
3. $s, t \Vdash \phi \vee \psi$ iff $s, t \Vdash \phi$ or $s, t \Vdash \psi$
4. $s, t \Vdash \phi \rightarrow \psi$ iff $\forall t_1 \geq t, s_1 \sqsupseteq_{t_1}^{\phi,t} s : \exists t_2 \geq t_1$ such that $s_1, t_2 \Vdash \psi$
5. $s, t \Vdash \neg\phi$ iff $\forall t_1 \geq t, s_1 \sqsupseteq_{t_1}^{\phi,t} s : \exists t_2 \geq t_1$ such that $s_1, t_2 \Vdash \perp$
6. $s, t \Vdash \sim \phi$ iff $\forall t_1 \geq t : s, t_1 \nVdash \phi$

where minimal information extensions are defined as

- $s_1 \sqsupseteq_{t_1}^{\phi,t} s$ iff
  a) $s_1 \sqsupseteq_{t_1} s$
  b) $s_1, t_1 \Vdash \phi$, and
  c) $\nexists t_2, s_2$ such that $t \leq t_2 < t_1$, $s \sqsupseteq_{t_2} s_2 \sqsupseteq_{t_2} s_1$, and $s_2, t_2 \Vdash \phi$

Monotonicity of Stable Formulas

The out-of-state monotonicity of stable formulas is established by induction on the complexity of the formula.

Out-of-State Monotonicity:
If $s, t \Vdash \phi$ then $\forall s_1 \sqsupseteq_t s, t_1 \geq t : s_1, t_1 \Vdash \phi$

Base case: If ϕ is atomic, it is stable and by the monotonicity of information growth it also obeys out-of-state monotonicity.

Conjunction and disjunction: Let ϕ and ψ be stable, and by the induction hypothesis, monotonic. By definition of stability, $\phi \wedge \psi$ and $\phi \vee \psi$ are stable, and by semantics of $\wedge$ and $\vee$ they are monotonic.

Implication: Let ψ be stable, and by induction hypothesis monotonic. Then $\phi \rightarrow \psi$ is stable. By a simple variant of the proof of the monotonicity of conditionals given in Appendix B, $\phi \rightarrow \psi$ is monotonic.

Out-of-state negation: Define $\neg\phi$ as $\phi \rightarrow \perp$. $\perp$ is monotonic, and so $\phi \rightarrow \perp$ is to. Alternatively, observe that $\neg\phi$ means that there are no minimal ϕ-extensions. The universal quantification over extending states ensures monotonicity.

Double in-state negation: For $\sim\sim \phi$, assume that ϕ is stable, and by induction, monotonic. If the double negation holds in s at t, this means that ϕ holds in s at some time later than t. Since ϕ is monotonic, it will also continue to hold in all subsequent states and times. ϕ holding at some time in a state s is enough to ensure that $\sim\sim \phi$ holds at any time in that state. Thus if $\sim\sim \phi$ holds in s at t, it also holds in all subsequent states at all subsequent times.

As mentioned in Chapter 5, the converse result does not hold: some unstable formulas are monotonic. However, the soundness and completeness proofs below shows that this does no harm (see p. 245).

C.2 Soundness

To prove soundness we must show that

$\Gamma \vdash \phi \Rightarrow \Gamma \models \phi$

where $\Gamma \models \phi$ is a short hand for saying that
for any model M, state s and time t such that $M, s, t \Vdash \sigma$ for every $\sigma \in \Gamma$, $M, s, t \Vdash \phi$

and $\Gamma \vdash \phi$ means that ϕ is derivable from premises Γ in the inference system

The proof proceeds by induction on the complexity of the derivation $\Gamma \vdash \phi$. We will assume a fixed model M.

Atomic Derivations Soundness of atomic derivations, $\phi \vdash \phi$, is trivial. Any model of the premise is a model of the conclusion. This provides the base case for the induction hypothesis used below.

Conjunction Introduction Suppose we have $\Gamma \vdash \phi$ and $\Gamma \vdash \psi$. By the induction hypothesis, $\Gamma \models \phi$ and $\Gamma \models \psi$. Choose any state s and time t such that $s, t \Vdash \Gamma$, from which it follows that $s, t \Vdash \phi$ and $s, t \Vdash \psi$. By the semantic definition of conjunction, $s, t \Vdash \phi \wedge \psi$. Hence it is safe to assume that $\Gamma \vdash \phi \wedge \psi$.

Conjunction Elimination Trivial.

Disjunction Introduction Trivial.

Disjunction Elimination We have three induction hypotheses: $\Gamma \models \phi \vee \psi$, $\Gamma, \phi \models \chi$, and $\Gamma, \psi \models \chi$.

Let s and t be such that $s, t \Vdash \Gamma$. Hence $s, t \Vdash \phi \vee \psi$. Hence either $s, t \Vdash \phi$ or $s, t \Vdash \psi$. In the first case, $s, t \Vdash \Gamma, \phi$, and so by induction hypothesis, $s, t \Vdash \chi$. Similarly in the second case. Either way, $s, t \Vdash \chi$, and so $\Gamma \models \chi$. Hence it is safe to assume that $\Gamma \vdash \chi$.

In-State Negation Introduction Suppose $\Gamma, \phi \vdash \bot$. By the induction hypothesis, $\Gamma, \phi \models \bot$. Since $\bot$ is never verified in any state, we may therefore assume that for no state and time does $s, t \Vdash \Gamma, \phi$.

Suppose that $s, t \Vdash \Gamma$. By monotonicity of direct verification, for all $t_1 \geq t$, $s, t_1 \Vdash \Gamma$. Since no state and time force Γ and ϕ, it follows that $\forall t_1 \geq t$, $s, t_1 \nVdash \phi$. By the semantics of in-state negation, this means that $s, t \Vdash \sim \phi$. Hence $\Gamma \models \sim \phi$ and so it is safe to assume that $\Gamma \vdash \sim \phi$

In-State Negation Axiom This is not a rule of inference but an axiom. We must show that $s, t \Vdash \sim \phi \vee \sim\sim \phi$ for all s, t.

According to semantic definitions,
(1) $s, t \Vdash \sim \phi$ iff $\forall t_1 \geq t : s, t_1 \nVdash \phi$, and
(2) $s, t \Vdash \sim\sim \phi$ iff $\forall t_1 \geq t, \exists t_2 \geq t_1 : s, t_2 \Vdash \phi$
Since the temporal order is linear, (2) is equivalent to
(2a) $s, t \Vdash \sim\sim \phi$ iff $\exists t_2 \geq t : s, t_2 \Vdash \phi$
It can be seen that (1) and (2a) are complements of one another, and so it follows that either $s, t \Vdash \sim \phi$ or $s, t \Vdash \sim\sim \phi$, for any s and t. Hence by the semantic definition of disjunction, $s, t \Vdash \sim \phi \vee \sim\sim \phi$ for all s, t.

Note The in-state negation axiom is in fact characteristic of ordinary intuitionistic logic with linear frames. The logic **KC** (Gabbay 1981:66) adds $\vdash \neg\phi \vee \neg\neg\phi$ as a single extra axiom governing intuitionistic negation. **KC** is sound and complete with respect to intuitionistic frames where $\sqsubseteq$ is reflexive, transitive and linear. Since the temporal order $\leq$ is reflexive, transitive and linear, the in-state negation axiom is inevitable.

Implication Introduction Suppose that $Stable(\Gamma), \phi \vdash\sim\sim \psi$. By the induction hypothesis, $Stable(\Gamma), \phi \models\sim\sim \psi$. That is, if $s, t \Vdash Stable(\Gamma), \phi$, then $\exists t_1 \geq t : s, t_1 \Vdash \psi$.

Let $s_0, t_0 \Vdash \Gamma$. Let s and t be a state and time forming a minimal extension of s_0 and t_0 with respect to ϕ. By out-of-state monotonicity, $\forall t' \geq t_0, s' \sqsupseteq_{t'} s_0 : s', t' \Vdash Stable(\Gamma)$,and so $s, t \Vdash Stable(\Gamma)$. By the induction hypothesis $s, t \Vdash \phi$ and so $\exists t_1 \geq t : s, t_1 \Vdash \psi$.

These are the semantic conditions required for $\phi \rightarrow \psi$ to hold in s_0, t_0. So it is safe to assume that $\Gamma \vdash \phi \rightarrow \psi$.

Implication Elimination Suppose $\Gamma \vdash \phi$ and $\Gamma \vdash \phi \rightarrow \psi$. By the induction hypothesis, $\Gamma \models \phi$ and $\Gamma \models \phi \rightarrow \psi$. According to the semantic definition of implication, if $s, t \Vdash \phi$ and $s, t \Vdash \phi \rightarrow \psi$, then there is a $t_1 \geq t$ such that $s, t_1 \Vdash \psi$. From which we conclude that $s, t \Vdash\sim\sim \psi$. Hence, $\Gamma \models\sim\sim \psi$, and so it is safe to assume that $\Gamma \vdash\sim\sim \psi$.

Out-of-State Negation Introduction It is possible to define out-of-state negation as $\neg\phi =_{df} \phi \rightarrow \bot$, eliminating the need for a separate introduction rule. Its introduction rule can be shown to be sound analogously to implication introduction, noting that no state or time is ever such that $s, t \Vdash\sim\sim \bot$.

Absurdity Suppose $\Gamma \vdash \bot$, and so by induction hypothesis $\Gamma \models \bot$. Since there are no states s or times t such that $s, t \Vdash \bot$, one can safely assume that $\Gamma \models \phi$ for any ϕ.

Future Absurdity Suppose $\Gamma \vdash\sim\sim \bot$, and so by induction hypothesis $\Gamma \models\sim\sim \bot$. By the semantic of $\sim$ if $s, t \Vdash\sim\sim \bot$, then $\exists t_1 \geq t$ such that $s, t_1 \bot$. Since there are no states or time that force $\bot$, there can be no states of times that force $\sim\sim \bot$. So it is safe to conclude that $\Gamma \models \bot$ and so $\Gamma \vdash \bot$.

Future Implication Elimination Suppose $\Gamma \vdash\sim\sim \phi$ and $\Gamma \vdash \phi \rightarrow \psi$. By the induction hypothesis, $\Gamma \models\sim\sim \phi$ and $\Gamma \models \phi \rightarrow \psi$. According to the semantics of $\sim$, if $s, t \Vdash\sim\sim \phi$, then $\exists t_1 \geq t$ such that $s, t_1 \Vdash \phi$. By in-state monotonicity, if $s, t \Vdash \phi \rightarrow \psi$ then $s, t_1 \Vdash \phi \rightarrow \psi$ as well. By semantics of implication, $\exists t_2 \geq t_1$ such that $s, t_2 \Vdash \psi$. And so $\exists t_2 \geq t$ such that $s, t_2 \Vdash \psi$. Thus $s, t \Vdash\sim\sim \psi$. Hence $\Gamma \models\sim\sim \psi$, and so $\Gamma \vdash\sim\sim \psi$.

C.3 Completeness

The proof of completeness is a variation on the Henkin-style completeness proofs for intuitionistic used by Aczel, Fitting and Thomason (van Dalen 1984), and also by Veltman (1985) for Data Semantics. The method involves constructing information states that are represented as saturated sets of formulas. A number of definitions and proofs concerning saturated sets are therefore required before proving completeness.

Saturated Theories

Suppose that we have a language L containing a set of sentences Γ, such that

$$\Gamma \nvdash \phi$$

We can construct a *saturated extension* of Γ relative to (an enumeration of) the sentences in L and ϕ, $Sat(\Gamma, \phi, L)$, such that

(a) $Sat(\Gamma, \phi, L)$ is closed under $\vdash$

(b) $\Gamma \subseteq Sat(\Gamma, \phi, L)$

(c) If $\psi \vee \chi \in Sat(\Gamma, \phi, L)$, then either $\psi \in Sat(\Gamma, \phi, L)$ or $\chi \in Sat(\Gamma, \phi, L)$

(d) $\phi \notin Sat(\Gamma, \phi, L)$

The construction proceeds inductively as follows.

1. Stage 0:
   Let $\Delta_0 = \Gamma$

2. Stage $k+1$:
   Let $\psi_1 \vee \psi_2$ be the first disjunction in the enumeration of L that has not so far been dealt with, where $\Delta_k \vdash \psi_1 \vee \psi_2$.

   Then $\Delta_{k+1} = \Delta_k \cup \{\psi_1\}$ if $\Delta_k \cup \{\psi_1\} \nvdash \phi$
   $\Delta_k \cup \{\psi_2\}$ otherwise.

3. $Sat(\Gamma, \phi, L) = \bigcup_{k \geq 0} \Delta_k$

The definition and construction, as well as the proofs that (a)–(d) hold, are taken from van Dalen 1984. We start by establishing the disjunction property, (c).

(c) Let $\psi \vee \chi \in Sat(\Gamma, \phi, L)$, and let k be the least number such that $\Delta_k \vdash \psi \vee \chi$. Then $\Delta_i \vdash \psi \vee \chi$ for all $i \geq k$. $\psi \vee \chi$ cannot be treated before stage k, so suppose that it is treated at state $j \geq k$. Then either $\psi \in \Delta_{j+1}$ or $\chi \in \Delta_{j+1}$. And $\Delta_{j+1} \subseteq Sat(\Gamma, \phi, L)$

(a) $Sat(\Gamma, \phi, L)$ is closed under $\vdash$. For is $Sat(\Gamma, \phi, L) \vdash \psi$, then $Sat(\Gamma, \phi, L) \vdash \psi \vee \psi$. So by property (c), $\psi \in Sat(\Gamma, \phi, L)$.

(b) That $\Gamma \subseteq Sat(\Gamma, \phi, L)$ is evident from the construction.

(d) Clearly, $\Delta_0 \not\vdash \phi$. If $\Delta_k \not\vdash \phi$, then $\Delta_{k+1} \not\vdash \phi$. For let $\phi \vee \psi$ be the $k+1$th disjunction treated, i.e. $\Delta_k \vdash \phi \vee \psi$. Clearly, $\Delta_k \cup \{\phi\} \vdash \phi$, so ψ must be added. Moreover, $\Delta_k, \psi \not\vdash \phi$, for if it did and $\Delta_k \vdash \phi \vee \psi$, then it could not be that $\Delta_k \not\vdash \phi$.

Which saturated extension of Γ one gets is dependent on the order in which the sentences in L are enumerated. When constructing sets of saturated extensions for Γ, as we will do below, it is assumed that all possible enumerations are employed.

Saturated theories will form the state snapshots used to build canonical models used to establish completeness.

Statement of Completeness Problem

To prove completeness, we must show that

$$\Gamma \models \phi \Rightarrow \Gamma \vdash \phi$$

Put in contrapositive form, and expanding out the meaning of $\models$, this amounts to

$\Gamma \not\vdash \phi \Rightarrow \Gamma \not\models \phi$, i.e.

$\Gamma \not\vdash \phi \Rightarrow$ there is a model, state and time such that $M, s, t \Vdash \Gamma$ but $M, s, t \not\Vdash \phi$.

The general strategy for proving completeness is to construct a saturated extension relative to Γ, ϕ and some enumeration of the language. This forms the state and time that forces Γ. We then build up temporal extensions of the state, as well as new states plus their temporal extensions, all of which are also saturated theories. We show that for any such theory, Σ_s^t, (i.e. state s taken at time t), $s, t \Vdash \psi$ iff $\psi \in \Sigma_s^t$. Since Σ_0^0 is just a saturated extension of Γ relative to ϕ, we have $\Gamma \subseteq \Sigma_0^0$, but $\phi \notin \Sigma_0^0$. This establishes the contrapositive of completeness.

Constructing State Snapshots

Suppose that we have some saturated extension of a set of sentences Γ relative to ϕ and a given enumeration of L. Call it Σ^0, and take it as the initial snapshot of the state. Subsequent snapshots, Σ^i are constructed relative to the same enumeration of the language as follows:

- Σ^{i+1}: Let $\sim\sim \psi$ be the first double negation not so far dealt with, such that $\Sigma^i \vdash\sim\sim \psi$.
- Construct a saturated theory from $\Sigma^i \cup \{\psi\}$ as follows
  - Step 0:
    Let $\Delta_0 = \Sigma^i \cup \{\psi\}$
  - Step $k+1$:
    Let $\psi_1 \vee \psi_2$ be the first disjunction not so far dealt with such that $\Delta_k \vdash \psi_1 \vee \psi_2$.

    $$\begin{aligned}\Delta_{k+1} = {} & \Delta_k \cup \{\psi_1\} \text{ iff } \Sigma^i \vdash\sim\sim \psi_1 \text{ or } \Sigma^i \vdash \psi_1 \\ & \Delta_k \cup \{\psi_2\} \text{ otherwise.}\end{aligned}$$
  - $\Sigma^{i+1} = \bigcup_{k \geq 0} \Delta_k$

It is straightforward to show that the conditions (a)–(c) of saturated sets apply to these temporal snapshots. The final property (d), that ϕ is not contained in any Σ^i, where ϕ is the sentence used in the construction of the saturated extension to Γ, does not always hold; if ϕ is a stable sentence, it may well get added to some Σ^i. But in exchange for this condition, it can be shown that the following holds:

For all Σ^i, if $\psi \in \Sigma^i$, then $\sim\sim \psi \in \Sigma^0$

Constructing an Information Model

To construct an information model where $\Gamma \not\models \phi$ given that $\Gamma \nvdash \phi$

1. Construct a saturated extension, Σ_0^0, of Γ relative to ϕ and some enumeration of L, and construct a sequence of temporal snapshots, Σ_0^i relative to the same language enumeration.
2. Take $Stable(\Sigma_0^0)$ and construct saturated extensions from it relative to the absurd proposition $\perp$, and temporal extensions of these saturated extensions. These form the information states of the canonical information model.
3. For any Σ_s^t and $\Sigma_{s'}^t$ formed in this, set

$s \sqsubseteq_t s'$ iff
(a) $Stable(\Sigma^t_s) \subseteq Stable(\Sigma^t_{s'})$, and
(b) for any semi stable implication $\phi \rightarrow \psi \in \Sigma^t_s$, if $\neg\psi \in \Sigma^t_{s'}$ then $\phi \rightarrow \psi \notin \Sigma^t_{s'}$

4. The valuation function for atomic sentence letters coincides with the atomic sentences contained in a given state s at time t, i.e. Σ^t_s.

We need to establish the models constructed in this way are legitimate information models. It is easy to show that the relation $\sqsubseteq_t$ is transitive and reflexive from the transitivity and reflexivity of the $\subseteq$ relation. That the relation converges (i.e if $s_0 \sqsubseteq_{t_0} s_1 \sqsubseteq t_1 s_2$, then $s_0 \sqsubseteq_{t_2} s_2$ for some t_2 such that $t_2 \geq t_0$ and $t_2 \geq t_1$) can be seen (for a denumerable language L) by picking a time large enough that all the temporal extensions of s_0, s_1 and s_2 become fixed and unchanging. The monotonicity of the valuation function with respect to subsequent verification times follows from the fact that all the sentences contained in Σ^t_s are also included in Σ^{t+1}_s. Monotonicity of information growth for atomic sentences arises from the fact that all atomic sentences are classed as stable, and so will be included in all states extending a state that contains them.

Before showing that complex formulas also satisfy monotonicity of information growth where appropriate, it is necessary to establish one other thing about these canonical information models.

Partial monotonicity of semi-stable implications:
If $\phi \rightarrow \psi \in \Sigma^t_s$, then $\forall s', t'$ if $s' \sqsupseteq_{t'} s$, $\phi \in \Sigma^{t'}_{s'}$ and $\neg(\exists s'', t''$ such that $t \leq t'' < t'$, $s \sqsubseteq_{t''} s'' \sqsubseteq_{t''} s'$ and $\phi \in \Sigma^{t''}_{s''}$,
then $\phi \rightarrow \psi \in \Sigma^{t'}_{s'}$

Proof: (a) Suppose that $\neg\psi \in \Sigma^t_s$. It can be proved that $\{\phi \rightarrow \psi, \neg\psi\} \vdash \neg\phi$ and so Σ^t_s has no ϕ-supporting extensions, minimal or otherwise. And so the condition is vacuously true. (b) Suppose $s' \sqsupseteq_{t'} s$, $\phi \in \Sigma^{t'}_{s'}$ and $\phi \rightarrow \psi \notin \Sigma^{t'}_{s'}$. By the construction of the model it follows that $\neg\psi \in \Sigma^{t'}_{s'}$. By (a) it can be seen that $\neg\psi \notin \Sigma^t_s$, which means that $\Sigma^{t'}_{s'}$ cannot be a minimal ϕ-extension of Σ^t_s. That is, by the construction of the model, there will be some $\Sigma^{t''}_{s''}$ preceding $\Sigma^{t'}_{s'}$ and extending Σ^t_s in which ϕ holds, and where $\neg\psi$ does not.

In other words, semi-stable implications persist until their consequents become false. Any minimal extension of a state in which a semi-stable implication holds to support the implication's antecedent will lead to a state in which the implication still holds.

This result helps to nullify problems caused by the monotonicity of some unstable formulas. Recall that unstable formulas can sometimes be valid or invalid (and hence monotonic), conjunctions or disjunctions of stable formulas with valid or invalid unstable formulas, in-state negations of these formulas, or implications into them. Taking $Stable(\Sigma)$

to form the basis of extending information states might seem to miss out on some of the monotonic formulas holding in Σ.

Only implication is of any concern. Since Σ is saturated, all conjunctions and disjunctions and Σ are broken down into their component parts, and will be included in $\mathit{Stable}(\Sigma)$, and so will be included in states extending Σ. Valid (or invalid) unstable formulas will be included (or excluded) in extensions of Σ (by the logical closure of saturated sets), allowing monotonic conjunctions and disjunctions in Σ to be reformed by logical closure. In-state negation can be pushed in so that it has scope over atomic formulas, or out-of-state negations ($\sim (\phi \rightarrow \psi) \equiv \sim \neg(\phi \wedge \sim \psi)$). Out of state negations, $\neg\phi$ are monotonic, and may either be valid or invalid, in which case $\sim \neg\phi$ is monotonic, or contingent, in which case $\sim \neg\phi$ is non-monotonic. In the former case, validities will be included and invalidities included in states built on $\mathit{Stable}(\Sigma)$ by logical closure, causing $\sim \neg\phi$ to be excluded or included in saturations of $\mathit{Stable}(\Sigma)$ as required. When $\neg\phi$ is contingent, it will not and should not be included in $\mathit{Stable}(\Sigma)$.

With implication, note that $\phi \rightarrow (\psi \wedge \chi)$ is equivalent to $(\phi \rightarrow \psi) \wedge (\phi \rightarrow \chi)$, so problematic unstable monotonic conjunctions in the consequent get broken down into their component parts. This leaves $\phi \rightarrow (\psi \vee \chi)$, where χ is a valid or invalid unstable formula. This is where the partial monotonicity of semi-stable implications comes to the rescue. First, let χ be valid. Then $\psi \vee \chi$ will never become false, and so any extension of Σ will include $\phi \rightarrow (\psi \vee \chi)$ by the partial monotonicity of semi-stable implications. Now let χ be invalid, and note that $\phi \rightarrow (\psi \vee \bot)$ is equivalent to $\phi \rightarrow \psi$, which is stable if ψ is stable, and so will be included in $\mathit{Stable}(\Sigma)$.

Forcing and Containment

To finish the completeness proof, we need to establish that

$$s, t \Vdash \psi \text{ iff } \psi \in \Sigma_s^t$$

To set the ball rolling, for atomic sentences we define the forcing relation as

$$s, t \Vdash p \text{ iff } p \in \Sigma_s^t, \text{ atomic } p$$

Then by an inductive argument we show that the same holds for complex formulas. Note that Σ_s^t is deductively closed, and so $\psi \in \Sigma_s^t$ iff $\Sigma_s^t \vdash \phi$.

Conjunction $s, t \Vdash \phi \wedge \psi$ iff $\phi \wedge \psi \in \Sigma_s^t$

Only if: Assume $\phi \wedge \psi \in \Sigma_s^t$.
By $\wedge$-E, $\phi \in \Sigma_s^t$, $\psi \in \Sigma_s^t$.

By inductive hyp., $s, t \Vdash \phi$, $s, t \Vdash \psi$.
By semantics of $\wedge$, $s, t \Vdash \phi \wedge \psi$.

If: Assume $s, t \Vdash \phi \wedge \psi$.
By semantics of $\wedge$, $s, t \Vdash \phi$, $s, t \Vdash \psi$.
By inductive hyp., $\phi \in \Sigma_s^t$, $\psi \in \Sigma_s^t$.
By $\wedge$-I, $\phi \wedge \psi \in \Sigma_s^t$

Disjunction $\quad s, t \Vdash \phi \vee \psi$ iff $\phi \vee \psi \in \Sigma_s^t$

Only if: Assume $\phi \vee \psi \in \Sigma_s^t$.
By saturation of Σ_s^t, $\phi \in \Sigma_s^t$, or $\psi \in \Sigma_s^t$.
By inductive hyp., $s, t \Vdash \phi$ or $s, t \Vdash \psi$.
By semantics of $\vee$, $s, t \Vdash \phi \vee \psi$.

If: Assume $s, t \Vdash \phi \vee \psi$.
By semantics of $\vee$, $s, t \Vdash \phi$, or $s, t \Vdash \psi$.
By inductive hyp., $\phi \in \Sigma_s^t$, or $\psi \in \Sigma_s^t$.
By $\vee$-I, $\phi \vee \psi \in \Sigma_s^t$

In-State Negation $\quad s, t \Vdash \sim \phi$ iff $\sim \phi \in \Sigma_s^t$

Only if: Assume $\sim \phi \in \Sigma_s^t$
By construction of model, for all $t' \geq t$, $\sim \phi \in \Sigma_s^{t'}$
By $\sim$-I, $\phi \notin \Sigma_s^{t'}$
By inductive hyp., $s, t' \nVdash \phi$
By semantics of $\sim$, $s, t \Vdash \sim \phi$

If: Assume $s, t \Vdash \sim \phi$
By semantics of $\sim$, for all $t' \geq t$, $s, t' \nVdash \phi$
Suppose $\sim \phi \notin \Sigma_s^t$ and derive a contradiction:
By $\sim$-Ax, if $\sim \phi \notin \Sigma_s^t$, then $\sim\sim \phi \in \Sigma_s^t$
By semantics of $\sim\sim$, $\exists t' \geq t.\, s, t' \Vdash \phi$
Contradiction.

Implication $\quad s, t \Vdash \phi \rightarrow \psi$ iff $\phi \rightarrow \psi \in \Sigma_s^t$

Only if: Assume $\phi \rightarrow \psi \in \Sigma_s^t$
Consider a minimal extension of Σ_s^t, $\Sigma_{s'}^{t'}$, with respect to ϕ.
By the partial monotonicity of semi-stable implications, $\phi \rightarrow \psi \in \Sigma_{s'}^{t'}$ (If $\phi \rightarrow \psi$ is stable it will be in $\Sigma_{s'}^{t'}$ as a matter of course).

By $\rightarrow$-E, $\sim\sim\psi \in \Sigma_{s'}^{t'}$
By ind. hyp., $s', t' \Vdash \phi$, $s', t' \Vdash \sim\sim\psi$
By semantics of $\sim\sim$, $\exists t" \geq t'. s', t" \Vdash \psi$
Since $s' \approx_{t'} s'$, $\exists s" \sqsupseteq_{t'} s', \exists t" \geq t'. s", t" \Vdash \psi$
Since s' and t' were arbitrary, by semantics of $\rightarrow$, $s, t \Vdash \phi \rightarrow \psi$

If: Assume $s, t \Vdash \phi \rightarrow \psi$
Suppose $\phi \rightarrow \psi \notin \Sigma_s^t$ and derive a contradiction:
By $\rightarrow$-I, $Stable(\Sigma_s^t), \phi \not\vdash \sim\sim\psi$
Let $\Sigma_{s'}^{t'}$ be a saturated theory constructed from $Stable(\Gamma) \cup \{\phi\}$ relative to $\sim\sim\psi$ and some enumeration of L
By construction of model, there will be such a $\Sigma_{s'}^{t'}$, where $\Sigma_s^t \sqsubseteq_t \Sigma_{s'}^{t'}$
$\sim\sim\psi \notin \Sigma_{s'}^{t'}$, and likewise for any $\Sigma_{s''}^{t'} \approx_{t'} \Sigma_{s'}^{t'}$
By ind. hyp., $\exists s' \sqsupseteq_t s, t' \geq t. s', t' \Vdash \phi$ but not $\exists t'' \geq t'. s', t'' \Vdash \psi$
Contradiction.

Out-of-State Negation By analogy with implication: $\neg\phi =_{df} \phi \rightarrow \bot$.

Completeness To sum up: Constructing a model from $\Gamma \not\vdash \phi$, the state snapshot Σ_0^0 contains Γ but not ϕ. Consequently $0, 0 \Vdash \Gamma$ but $0, 0 \not\Vdash \phi$. So there is a model falsifying $\Gamma \models \phi$. Hence $\Gamma \not\models \phi$.