Chapter 6

Using the Taxonomy to Create Relation Definitions

6.1 Introduction

In this chapter, the linguistic data assembled in the taxonomy is finally given a theoretical interpretation. Until now, the investigation has focused on surface structures in text; actual words and phrases, and how they can be manipulated by writers. However, as Scott and Paris (1995) note, at some point it is necessary to go 'beyond the text', to produce some description of it which is independent of its surface linguistic structures, in terms of which these structures can be defined. This is the concern of the present chapter.

A great deal has already been written about the semantics and pragmatics of cue phrases, and it should be no surprise that many of the ideas in this chapter draw on or expand on existing work. There are, however, two novel elements in the study, which provide some interesting new perspectives. Firstly is the fact that the set of cue phrases under investigation is much larger than usual. Often, theorists concentrate on a small set of cue phrases, or even on a single phrase—for instance *but* (Spooren (1989), von Klopp (1993)) or *when* (Moens and Steedman (1988)). The present study, as a result of the requirement of productivity, is much broader in scope: we will be looking for parameters which are valid right across the space of cue phrases. The reasoning behind this approach is that much can be learned about the semantics of a given cue phrase by comparing it to a number of other quite different phrases. To investigate the semantics of *because*, for instance, it would be instructive to be able to point up the similarities and differences between it and phrases as diverse as *when*, *then*, *even if*, *or*, *to*, *on the other hand* and *but*.

A second difference in the present work is that each theoretical construct introduced will be motivated in exactly the same way; by examining a portion of the taxonomy of cue phrases, and noting a demand for features to represent the patterns of substitutability it contains.

It should be stressed that the aim of this chapter is not to come up with a complete and

watertight set of relation definitions. Producing a complete feature-theoretic account of the taxonomy is a huge task, and well beyond the scope of this thesis. The present aim is rather twofold:

- to demonstrate the utility of the substitution methodology, by noting some of the interesting theoretical constructs which emerge when it is pursued;
- to motivate a core set of features needed to describe the taxonomy, to serve as the basis for further investigations.

In Section 6.2, the set of feature definitions is presented. Each feature is motivated separately, using appropriate extracts from the taxonomy. Section 6.3 provides a brief account of how the individual features interact with each other. Section 6.4 takes another look at the mapping between the cue phrases for which definitions have been formulated and the coherence relations which they signal. A summary and discussion is given in Section 6.5.

6.2 Features Motivated by the Taxonomy

6.2.1 SEMANTIC and PRAGMATIC Relations

The first extract from the taxonomy to be considered is given in Figure 6.1.[1] Motivating examples are given in Texts 6.1 and 6.2:

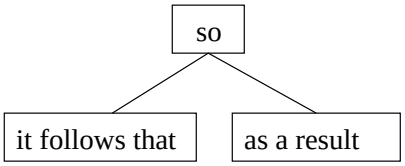

Figure 6.1: SEMANTIC and PRAGMATIC Phrases

$$\text{The footprints are deep and well-defined.} \left\{ \begin{array}{l} \textit{It follows that} \\ \surd \textit{So} \\ \# \textit{ As a result,} \end{array} \right\} \text{the thief was a heavy man.} \quad (6.1)$$

$$\text{I had a puncture on the M25 on my way back from work.} \left\{ \begin{array}{l} \textit{As a result,} \\ \surd \textit{So} \\ \# \textit{ It follows that} \end{array} \right\} \text{I missed most of the first half.} \quad (6.2)$$

As a result seems strange in the context of Example 6.1, because it suggests that the thief's heaviness is *caused* by the footprints being deep. Conversely, *it follows that* is odd in Example 6.2, because it suggests that the writer is *deducing* the fact that she missed most of the first half—while in fact she is reporting from her own experience. Note that *so* is acceptable in both cases.

[1] The extracts from the taxonomy given in this section will not be labelled with the features they serve to motivate. A labelled version of the taxonomy is given in Section 6.3; the reader might find it useful to refer forward to this section to see how feature values are assigned to cue phrases.

Both examples are commonly analysed as involving a causal/inferential relation of some kind, the difference between them being to do with what this relation holds between. Traditionally, the relation in examples like 6.2 is taken to be between the events in the world described by the two clauses: the puncture causes the missed first half. In examples like 6.1, the relation is taken to involve linguistic events themselves, not just the events they represent. For Martin (1992), Sanders *et al* (1992) and others, the writer's *statement* that the thief was heavy is caused by her belief that the footprints are deep. Commentators have used a variety of terms to represent these distinctions. Martin (1992) and Halliday and Hasan (1976) use the labels EXTERNAL and INTERNAL to refer to examples like 6.2 and 6.1 respectively; Redeker (1990) uses the labels IDEATIONAL and PRAGMATIC; Van Dijk (1979) and Sanders *et al* (1992) talk of SEMANTIC and PRAGMATIC relations.

A useful modification of the notion of PRAGMATIC relations is introduced by Sweetser (1990). For her, the relation in examples like 6.1 primarily describes the cause of the writer's *conclusion* that the thief must have been heavy, and only indirectly describes the cause of her *statement* to this effect. The important relation in the example is the logical one, between two of the writer's beliefs. Sweetser calls relations involving the writer's beliefs EPISTEMIC, and defines a further category of SPEECH ACT relations which make reference to actual writer utterances. Prototypical of SPEECH ACT relations are examples like the following:

(6.3) What are you doing tonight, because there's a good movie on.

Here it is certainly appropriate to analyse the relation as describing the cause of the writer's utterance *What are you doing tonight?*.

Sweetser's definition of EPISTEMIC relations is an improvement on the previous definitions. However, it still leaves something to be desired. Consider again Sweetser's EPISTEMIC analysis of Example 6.1, as a statement about the writer's conclusions and how they were reached. The text, according to the new analysis, is still fundamentally descriptive; instead of describing the external world, it now contains a description of the writer's own thought processes. What is missing is an account of how an argumentative text like this one achieves a rhetorical *effect* on the reader—how it *persuades* the reader that the thief was heavy, where a simple statement like *The thief was heavy* might not have sufficed. Of course, in offering the reader a trace of the writer's reasoning, the text suggests how the reader might come to the same conclusion. But while it is vital for the writer's purposes that the reader take this last step, it is not represented in Sweetser's analysis.

Thinking about utterances in terms of their intended effects on the reader suggests an alternative definition for EPISTEMIC relations. The new definition expresses the intended effect of a text containing two related utterances *as a relation between the intended effect of each individual utterance*. The suggestion is, for instance, that in a text containing an EPISTEMIC *so*, the writer's intended effect is not that the reader *believe* a statement about the causes of the writer's beliefs, but rather that a causal relation *actually does* hold, in the real world, between the intended effects of the two related utterances—in other words, between two reader beliefs.

Consider how this definition works in the case of Example 6.1. The text is presented

again below, and the intended effects of its two clauses are shown in italics:

(6.4)

The footprints are deep.		*So*	the thief was a heavy man.
R believes the footprints are deep.		*So*	*R believes the thief was a heavy man.*

On this interpretation, the intended effect of the text is that the reader's belief that the footprints are deep *causes* the reader to believe that the thief was a heavy man, where otherwise the reader would not have been inclined to believe this latter statement. It is because of *this* cause that the relation has persuasive force.

Note that the proposed new definition of EPISTEMIC relations actually extends to some relations Sweetser considers as SPEECH ACT. For Sweetser, "if an utterance is imperative (...) in form, then it cannot reasonably be causally conjoined to another utterance except at the speech act level" (p78). Thus Sweetser would interpret an example like 6.5 as SPEECH ACT:

(6.5) Hurry up, because we haven't much time!

But here again, there are advantages in expressing the relation in terms of intended effects. The writer's main intention in such a case is not to inform the reader about the cause of her utterance; but rather that the realisation that they haven't much time should motivate the reader to hurry up. Using the new definition, this is just what is expressed. The intended effect of the imperative *Hurry up* is that the reader hurry up; the intended effect of the statement *we haven't much time* is that the reader believe they haven't much time; and the intended effect of the whole utterance is that this belief causes the reader to hurry up.

Because the new definition encompasses examples such as this one, we have decided to revert to the label PRAGMATIC to refer to the relations it describes; and consequently to return to the label SEMANTIC for what Sweetser calls CONTENT relations. The definition for a feature with alternative values SEMANTIC and PRAGMATIC can now be given. After Sanders *et al*, we can call this feature SOURCE OF COHERENCE.[2]

SOURCE OF COHERENCE

SEMANTIC: the intended effect of the text containing the relation is that the reader believes some relation holds between two propositions A and C. A and C are the propositional contents of the two related text spans S_A and S_C.

PRAGMATIC: the intended effect of the text containing the relation is that some relation actually holds between two propositions A and C. A and C are the intended effects of the two related text spans S_A and S_C.

[2] It should be noted at this point that, as new features are motivated, changes are sometimes required to the definitions of features already introduced. In consequence, the feature definitions presented one by one in this chapter should not be regarded as final, but only as sufficient to account for the data so far introduced.

Note that the new notion of PRAGMATIC by no means covers all of Sweetser's SPEECH ACT relations. (For example, Text 6.3 is still much better analysed as describing the causes of the writer's speech act.) However, the class of SPEECH ACT relations has not yet been motivated from the taxonomy, as it is hard to find cue phrases which are specific to this class.

6.2.2 POSITIVE and NEGATIVE POLARITY Relations

A second portion of the taxonomy is given in Figure 6.2. Some motivating examples are provided in Texts 6.6 and 6.7.

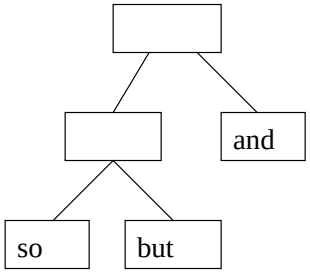

Figure 6.2: POSITIVE and NEGATIVE POLARITY Phrases

$$\text{Jim had just washed his car,} \left\{ \begin{array}{l} so \\ \surd\ and \\ \#\ but \end{array} \right\} \text{he wasn't keen on lending it to us.} \quad (6.6)$$

$$\text{It was odd. Bob shouted very loudly,} \left\{ \begin{array}{l} but \\ \surd\ and \\ \#\ so \end{array} \right\} \text{nobody heard him.} \quad (6.7)$$

The fact that the phrases *but* and *so* can never be substituted for one another is clear from consideration of examples like these. But the examples also show that the phrase *and* is contingently substitutable both for *but* and (in other contexts, of course) for *so*. In feature-theoretic terms, we can conclude that *but* and *so* are defined for different values of some feature; and that *and*, being contingently substitutable for both, is undefined for this feature. It remains now to decide what the feature is.[3]

Many different suggestions have been made as to the similarities and differences between phrases like *but* and *so*. It is uncontroversial (as far as it goes) that *A, so C* signals some kind of implication or cause, with *A* as the antecedent/cause and *C* as the consequent/result. And it is likewise uncontroversial to say that *A, but C* signals (or at least can signal) a violation of the type of relation signalled by *so*. To illustrate with reference to the above examples: in Text 6.6, *so* signals that it follows from the fact that Jim had just washed his car that he was unwilling to lend it to us. In Text 6.7, *but* signals that it *normally* follows from the fact that Bob shouts loudly that people

[3] Note that in order to motivate a feature which distinguishes between *so* and *but*, we are not obliged to find a common hypernym of the two phrases, as we did in the previous section. The fact that *and* can sometimes be substituted for both *but* and *so* is sufficient to show that it *cannot* be defined for any feature which takes alternative values for these two phrases. If it were, there would have to be an exclusive relationship between it and one or other of the phrases.

hear him, but in this case no-one does. Both phrases can thus be thought of as having a consequential component: for *so*, the consequence relation is specified as succeeding; while for *but*, an expected consequence is not forthcoming. With *and*, it is simply not specified whether or not the consequence relation succeeds—the information is left to be inferred by the reader.

The important question is how to express the above ideas more precisely. It has been common to begin formalising the difference between relations signalled by *so* and those signalled by *but* by making reference to a 'statement of implication' $P \rightarrow Q$ which underlies both types of relation. The difference between the two relations is then expressed in terms of the relationship between P and Q and the propositions in the related spans of text. For *so*, P relates to the proposition in the first span and Q to that in the second span. For *but*, P relates to the proposition in the first span and Q to the *negation* of that in the second span. This story is roughly that given by Longacre (1983) in distinguishing between 'consequence' and 'frustrated consequence' relations. Sanders *et al* (1992) give a similar story to distinguish between POSITIVE and NEGATIVE POLARITY relations; I shall use these latter terms in what follows.

Central to the distinction between *so* and *but* is the notion that causal or consequential rules can be **defeated**. A number of recent accounts of concessive relations have employed the notion of defeasible rules; in particular Oversteegen (1995) and Grote, Lenke and Stede (1995). Defeasible rules provide a useful method for representing the kind of common-sense generalisations which people rely on in order to make up for their partial knowledge of the world. They are becoming increasingly popular in computational linguistics, as a tool for modelling the influence of the reader's world knowledge on the resolution of ambiguities. For instance, Lascarides and Asher (1991), Lascarides, Asher and Oberlander (1992) use a system of defeasible rules to develop a framework for deciding which coherence relation is present at a particular point in a text when this is not signalled explicitly. Hobbs *et al* (1993) use defeasible rules to model a range of processes in text interpretation, including the resolution of anaphora, lexical ambiguities and compound nominals. The use of defeasible rules in the present context is somewhat different, however. They are not being proposed as a way of *deciding* about the interpretation of some part of a text, but as a part of the interpretation *itself*—the defeasible rules used by a reader and writer to model the world are actually implicit in the semantics of phrases like *but*. As an initial model, then, we might propose that the phrases *so* and *but* are each associated with a defeasible rule, which in the case of *so* succeeds and in the case of *but* is defeated.

A number of questions still remain, however. Most importantly, what is the communicative status of the defeasible rule? Is it something which the reader must already know as a precondition to understanding the text, or is it something which the reader is told in the text? Along with Oversteegen (1995), it is here proposed that the existence of the defeasible rule should be seen as a precondition. It is problematic to suggest that the rule itself is part of the information communicated to the reader by the writer. For one thing, the reader is only given one instance of the rule—it would then be necessary to abstract away from this to the rule itself; a process which is very underconstrained. Moreover, it is questionable whether causal or inferential *rules* constitute the kind of information that a reader will accept 'on authority' from a writer in any circumstance. Consider again the statement in Text 6.6: *Jim had just washed his*

car, so he wasn't keen on lending it to us. This may certainly provide new information, but it is implausible to suggest that the writer is *informing* the reader that 'if a person with temperament T has just washed his car, he normally doesn't like to lend it to others'. It is more plausible to suggest that a rule along these lines is already known by the reader, and what is being communicated is the fact that the rule succeeds in this instance.

Note that the information conveyed by such a statement might be more than the bare assertion that 'there is nothing unusual about the situation being described'. Knowing that the rule in question is *triggered* would allow the reader to infer that Jim is of type T, for instance, if this was not already known. But it is much easier to imagine the reader adding such facts to his database than whole causal rules.

To sum up: we can hypothesise a feature called POLARITY, with alternative values NEGATIVE and POSITIVE. It is assumed that each relation presupposes the presence of a defeasible rule $P \rightarrow Q$.[4] The relationship between P and Q and the propositions A and C (defined in the SOURCE OF COHERENCE feature) is determined by the different values of the POLARITY feature, as follows:

POLARITY

POSITIVE: $A = P$; $C = Q$. The rule is specified to succeed.

NEGATIVE: $A = P$; C is inconsistent with Q. The rule is specified to fail.

Conditional NEGATIVE and POSITIVE POLARITY

Another portion of the taxonomy which can be used to motivate the POLARITY parameter is given in Figure 6.3. Motivating examples are given in Texts 6.8 and 6.9:

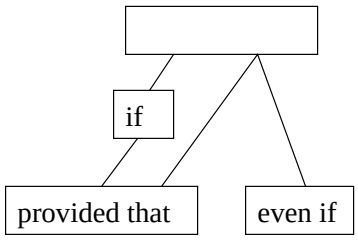

Figure 6.3: Conditional POSITIVE and NEGATIVE POLARITY Phrases

You can sit in the front seat, { *provided that* / ✓ *if* / # *even if* } you put your seatbelt on. (6.8)

I wouldn't vote for Major { *even if* / ✓ *if* / # *provided that* } you gave me a thousand pounds. (6.9)

[4] We will not at this point buy into any particular formalism for representing defeasible rules. However, some of the requirements for the formalism eventually to be adopted will emerge from the discussion in following sections.

The phrases in Figure 6.3 can be compared to those in Figure 6.2, a crucial difference being the fact that the latter group of phrases relate hypothetical eventualities while the former phrases relate actual ones. Motivation for a feature representing this dimension of variation will be provided below, in Section 6.2.8; for now the important thing to note is the variation in polarity exhibited by the phrases. For each phrase, an underlying defeasible rule $P \rightarrow Q$ can be identified. For *C, provided that A*, *A* and *C* map onto *P* and *Q* respectively and the rule is represented as succeeding. For *C, even if A*, *A* and *C* map onto *P* and $\neg Q$ respectively, and the rule is represented as failing. Thus in Example 6.8, the rule that putting a seatbelt on causes being allowed to sit in the front seat is asserted to succeed, while in Example 6.9 the rule that giving people lots of money causes them to vote against their will is asserted to fail in the case of the writer.

SEMANTIC and PRAGMATIC NEGATIVE POLARITY Relations

The distinction between POSITIVE and NEGATIVE POLARITY also cuts across the SEMANTIC/PRAGMATIC distinction. Consider Figure 6.4, for which motivating examples are given below:

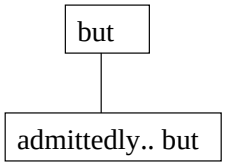

Figure 6.4: SEMANTIC and PRAGMATIC NEGATIVE POLARITY Phrases

$$\text{United have some key players injured;} \left\{ \begin{array}{l} \textit{admittedly... but} \\ \surd\ \textit{but} \end{array} \right\} \text{they're still bound to win.} \qquad (6.10)$$

$$\text{Mary was behaving oddly. She ordered a pizza,} \left\{ \begin{array}{l} \textit{but} \\ \#\ \textit{admittedly... but} \end{array} \right\} \text{she didn't eat any of it.} \qquad (6.11)$$

(The construction *admittedly... but* in these examples is to be read as distributed between the two clauses in the relation. The first example should thus read '*Admittedly,* United have some key players; *but*...', and the second example should read '*Admittedly,* she ordered a pizza, *but*...'.)

The point is that *Admittedly... but* signals the breaking of a defeasible rule just as *but* does; yet it has a specifically argumentative flavour. In Text 6.10, *admittedly* introduces a proposition which suggests one conclusion, and the negation of that conclusion is then asserted. *But* on its own can also be used in the absence of any argument, as in Text 6.11: here, the writer is simply informing the reader about an unusual state of affairs, and *admittedly* is quite out of place.

The SEMANTIC/PRAGMATIC distinction is useful in capturing the difference between these NEGATIVE POLARITY phrases. *Admittedly... but* can be defined as signalling the value PRAGMATIC, and *but* can be thought of as undefined for the feature. The

difference between SEMANTIC and PRAGMATIC NEGATIVE POLARITY relations can be thought of as follows. In the SEMANTIC case, the writer's aim is to inform the reader that some general rule in the reader's model of the world is defeated in the situation being described. Thus in the above example, the defeated rule is that people who order food generally eat it. In the PRAGMATIC case, the writer's aim is that some general rule in the world itself *actually* fails in the present instance. This rule holds between two reader beliefs—the intended effects of the first and second clauses taken individually. In the above example, the intended effect of the first span is that the reader believe that United has some key players injured; that of the second span is that the reader believe that United will win. Normally, if the reader believes a team has several players injured, he will believe they will lose; but in this case, the writer's intention is that this conclusion is not drawn.

6.2.3 UNILATERAL and BILATERAL Relations

The next portion of taxonomy to be considered, given in Figure 6.5, also involves NEGATIVE POLARITY phrases. Motivating examples are given in Texts 6.12 and 6.13:

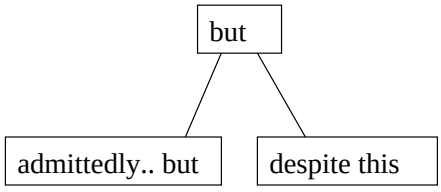

Figure 6.5: UNILATERAL and BILATERAL Phrases

Bill lost the 400m last year. { *Admittedly... but* / ✓ *But* / ✓ *Despite this,* } He should win it this year. (6.12)

Bill should win the 400m. He lost last year; { *Admittedly... but* / ✓ *but* / # *despite this,* } they're running at altitude this time. (6.13)

The relation in each of these examples can be signalled by *admittedly... but*, and can thus be considered as PRAGMATIC NEGATIVE POLARITY. However, the phrase *despite this* is only acceptable as a substitute in Text 6.12. In Text 6.13 it appears odd; it suggests that Bill's losing last year would normally lead to them not running at altitude this time—an implausible assumption.

The difference between the two examples appears to be to do with the status of the second span in the relation (the one introduced by *but*). In each case, the first span presents a premise P_1 which suggests a conclusion C. In Text 6.12, the second span presents the negation of this conclusion, $\neg C$. In Text 6.13, the second span presents another premise P_2, which is more telling than P_1, and suggests an alternative conclusion. This dimension of variation suggests another parameter, which we can call PATTERN OF INSTANTIATION.

To allow for multiple premises in the rule underlying the relation, we need to review the definition of POLARITY presented in the previous section. The assumption must

now be that each relation presupposes a rule of the form $P_1 \wedge \ldots \wedge P_n \rightarrow Q$. We need to map elements from this rule onto the variables A and C, defined in the feature SOURCE OF COHERENCE. Let us assume that A is always on the left-hand side of the rule.[5] The POLARITY feature presented before assumed that C always mapped onto the conclusion of the rule (Q); but we must now abstract away from this idea. We must define a new variable—call it C'—whose relationship to C is determined by the value of the POLARITY feature. For symmetry, we will also introduce a variable A', which always equates directly with A.[6] The POLARITY feature can now be thought of as specifying a function from A and C to A' and C'. Its revised definition now looks like this:

POLARITY (2nd definition)

POSITIVE: $A' = A$; $C' = C$. The rule is specified to succeed.

NEGATIVE: $A' = A$; C' is inconsistent with C. The rule is specified to fail.

The mapping between A' and C' and the rule $P_1 \wedge \ldots \wedge P_n \rightarrow Q$ is now given by the new feature:

PATTERN OF INSTANTIATION

UNILATERAL: C' is on the same side of the rule as A' (but not the same as A).

BILATERAL: C' is on the opposite side of the rule to A'.

The phrase *despite this* can now be thought of as defined as BILATERAL, while both *but* and *admittedly... but* are undefined for the feature. Thus in Example 6.12, for instance, A' is the proposition 'Bill lost the 400m last year' and C' is the proposition 'It is *not* the case that Bill should win this year'. These two propositions can be thought of as premise and conclusion of a defeasible rule. The relation can thus be thought of as BILATERAL, (and thus *despite this* is appropriate). Note that since the relation is also NEGATIVE POLARITY, the point is that the rule is defeated.

Several examples of phrases defined as UNILATERAL will be given in the following sections.

6.2.4 CAUSAL and INDUCTIVE Relations

Consider next the extract of the taxonomy in Figure 6.6, again featuring NEGATIVE POLARITY phrases. Motivating examples are given below:

[5] This assumption itself changes in Section 6.2.5, for reasons developed in that section.

[6] This is another assumption which will be re-examined as further portions of the taxonomy are considered; see Section 6.2.6 for details.

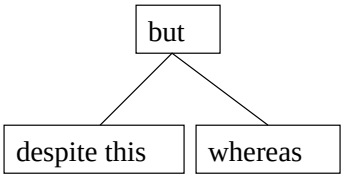

Figure 6.6: CAUSAL and INDUCTIVE Phrases

Bill and Jill are like chalk and cheese. Bill lives for his books; { *whereas* / ✓ *but* / # *despite this,* } Jill is only interested in Tae Kwan Do. (6.14)

Bob was out of training; { *despite this* / ✓ *but* / # *whereas* } he completed the marathon in record time. (6.15)

These phrases draw attention to another dimension of variation in the phrase *but*. Text 6.15 is of a kind we have already seen, where a causal or inferential rule is defeated: the rule in this case is that people who are out of training do not normally break records. But no such rule appears to underlie Text 6.14: knowing that Bill lives for his books gives no grounds for thinking that Jill is not just interested in Tae Kwan Do.

Whereas intuitively signals a contrast between two propositions. On a conventional analysis (see e.g. Spooren (1989)), we are able to say S_1, *whereas* S_2 if S_1 and S_2 allow the inference of $p(A)$ and $\neg p(B)$ respectively, for some predicate p and two entities A and B.[7] The question posed by the extract given in Figure 6.6 is: what does this have in common with the violated expectation analysis required for *despite this*? The feature-theoretic interpretation of the diagram requires us to find some feature or features of *whereas* that are shared by *despite this*; *but* is defined for this common component, and undefined for the feature(s) which distinguish the two phrases.

A point to note about $p(A)$ *whereas* $\neg p(B)$ is that A and B are required to belong in some sense to the same category of entities. Constrasts are not made between objects which have nothing in common at all. Thus Bill and Jill might be brother and sister, or friends, or two candidates for some job that needs doing. Put another way: relations signalled by *whereas* highlight an inability to *generalise* over the objects in a given class as regards some property p. This idea prompts the suggestion that a different type of rule underlies such relations; namely **inductive rules**.

An inductive rule is of the following general form: if property p is true of a certain finite number of elements from a particular class, then it follows that p is true of *all* the elements in that class. Clearly, as rules of inference, such statements are not sound. But inductive rules are nonetheless a mainstay of human reasoning: in the end, all our generalisations about the world are arrived at by inducing from particular instances. They can in fact be considered as another kind of *defeasible* rule, albeit quite different from those which we have so far been considering. Both kinds of rules

[7] In Text 6.14, for instance, we can infer from *Bill lives for his books* that *Bill is not only interested in Tae Kwan Do*, which is an explicit negation of the predicate in the second span of the relation.

are to be relied on when no information to the contrary is found, but are overruled if contrary information is forthcoming.

A proposition like $p(A)$ can thus be thought of as forming part of the left hand side of an inductive rule. If enough other objects in the same class as A have property p, then the inductive generalisation that $\forall X(X \in C \rightarrow p(X))$ will be triggered. Even the single instance of $p(A)$, in the absence of any other information, is presumably enough to trigger some very weak hypotheses about other similar objects. However, the inductive rule is defeated as soon as $\neg p(B)$ is presented.

In summary, what *whereas* and *despite this* seem to have in common is that they both presuppose a defeasible rule of some kind, and both signal its defeat. In one case the defeasible rule is causal and in the other case, inductive; *but* can then be seen as undefined with respect to the type of rule which is defeated. A new feature RULE TYPE is now motivated:

RULE TYPE

CAUSAL: the defeasible rule $P_1 \wedge \ldots \wedge P_n \rightarrow Q$ is a causal rule.

INDUCTIVE: the defeasible rule $P_1 \wedge \ldots \wedge P_n \rightarrow Q$ is an inductive rule.

As might have been noticed, *whereas* and *despite this* differ not only as regards the feature RULE TYPE, but also as regards the feature PATTERN OF INSTANTIATION. *Despite this* is defined as BILATERAL, as we have already seen in Section 6.2.3. *Whereas* must be defined as UNILATERAL: its two spans present two propositions from which generalisations can be drawn, and the generalisation itself (or rather the lack of it) remains implicit. Consider Example 6.14 in more detail. A and C can be identified as 'Bill lives for his books' and 'Jill is only interested in Tae Kwan Do' respectively. The relation is NEGATIVE POLARITY, so while A' is identical to A, C' is inconsistent with C. *Whereas* is defined as UNILATERAL; so A' and C' are both on the left-hand side of some defeasible rule. The relation is INDUCTIVE, so each item on the left-hand side of the rule will be attributing the same predicate to a different member of a given class. In this case, then, C' must be defined as 'Jill lives for her books', or something to that effect. As this is inconsistent with the value of C, the inductive rule fails.

SEMANTIC and PRAGMATIC INDUCTIVE Phrases

The distinction between CAUSAL and INDUCTIVE relations cuts across that between SEMANTIC and PRAGMATIC relations. Consider the extract in Figure 6.7, motivated by the examples below:

Bill and Jill are like chalk and cheese. Bill lives for his books; { *whereas* / ✓ *on the other hand,* / ✓ *but* / # *then again,* / # *despite this,* } Jill is only interested in Tae Kwan Do. (6.16)

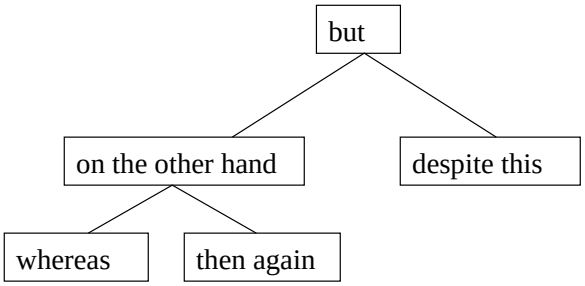

Figure 6.7: SEMANTIC and PRAGMATIC INDUCTIVE Phrases

I don't know where we should eat tonight. The Kalpna is great value for money; { *then again*, / ✓ *on the other hand*, / ✓ *but* / # *whereas* / # *despite this* } Sid isn't crazy about Indian food. (6.17)

Example 6.16 is the kind of contrast we have already considered, between objects in the world with inconsistent predicates. Example 6.17 can also be analysed as a contrast of sorts, but here the inconsistency is between the argumentative force of two propositions. The fact that the Kalpna is good value for money suggests that we should eat there. But the fact that Sid doesn't like Indian food suggests that we should not eat there.

The interesting relationship in this diagram is between *whereas*, *then again* and *on the other hand*. *Then again* seems quite wrong in the first example, and *whereas* is out of place in the second one. But note that *on the other hand* is acceptable in both cases. It will be suggested in the remainder of this section that all three phrases are signallers of NEGATIVE POLARITY INDUCTIVE relations; that the exclusivity between *whereas* and *then again* is due to the SEMANTIC/PRAGMATIC distinction; and that *on the other hand* is undefined for this latter feature and hence substitutable for both phrases.

The contrast signalled by *whereas* relates to the propositional content of the related spans, and the writer's objective in presenting it is to make the reader aware of the generalisation which fails. It can thus be thought of as a SEMANTIC contrast. But clearly, the contrast signalled by *then again* does not have to relate to the propositional content of the spans. No inconsistent predicates are present in the contents of the two spans in Example 6.17, for instance. In such cases, different kinds of objects are apparently being compared. A PRAGMATIC analysis of these cases will be suggested here: in this analysis, the objects are reader beliefs (rather than objects in the world), the predicates about the objects concern the different conclusions supported by different beliefs, and the classes into which beliefs fall concern the conclusions to which they are relevant.

Consider what happens in the process of 'reaching a conclusion by examining premises'. There must first be an attempt to delineate those premises which will be relevant; an exhaustive search will not be feasible in any system with a reasonably sized set of facts and rules. Even the set of relevant propositions is likely to be too large to be exhaustively searched, and further heuristics will need to be used to consider these selectively. It is reasonable to suppose that inductive principles play a part in these heuristics: if we consider a certain number of relevant premises, and each one supports

the same conclusion, there will come a point at which we decide that the conclusion is true, and stop looking for additional premises. The notion of 'reaching a conclusion' can be thought of as the moment when an inductive generalisation is made, and we assume that *all* premises relevant to the conclusion in fact support the conclusion.

Now consider what happens with *then again*. Here two relevant premises are presented which support opposite conclusions. This has the effect of blocking the generalisation, making the reader unable to reach a decision. In Lascarides and Asher's (1991) terminology, the reader has encountered a 'Nixon Diamond', where two defeasible rules are triggered, and neither takes precedence. Note that as this is a PRAGMATIC relation, the effect is not simply that the reader *realises* that inconsistent premises have been presented, but that a Nixon Diamond *actually happens* in the reader's theorem proving system. However, we must also note that the system does not freeze up altogether as a result of this impasse: it is only as regards one particular conclusion that no decision can be reached. Inductive generalisations can still be made to reach other conclusions.

POSITIVE and NEGATIVE POLARITY INDUCTIVE Phrases

Finally, consider the diagram in Figure 6.8, motivated by Texts 6.18 and 6.19.

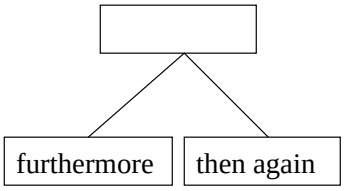

Figure 6.8: POSITIVE and NEGATIVE PRAGMATIC INDUCTIVE Phrases

United are bound to win. They have a great team; { *furthermore,* / # *then again,* } they're playing at home. (6.18)

I don't know whether United will win. They have a great team; { *then again,* / # *furthermore,* } they're playing away from home. (6.19)

If *then again* signals the defeat of a PRAGMATIC INDUCTIVE rule, then *furthermore* can be regarded as signalling the success of such a rule. In Example 6.18, a conclusion is reached: that United will win. Two relevant premises needed to be considered in order for the inductive rule to fire in this instance.

In other cases, more than two premises need to be advanced. Lists of premises signalled by phrases like *furthermore* can in principle be of any length. In the present model, these lists are analysed as nested applications of a binary relation, as in Figure 6.9 (i). (Premises are marked with a P; the conclusion with a C.) As regards the relations between premises, the topmost relation links the first premise with a complex span consisting of another relation between two further premises. It is easy to see how this pattern could be extended.

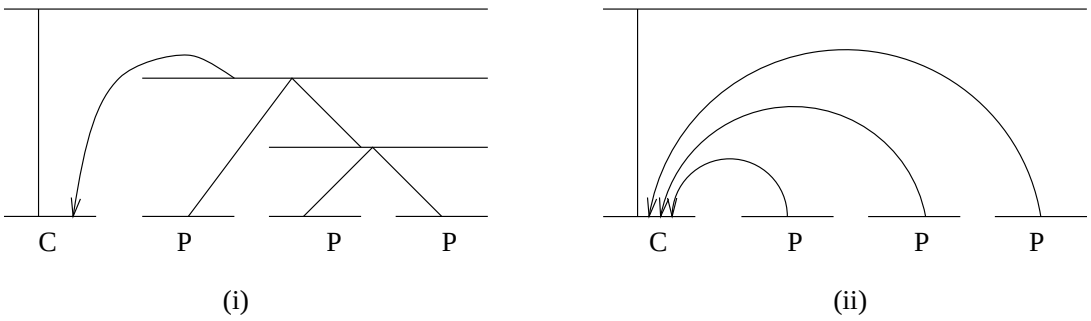

Figure 6.9: Alternative Structural Analyses of *Furthermore*

Figure 6.9 (ii) shows the structural analysis which RST would give for a text containing a sequence of premises. As already outlined in Section 2.4.3, the notion of **multiple schema applications** does service in such texts: the premises are not represented in relation to each other, but in relation to the conclusion they support. The conclusion span is thus related to several adjacent text spans. However, a disadvantage of this approach is that it can only be used if the conclusion is represented explicitly in the text; and this is far from always the case. If a conclusion is implicit, RST would have to analyse the premises using a much less informative relation, LIST. Maier and Hovy (1991) counter this problem by adding a separate level of 'textual' relations to the diagram in Figure 6.9 (ii), which link the adjacent premises (see Section 2.5.3). But in the present system, two levels of relations are not necessary: the relation defined by POSITIVE PRAGMATIC INDUCTIVE features is sufficiently abstract to capture both the relationship between two premises and that between the premises and the conclusion.[8]

6.2.5 CAUSE and RESULT-DRIVEN Relations

A further extract from the NEGATIVE POLARITY portion of the taxonomy is given in Figure 6.10. Motivating examples appear below.

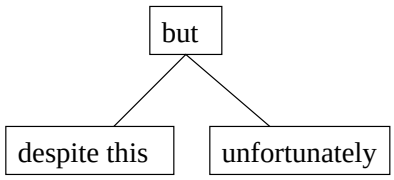

Figure 6.10: CAUSE-DRIVEN and RESULT-DRIVEN Phrases

Bill took the lid off the pot. $\left\{ \begin{array}{l} \textit{Unfortunately,} \\ \surd \textit{But} \\ \# \textit{ Despite this,} \end{array} \right\}$ there was nothing inside. (6.20)

[8] A question remains as to the nature of the relation which links the set of premises to the conclusion: it could either be PRAGMATIC POSITIVE CAUSAL or PRAGMATIC POSITIVE INDUCTIVE. I favour the former suggestion; note, for instance that *it follows that* can be used to introduce a conclusion supported by several premises.

Sue had been up all night; { *despite this,* / ✓ *but* / # *unfortunately,* } she was looking fresh as a daisy. (6.21)

The text in Example 6.20 can be thought of as presenting an unsatisfied desire. Let us say Bill wants to eat something; one way of achieving this goal would be if the pot contained food and its lid were removed. He lifts the lid, but there is nothing inside. *But* is certainly appropriate to describe the circumstance which blocks fulfilment of the goal; as is *unfortunately* (understood as 'unfortunately *for Bill*'). But *despite this* is quite wrong—it suggests that lifting the lid is expected to cause or entail that the pot is not empty. In Example 6.21, a text with a more conventional violated expectation is presented. Here *but* and *despite this* are acceptable, but *unfortunately* is inadmissible as a substitute.

These two different uses of *but* have been noted by others;[9] the following two examples of the unsatisfied-desire type are respectively from Longacre (1983) and Spooren (1989):

(6.22) I intended to go, but we had visitors that night.

(6.23) I went to the church, but the vicar was not there.

Both of these commentators attempt an explanation of such texts by proposing more abstract ways in which expectations are violated. For Longacre,

> something is presupposed here like the Newtonian assumption (inertia) that a body in motion in a given direction will keep moving in that direction unless some force deflects or stops it...

For Spooren, the expectation arises as a result of implicatures that follow from the statement of intention:

> part of our world knowledge is that going to church probably means that the vicar is in the church.

Neither of these explanations is very convincing. It is far from the case that intentions are normally achieved. It is certainly possible to set up contexts where an intention has no chance of being achieved and yet *but* is still appropriate. Imagine Jim is in a prison cell from which he cannot escape. We could still say that

(6.24) Jim looked around for food, but there was none to be found.

No amount of wanting or looking for food is going to satisfy Jim's intention, so the 'inertia' explanation is ruled out. His looking for food does not probably mean that there is food to be found, so the introduction of implicatures is similarly inadmissible. Yet the phrase *but* does not seem out of place.

[9] However, no-one to my knowledge has suggested using *unfortunately* as a diagnostic for the unsatisfied-desire use.

A preferable explanation—and one that is motivated by the pattern of substitutability in Figure 6.11—is that *but* is undefined for a further feature, for which *despite this* and *unfortunately* signal different values. The values of the feature relate to the manner in which the rule $P_1 \wedge \ldots \wedge P_n \rightarrow Q$ is used: are we *predicting* Q from our knowledge of $P_1 \wedge \ldots \wedge P_n$, or are we seeking to *achieve* Q, and thus investigating whether $P_1 \wedge \ldots \wedge P_n$ are true or themselves achieveable? We can call this feature ANCHOR, to reflect whether the 'certainty' relates to the knowledge of the premises, or the desirability of the conclusion. In each case, the 'certain' thing will be identified as A (for **anchor**). The feature can be defined as follows:

ANCHOR

CAUSE-DRIVEN: $A \in P_1 \ldots P_n$; $P_1 \wedge \ldots \wedge P_n$ is true.

RESULT-DRIVEN: A corresponds to Q; and A is desired by the protagonist.

To take an example, consider again Text 6.22: *I intended to go, but we had visitors that night.* This is a RESULT-DRIVEN relation: A, the first clause, which presents the writer's intention to go, corresponds to the right-hand side of the rule $P_1 \wedge \ldots \wedge P_n \rightarrow Q$. C is the second clause, presenting the fact which prevents the intention being achieved. The relation is BILATERAL, since C relates to a fact on the left-hand side of the rule $P_1 \wedge \ldots \wedge P_n \rightarrow Q$. It is NEGATIVE POLARITY, since the relevant fact in the rule (C') is inconsistent with C.

In one respect, the above definition of the ANCHOR feature is slightly fudged. Consider the original example of a RESULT-DRIVEN relation: *Bill took the lid off the pot; but there was nothing inside.* The first clause must still be considered A, but note that it does not *itself* present the intention. Rather, it presents an action performed to achieve the intention. There will be more to say about such cases in Section 6.2.7, which deals with presupposition. For the moment, note that the intention behind the action in the current example can be expressed in a subordinate clause:

(6.25) Bill took the lid off the pot *to* get some food; but there was nothing inside.[10]

Note that the revised definitions of POSITIVE and NEGATIVE POLARITY presented in Section 6.2.3 are still serviceable for both CAUSE-DRIVEN and RESULT-DRIVEN relations. In POSITIVE POLARITY relations, nothing is negated, and so no problems arise. And in NEGATIVE POLARITY relations, it is always C which is negated. For a CAUSE-DRIVEN NEGATIVE POLARITY relation, the anchor A is on the left-hand side of the rule, and C is the negation of the expected conclusion. For a RESULT-DRIVEN NEGATIVE POLARITY relation, the right-hand side of the rule is desired by the protagonist, and C is the negation of one of the conditions necessary for this desire to be brought about.

It should also be noted that the definition of PATTERN OF INSTANTIATION in Section 6.2.3 does not need amendment. This definition determines whether C is on the

[10] Strictly speaking, *to* is not a cue phrase, as it does not pass the test for relational phrases. But, as noted in Section 4.2, it is similar enough to a cue phrase to warrant attention.

same side of the rule as A (UNILATERAL), or whether they are on opposite sides (BILATERAL). In combination with the two values of the ANCHOR feature, four possible patterns of instantiation can now be expressed: A and C can both be on the left of the rule; or they can both be on the right; or A can be on the left and C on the right; or C can be on the left and A on the right.

Figure 6.11 provides some additions to the diagram in Figure 6.10. Again, motivating examples are provided:

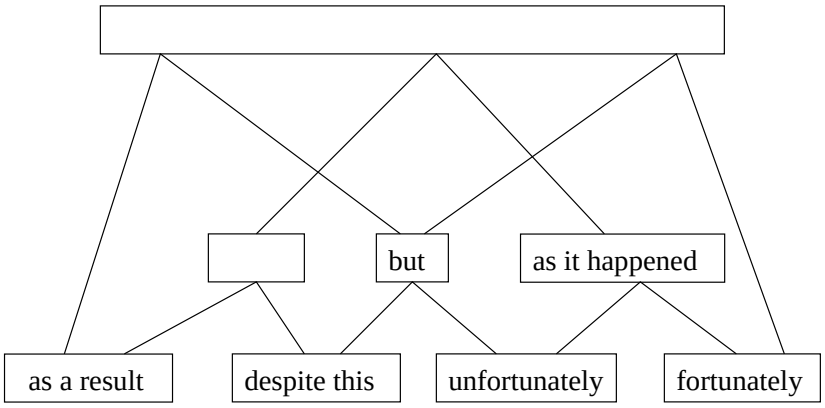

Figure 6.11: Additional CAUSE-DRIVEN and RESULT-DRIVEN Phrases

Bill took the lid off the pot. { *Unfortunately,* / ✓ *But* / ✓ *As it happened,* / # *Fortunately,* / # *Despite this,* / # *As a result,* } there was nothing inside. (6.26)

Bill took the lid off the pot. { *Fortunately,* / ✓ *As it happened,* / # *Unfortunately,* / # *But* / # *Despite this,* / # *As a result,* } there was something inside. (6.27)

Note that the exclusivity of *fortunately* and *unfortunately* depends crucially on the fact that the two phrases set up different implicatures about what Bill wants. Both of the above texts presuppose that Bill's plan requires there to be something in the pot. It might also have been that the plan required the pot to be empty—in which case the appropriateness of *fortunately* and *unfortunately* would be reversed. The two phrases cannot be substituted for one another, as to do so requires a change in the assumptions about what Bill's plan involves.

The most significant fact about the diagram in Figure 6.11 is that *as it happened* is substitutable for both *fortunately* and *unfortunately*, while these two latter phrases are exclusive. All three phrases are exclusive with CAUSE-DRIVEN phrases like *as a result* and *despite this*, and it is plausible that they should all be labelled as RESULT-DRIVEN. The variation within the phrases can then be traced to the POLARITY feature. *Fortunately* has POSITIVE POLARITY: the intention behind the first clause is achieved if the second clause is true. *Unfortunately* has NEGATIVE POLARITY: the intention

behind the first clause is achieved if the *negation* of the second clause is true. And *as it happened* is undefined with respect to the POLARITY feature.

Conditional Variants of CAUSE-DRIVEN and RESULT-DRIVEN Phrases

Evidence for the distinction between *cause-driven* and *result-driven* relations also shows up in the conditional phrases in the taxonomy. Consider the two examples below:

(6.28) If you feel like cooking tonight, make something spicy.

(6.29) If you feel like cooking, there's a kitchen on the first floor.

Both of these examples are of POSITIVE POLARITY PRAGMATIC relations. In Example 6.28, the rule which the writer wants to succeed is that people who feel like cooking and who have a certain disposition make something spicy. (The writer's aim is thus that the reader takes on this disposition.) But we cannot envisage a similar rule for Example 6.29. The kitchen is on the first floor whether or not the reader feels like cooking. The point is rather that it is only *relevant* for the reader to know where the kitchen is if he feels like cooking. Treating the *if* in this text as RESULT-DRIVEN thus provides a better analysis. We assume there is a goal underlying the first clause; namely that the reader cooks. This will happen if the reader feels like cooking, *and* knows where the kitchen is.[11]

A similar story can be told for NEGATIVE POLARITY relations.

(6.30) Even if you manage to break out of the prison, you'll never make it home.

(6.31) You'll never make it home. Even if you manage to break out of the prison, the jungle round here is impenetrable.

In Example 6.30, the presupposed rule is that if one breaks out of prison, one can normally get home. The rule is defeated in the present case. In Example 6.31, however, there is no rule stating that breaking out of prison normally entails the jungle not being impenetrable. Rather, we must assume a goal behind the breaking out of prison, and a rule stating that the goal will be achieved if the outbreak occurs and the jungle is *not* impenetrable.

[11] An alternative analysis for this kind of text is given by Sweetser (1990). She considers such a text to be an example of a SPEECH ACT conditional, to be read as 'if you feel like cooking, then (let us consider that) I *inform* you that there's a kitchen on the first floor'. According to this analysis, the speech act of informing is only to be understood as having occurred if the information about the kitchen is considered relevant. However, it is odd to suggest that the speech act simply disappears if its content is not relevant. The information about the kitchen is conveyed to the hearer under any circumstances, even if it is not relevant; and it is hard to see what there is to an informative speech act beyond the deliberate conveying of information. The problem is exacerbated if the analysis is extended to cover examples such as *Whenever you feel like cooking, there's a kitchen on the first floor.* If we interpret this as a SPEECH ACT conditional, we must envisage a whole series of informative speech acts, one for each time the hearer feels like cooking. At this point, we are clearly stretching the notion of a 'speech act' beyond its normal use.

The distinction between CAUSE-DRIVEN and RESULT-DRIVEN conditional relations does not emerge systematically from the taxonomy. There are not always phrases which differ only with regard to this feature (although see Section 6.2.8 for a discussion of the phrase *in case*). The only systematic evidence for the feature in hypothetical relations is thus that the phrases *if* and *even if* can take either of its values. However, indirect evidence can be obtained by converting 'hypothetical' texts to 'actual' ones, and observing which cue phrases are now appropriate. The following two texts are 'actual' versions of Texts 6.31 and 6.29; the patterns of substitutability for the phrases *fortunately*, *unfortunately*, *despite this* and *so* are what we would expect for CAUSE-DRIVEN and RESULT-DRIVEN relations.

Bill managed to escape from prison. { *Unfortunately,* / ✓ *but* / # *despite this,* } the jungle was impenetrable. (6.32)

Bill felt like cooking. { *Fortunately,* / # *So* } there was a kitchen on the first floor. (6.33)

6.2.6 ANCHOR-BASED and COUNTERPART-BASED Relations

Another portion of the taxonomy is given in Figure 6.12. The motivating examples are as follows:

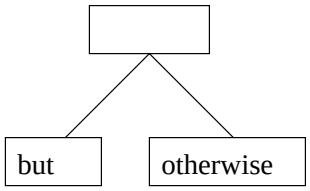

Figure 6.12: ANCHOR-BASED and COUNTERPART-BASED Phrases

Bob put his hands up, { *otherwise* / # *but* } Jill would have shot him. (6.34)

Bob kept his hands by his sides, { *but* / # *otherwise* } Jill didn't shoot him. (6.35)

It seems that both of these texts involve a rule along the following lines:

> If Bob doesn't put his hands up, Jill will shoot him.

For both texts, the anchor relates to the left-hand side of this rule, and the counterpart to the right-hand side. The texts are similar, in that the conclusion of the rule is avoided in each case. However, the reason for this is different in the two cases. In Example 6.35, the premise of the rule is true, but the rule is defeated: some stronger conflicting rule must therefore be supposed to have taken precedence. In Example 6.34, the premise of the rule does not even occur: the protagonist takes action to avoid a conclusion which is inconsistent with his goals.

In order to link the propositions related by the cue phrases onto the premise and conclusion of the relevant rule, another dimension of variation must be introduced. At present, in NEGATIVE POLARITY relations, it is always the *counterpart* span which needs to be negated to map back onto the rule $P_1 \wedge \ldots \wedge P_n \rightarrow Q$. This is still the case for Example 6.35; the counterpart is the second span, which is the negation of the expected conclusion. But for Example 6.34, it is rather the anchor (*Bill put his hands up*) which must be negated to map onto the rule, and the counterpart (*Jill would have shot him*) which maps onto the rule without being negated.

Examples such as this suggest the need for a new feature, which we can call FOCUS OF POLARITY, to specify whether the POLARITY transformation (if there is one) operates on the **anchor** or on the **counterpart**.

The definition of FOCUS OF POLARITY will require another modification to the definition of the POLARITY feature. In the current definition in Section 6.2.3, it is always the counterpart of the rule (C) which is a candidate for negation: the anchor A is never negated. We now need to express the definition so as to allow the candidate for negation to be determined by FOCUS OF POLARITY. To this end, we must introduce some new variables: the **focus of polarity** (F), which is the candidate for negation, and the **invariant** (I), which is never negated. The variables F' and I' will be used to represent F and I after the polarity transformation has taken place. The new definition for POLARITY now looks like this:

POLARITY (3nd definition)

POSITIVE: $F = F'$; $I = I'$;

NEGATIVE: $F = \neg F'$; $I = I'$.

The FOCUS OF POLARITY feature now identifies F and F' with A and A' (and I and I' with C and C') or F and F' with C and C' (and I and I' with A and A'.

FOCUS OF POLARITY

ANCHOR-BASED: $F = A$; $F' = A'$; $I = C$; $I' = C'$.

COUNTERPART-BASED: $F = C$; $F' = C'$; $I = A$; $I' = A'$.

Consider how these new definitions work with the phrases *but* and *otherwise*. For *Bob kept his hands by his sides, but Jill didn't shoot him*, the first clause is A and the second clause C. It is BILATERAL CAUSE-DRIVEN, so A' is part of the left-hand side of $P_1 \wedge \ldots \wedge P_n \rightarrow Q$, and C' is Q. The relation is COUNTERPART-BASED, so F is C and F' is C'. It is NEGATIVE POLARITY, so F is $\neg F'$; which means that C' is $\neg C$. A is the invariant, and hence maps straight onto A'. The effect is a violated expectation.

For *Bob put his hands up; otherwise Jill would have shot him*, the first clause is again A and the second clause C. The relation is BILATERAL CAUSE-DRIVEN, so A' is part of the left-hand side of $P_1 \wedge \ldots \wedge P_n \rightarrow Q$, and C' is Q. This time the relation is ANCHOR-BASED, so F is A and F' is A'. It is NEGATIVE POLARITY, so F is $\neg F'$; which means that A' is $\neg A$. This means that the rule does not trigger, and the right-hand

side of the rule (C) does not occur.[12]

SEMANTIC and PRAGMATIC ANCHOR-BASED Phrases

As it is expressed in terms of A and C, the definition of FOCUS OF POLARITY serves equally well for SEMANTIC relations and for PRAGMATIC ones. Consider the following two cases:

(6.36) Bob put his hands up, *otherwise* Jill would have shot him.

(6.37) Put your hands up, *otherwise* I'll shoot you.

Example 6.36 requires that the reader believe that a protagonist in the world being described (Bill) does not want some eventuality (being shot) to occur. It is thus a SEMANTIC relation, holding between the propositional contents of the related spans. However, for Example 6.37 it is a precondition that someone in the real world (namely the reader) does not want to be shot. The intended effect of the relation is that the reader actually put his hands up in order to avoid this eventuality. (Note that it is not specified whether the eventuality is *in fact* avoided; or even whether it is true that the writer would have shot the reader if he had not obeyed her instruction.)

Distinctions Amongst ANCHOR-BASED Relations

Otherwise is only one of a number of ANCHOR-BASED phrases. There are many others that seem to fall into this category: for example, *or*, *unless*, *until*, and *before* (see Section 6.2.7 for a discussion about this latter phrase). There is not room to talk about them all here, but the dimensions of variation between them look likely to correspond to those identified by other features. Deciding whether this is indeed the case is a matter for further research.

6.2.7 PRESUPPOSED and NON-PRESUPPOSED Relations

Another informative extract from the taxonomy is given in Figure 6.13. Motivating examples are given in Texts 6.38 and 6.39:

$$\text{I haven't always been unfit. I played a lot of rugby} \left\{ \begin{array}{l} \textit{while} \\ \checkmark\ \textit{when} \\ \#\ \textit{meanwhile} \end{array} \right\} \text{I was at college.} \tag{6.38}$$

$$\text{They set about preparing the meal. Bill marinated the meat;} \left\{ \begin{array}{l} \textit{meanwhile,} \\ \checkmark\ \textit{while} \\ \#\ \textit{when} \end{array} \right\} \text{Bob lit the barbecue.} \tag{6.39}$$

All three of the phrases in the diagram convey information about temporal simultaneity, among other things. However, there is a syntactic difference between *when*, which

[12] In fact, the story is likely to be more complicated than this. The reason why C does not occur is *because it is not desired by Bob*, and he takes action to avoid it. There thus seems likely to be a result-driven component to the relation which has not so far been captured.

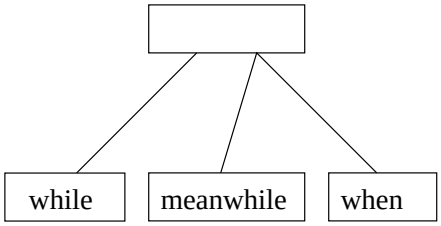

Figure 6.13: PRESUPPOSED and NON-PRESUPPOSED Phrases

is a temporal subordinator, and *meanwhile*, which is a sentential adverb. *While* is appropriate as a substitute for either phrase, as the examples show.

The clauses introduced by temporal subordinators are conventionally thought of as **presupposed** (Karttunen (1973), Keenan (1971), Lascarides and Oberlander (1993)). They describe an eventuality with respect to which the material in the main clause is temporally situated: this eventuality must either be already known to the reader, or must be **accommodated** prior to the addition of the temporal relation. The asymmetry of these sentences can be demonstrated by swapping the main and subordinate clauses. This typically results in incoherence, although the relation of 'temporal overlap' between the clauses is presumably unchanged:

(6.40) I haven't always been unfit. ?? I was at college *when* I played a lot of rugby.

On the other hand, no harm is done by changing the clause introduced by *meanwhile*:

(6.41) They set about preparing the meal. Bob lit the barbecue; *meanwhile* Bill marinated the meat.

The question of when accommodation is possible is addressed by Lascarides and Oberlander (1993). They propose two different mechanisms, which apply in different situations. If the subordinate clause functions simply as a temporal adverbial, as in Example 6.38, then there is no need to find a coherence relation between it and the preceding context; the important relation ('BACKGROUND', in this case) is between the main clause and the preceding context. But in other cases, as in the following example, the subordinate clause has an important narrative function:

(6.42) The backbenchers were in revolt. They were pacified after Major launched a charm offensive.

Here, a coherence relation between the subordinate clause and the preceding context must be found to allow accommodation, and only after the subordinate clause has been attached is the main clause considered. The explanation of the asymmetry introduced by the subordinator thus turns on the *order of attachment* of the two clauses.

A similar explanation of temporal subordinators must be sought for the present theory. Here, however, it must emerge from the definitions of the relations marked by these phrases, rather than from an account of the algorithm used to interpret them. We

therefore need to think about how to express the preconditions for a relation between two spans in terms of how these link to the immediately preceding context. We need a feature which takes one value for subordinating phrases like *when*, and another for non-subordinating phrases.

A feature with alternative values PRESUPPOSED and NON-PRESUPPOSED is used here to capture this difference. The feature introduces another defeasible rule, $X_1 \wedge \ldots \wedge X_n \rightarrow Y$, similar in structure to the one which the feature definitions have so far referred to. This rule makes a link between the span in the preceding context (which we will call *Precond*) and the anchor A. We can define the new feature as follows:

PRESUPPOSITIONALITY

PRESUPPOSED: *Precond* is part of the left-hand side of the rule $X_1 \wedge \ldots \wedge X_n \rightarrow Y$, and A is Y.

NON-PRESUPPOSED: *Precond* is A itself.

We can now give an account of what happens in Example 6.42. The first clause *the backbenchers were in revolt* is *Precond*. The rule $X_1 \wedge \ldots \wedge X_n \rightarrow Y$ has this clause as part of its left-hand side, and anchor clause *Major launched a charm offensive* as its right-hand side. When *Precond* occurs, the rule is thus triggered.

A similar story can be told for other temporal subordinators. Consider the following example:

(6.43) Bob heated the water. When it boiled, he stirred in the sugar.

The first clause *Bob heated the water* is *Precond* here. It can be seen as triggering a causal rule whose right-hand side is *the water boiled*, which is the anchor of the relation signalled by *when*.

We can now consider what happens in NON-PRESUPPOSED relations. An example of such a relation appears in the text below.

(6.44) Jill was curious, so she pulled the lever. Instantly, an alarm went off.

In fact, there appear to be two overlapping relations in this example: one between the first and second clauses (signalled by *so*); the other between the second and third clauses (signalled by *instantly*). It seems quite plausible that the counterpart of the first relation actually *identifies with* the anchor of the second relation. This is what is captured in the definition of NON-PRESUPPOSED: an identity is specified to hold between A and *Precond*.

RESULT-DRIVEN PRESUPPOSED Relations

It was noted above, in connection with Example 6.40, that swapping the main and subordinate clauses in a PRESUPPOSED relation leads to incoherence. However, it is

interesting that where the subordinate clause has a narrative function (rather than just acting as a temporal adverbial), coherent texts can be created by swapping the main and subordinate clause, provided that the right subordinating phrase is chosen. Consider the following variation on Example 6.42:

(6.45) The backbenchers were in revolt. Major launched a charm offensive to pacify them.

The anchor for the presuppositional relation is now 'The backbenchers are pacified', and this is what must be attached first to the preceding context. However, it is no longer the pacification *itself* which is caused by the context. Rather it is the *intention* that the backbenchers be pacified. We have now set up a context where *Major launched a charm offensive* can act as the anchor for a RESULT-DRIVEN rule of the kind discussed in Section 6.2.5.

ANCHOR-BASED PRESUPPOSED Phrases

Finally, it is interesting to note that the rule $X_1 \wedge \ldots \wedge X_n \rightarrow Y$ is defeasible, just like $P_1 \wedge \ldots \wedge P_n \rightarrow Q$. Consider this example:

(6.46) Tidy your room, before I lose my temper.

The relation signalled by the subordinator *before* is PRESUPPOSED; in other words, the subordinate clause 'W loses her temper' is A, and the main clause 'Reader tidies his room' is C. Since it is PRESUPPOSED, there is a rule running from the preconditions of the text (i.e. the situation that is currently true) to A. In other words, A is predicted to happen as things stand. However, the relation is also RESULT-DRIVEN, NEGATIVE POLARITY, and ANCHOR-BASED. These parameters describe the relationship between A and C. They specify that there is a rule $P_1 \wedge \ldots \wedge P_n \rightarrow Q$, whose left hand side includes C, the action which the writer intends the reader to perform, and whose right-hand side is the *negation* of A. 'Not A' is a desire of the reader's: he wants it not to be the case that the writer loses her temper. We thus have two conflicting rules: one leading to A and one leading to $\neg A$. The left-hand side of the former rule is currently true. The left-hand side of the latter rule contains what is currently true *plus* the desired reader action. The latter rule is the one which is intended to fire, and thus the rule leading to A is defeated.

The above account of counterfactual *before* still needs to be worked out in detail. However, it is interesting that the features being developed here seem well-suited for handling such cases.

6.2.8 HYPOTHETICAL and ACTUAL Relations

A final extract from the taxonomy is given in Figure 6.14. It is motivated by the following examples:

$$\text{We had a strict upbringing.} \left\{ \begin{array}{l} \textit{If} \\ \surd\ \textit{When} \end{array} \right\} \begin{array}{l} \text{we were naughty, we were} \\ \text{sent to bed with no supper.} \end{array} \qquad (6.47)$$

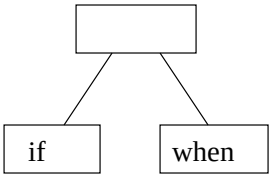

Figure 6.14: ACTUAL and HYPOTHETICAL Phrases

$$\text{You can stay up } \left\{ \begin{array}{l} \checkmark \textit{if} \\ \# \textit{ when} \end{array} \right\} \text{ you don't squabble.} \tag{6.48}$$

$$\left\{ \begin{array}{l} \textit{When} \\ \# \textit{ If} \end{array} \right\} \text{ Mary gets home, ask her to call me.} \tag{6.49}$$

In all of these examples, the suitability of the different cue phrases seems determined by whether the subordinate span (A) is known or unknown. A feature called MODAL STATUS can thus be proposed, with alternative values ACTUAL and HYPOTHETICAL.

The MODAL STATUS feature interacts productively with a number of other features. We have already talked about 'conditional' phrases in a number of other places; for instance in connection with the POLARITY feature (*if* versus *even if*) and the ANCHOR feature (for which *if* is undefined). The question is now how to define it.

A simple idea would be to identify the contexts where the anchor A is known as ACTUAL, and those where it is not known as HYPOTHETICAL. However, there are a number of problems with this approach. Consider the case of Text 6.49. It cannot be that the writer actually *knows* that Mary gets home in this example, as it is an event in the future. So what is it which makes *if* and *when* different in this case? It is plausible to suggest that *when* is sanctioned by the writer's ability to *predict* Mary's return before it has happened. It would thus be preferable to define the feature in terms of knowledge (or lack of knowledge) of the *cause* of A rather than of A itself. This being the case, we can thus make use of the variable *Precond* introduced by the PRESUPPOSITIONALITY feature in Section 6.2.7, which for PRESUPPOSED relations such as the above, represents the cause of A. The definition of MODAL STATUS would then be as follows:

MODAL STATUS

ACTUAL: *Precond* is known by the protagonist/writer.

HYPOTHETICAL: *Precond* is not known by the protagonist/writer.

It should also be possible to talk about the HYPOTHETICAL/ACTUAL distinction for NON-PRESUPPOSED phrases. Consider the following examples:

$$\text{Bob piled up the boxes underneath the bananas. } \left\{ \begin{array}{l} \textit{Then} \\ \# \textit{ In that case,} \end{array} \right\} \text{ he was able to reach them easily.} \tag{6.50}$$

$$\text{Bob might try piling up the boxes underneath the bananas. } \left\{ \begin{array}{l} \textit{Then} \\ \checkmark \textit{ In that case,} \end{array} \right\} \text{ he'll be able to reach them easily.} \tag{6.51}$$

Both *then* and *in that case* are NON-PRESUPPOSED in the above contexts, so *Precond* is identified with the anchor clause, which is the first clause in each case. *Then* seems to be undefined for MODAL STATUS, being appropriate in both contexts; but *in that case* seems to require an anchor which is unknown.

RESULT-DRIVEN HYPOTHETICAL Relations

A possible instance of a HYPOTHETICAL RESULT-DRIVEN phrase is *in case*. Consider the following text:

$$\text{Bill tidied the house, } \left\{ \begin{array}{l} \textit{in case} \\ \# \ \textit{because} \end{array} \right\} \text{ his parents came home early.} \tag{6.52}$$

This text requires an inference about one of Bill's goals; namely that he does not want his parents to come home early and find that the house is not tidy. It is not certain that his parents will come home early—hence the relation is HYPOTHETICAL—but it is sufficiently likely to make Bill tidy the house. We can therefore assume that the two related spans *Bill tidied the house* and *his parents came home early* are both part of the left-hand side of some rule whose right-hand side is a state of affairs desired by Bill, such as 'Bill's parents are not angry'. Note that the ACTUAL phrase *because* is inappropriate for signalling this relation.

6.2.9 SEMANTIC and PRAGMATIC Relations Revisited

This final section addresses a number of issues and problems that arise in connection with the SEMANTIC/PRAGMATIC distinction. These will first be outlined, and then some suggestions for a solution will be put forward.

A first observation is that the present definitions of SEMANTIC and PRAGMATIC bring together two quite different ideas: on the one hand, the issue of whether A and C represent the propositional content of the related clauses (SEMANTIC) or their intended effects (PRAGMATIC); and on the other, that of whether the intended effect of the whole relation is that the reader *believe* a relation between two propositions (SEMANTIC) or that a relation between two propositions is *actually the case* (PRAGMATIC). The latter distinction seems to define whether the relation is part of a description or narrative, which the reader accepts without question, or whether it takes place in the real world, where the writer's goals are not just communicative. There seems no reason a priori why these two dimensions should be related.

Indeed, it is not hard to find examples of prototypically PRAGMATIC cue phrases in purely narrative discourse. For instance, *furthermore* and *then again* can both feature in 'free indirect speech', where an agent's thought processes are being described:

(6.53) Sally couldn't decide who would win the match that evening. Spurs were at home; *furthermore*, they were on good form. *Then again*, they were playing the league champions...

But the intended effect of these relations is surely still descriptive rather than persuasive.

Conversely, going by the current definitions, it is hard to find any phrases which are purely SEMANTIC. A temporal phrase like *after this* is a plausible candidate, but such phrases can be used to link imperative clauses, which currently count as PRAGMATIC:

(6.54) Sweep the floors. *After this*, tidy the cupboards.

At the same time, *furthermore* and *after this* are exclusive phrases, as was noted in the very first extract to be presented from the taxonomy in Section 4.6.2. The motivating examples are reproduced below:

Television is bad for us. It kills creativity; { *furthermore,* / # *after this* } it promotes an unhealthy kind of 'crowd mentality'. (6.55)

Bob set about cleaning the house. He swept the floors; { *after this,* / # *furthermore,* } he tidied the cupboards. (6.56)

We need to find a feature to account for this exclusivity. Clearly, the present definitions of SEMANTIC and PRAGMATIC are not able to.

A suggestion for remedying the above problems comes from considering another of the distinctions motivated from the taxonomy; that between CAUSE-DRIVEN and RESULT-DRIVEN phrases (see Section 6.2.5). The point is that some of the work being done by the current SEMANTIC/PRAGMATIC distinction appears to overlap with work being done by this feature. The definition of PRAGMATIC talks about the 'intended effects' of utterances S_A and S_C; the definition of RESULT-DRIVEN also talks about an agent's intention, and an action or actions which are caused by this intention. It is interesting to speculate that the writer's utterances could be represented as goal-driven actions, just as are the actions of the agents which the writer talks about. To take just one piece of evidence for this line of reasoning: the notion of RESULT-DRIVEN NEGATIVE POLARITY relations seems to find useful application in analysing the kind of *but* which occurs in dialogues, between two speakers. Consider the following exchange:

(6.57) A: *Go to bed.*
B: *But I haven't done my homework yet...*

Just as with other RESULT-DRIVEN relations, it is odd to analyse the second span as 'violating an expectation' set up by the first span, and preferable to think of it as defeating a *goal* underlying the first span. In this case, however, the goal is that of the first speaker, rather than that of a protagonist being described in the text. This idea will not be pursued here, as the present work is not concerned with inter-speaker relations. But it is an interesting thought that such relations might eventually be netted in by the theory.

6.3 Summary of Features Motivated

Now that a number of features have been individually motivated, we can begin to put them together to build up the complex definitions required for cue phrases and relations. Until now, feature definitions have been presented individually. The complete

Name of Feature	Possible Values	
SOURCE OF COHERENCE	SEMANTIC	PRAGMATIC
ANCHOR	CAUSE-DRIVEN	RESULT-DRIVEN
PATTERN OF INSTANTIATION	UNILATERAL	BILATERAL
FOCUS OF POLARITY	ANCHOR-BASED	COUNTERPART-BASED
POLARITY	NEGATIVE	POSITIVE
PRESUPPOSITIONALITY	PRESUPPOSED	NON-PRESUPPOSED
MODAL STATUS	HYPOTHETICAL	ACTUAL
RULE TYPE	CAUSAL	INDUCTIVE

Figure 6.15: The Features So Far Motivated, and Their Alternative Values

set of definitions for the features so far motivated can be found in Appendix C; here, the interactions between the definitions can be more easily appreciated. For the moment, a summary of the eight features and their possible values is given in Figure 6.15.

Figure 6.16 presents a larger extract from the taxonomy, which draws together a selection of the phrases which have so far been discussed, labelled with the feature values which have so far been established. The feature-theoretic interpretation of the taxonomy is clearly illustrated here: exclusive phrases are defined for alternative values of at least one feature; hyponyms inherit all of the feature values associated with their hypernyms and are defined for other features in addition, and so on.

The diagram in Figure 6.16 is complex: it documents many of the substitutability relationships shown in the smaller diagrams in Section 6.2, as well as many relationships between phrases which appeared in different diagrams. The figure divides roughly into four exclusive groups of phrases:

- POSITIVE POLARITY CAUSAL ACTUAL relations (dominated by the phrase *so*);
- NEGATIVE POLARITY ACTUAL relations (dominated by the phrase *but*);
- HYPOTHETICAL relations (involving *if* and *even if*);
- POSITIVE POLARITY INDUCTIVE relations (the single phrase *furthermore*).

Much of the complexity in the diagram is due to high-level phrases such as *and* and *while*, which cut across these divisions.[13]

The features with which the phrases are labelled are unlikely yet to be sufficient as definitions, as many additional features have still to be motivated from the taxonomy. Even in this diagram—still just a small portion of the overall taxonomy—there remain relationships which are not yet explained by the features provided. (For instance, the contingent substitutability between *while* and *whereas* remains unexplained. So does

[13] The diagram in Figure 6.16 is already quite difficult to read—clearly, extending it to encompass *all* the phrases in the corpus would soon lead to problems. It is for this reason that the complete taxonomy in Appendix B is divided into a number of separate diagrams when it is presented; see Section 4.6.4.

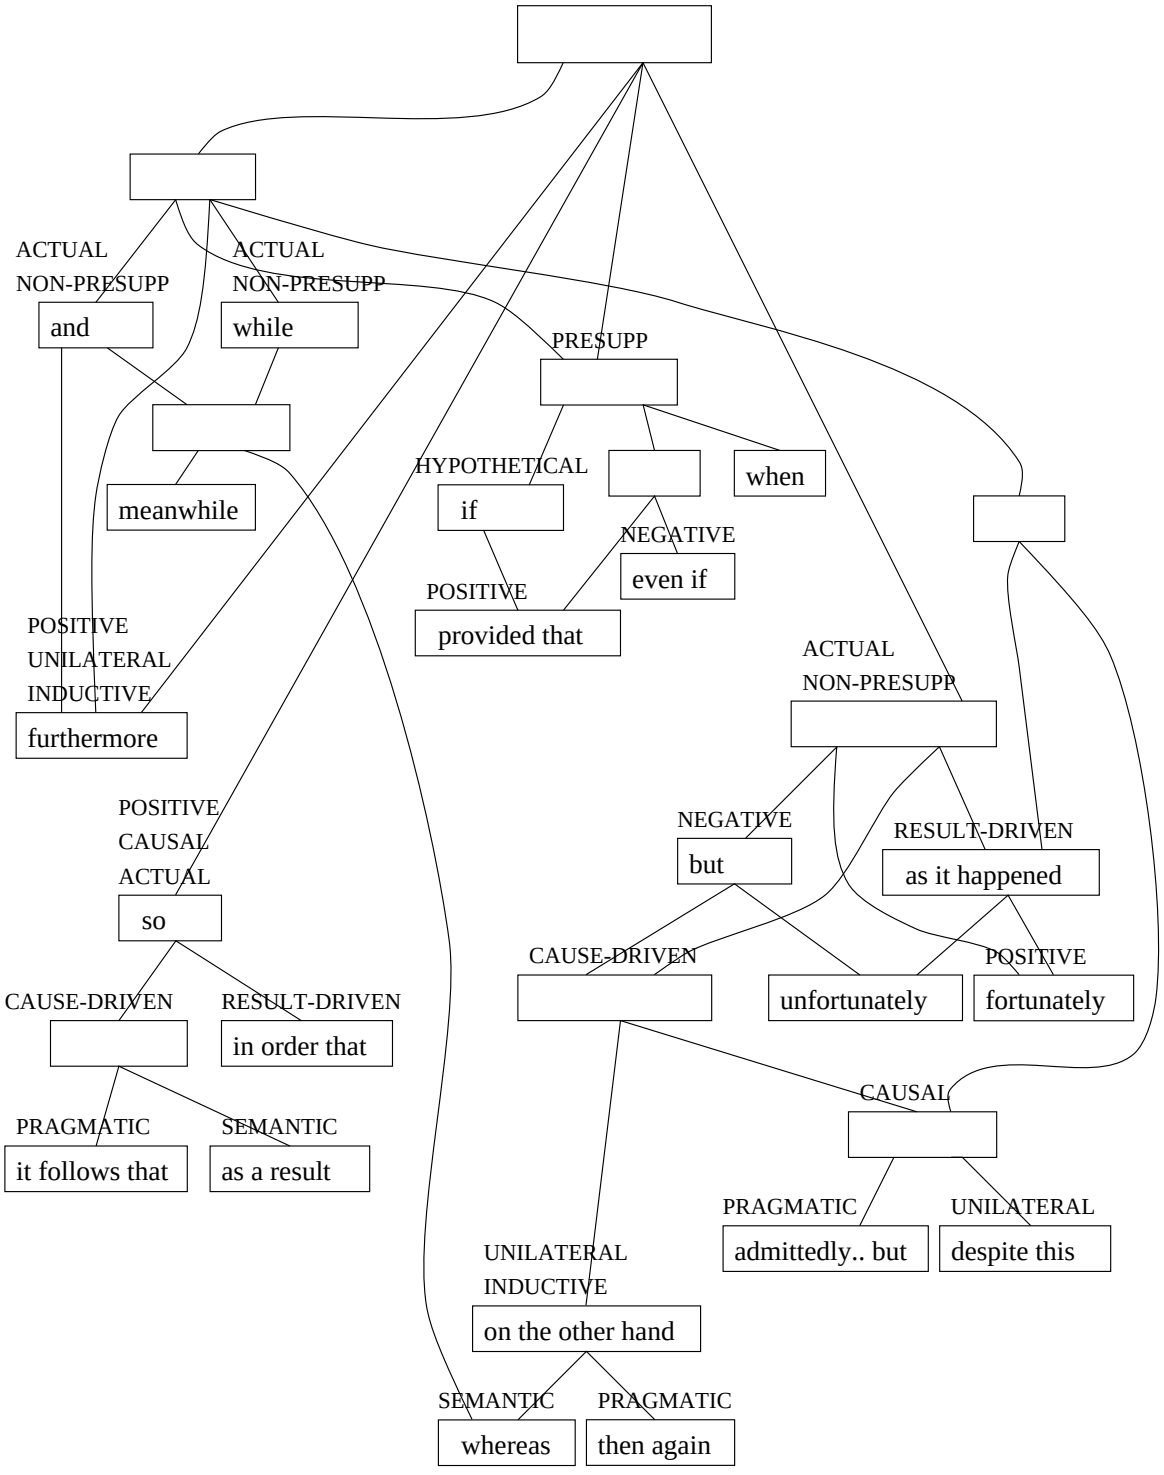

Figure 6.16: Summary of Motivated Features: A Labelled Extract from the Taxonomy

the exclusivity between *and* and *it follows that*.) However, the current set of features already give good approximate definitions in most cases, and at very least serve to indicate the flavour of the definitions which will eventually be reached.

A larger set of composite definitions is given in Appendix D, this time in tabular form. Again, the definitions are not always complete, but they already begin to provide a reasonable account of the variations between the phrases.

6.4 The Mapping between Cue Phrases and Relations

It is useful to sum up what has happened so far in the theoretical interpretation of the taxonomy. In Chapter 5, it was noted that the taxonomy lends itself well to a conception of relations as collections of independent features, and some general principles governing the motivation of features were put forward. In the first part of this chapter, a number of features were systematically motivated by considering small excerpts from the taxonomy one by one. In this section, the final step in the methodology is considered: how the features should be combined to give a set of coherence relation definitions.

Until now, the assumption has been that there will be a one-to-one correspondence between relations and the cue phrases at the *leaves* of the taxonomy. Cue phrases higher up in the taxonomy do not correspond to 'more general' relations, but are used to signal some components of a relation in circumstances where the reader is able to infer the others from context and background knowledge. We begin, therefore, by looking at the cue phrases at the leaves of the taxonomy.

6.4.1 An Uneven Distribution at the Leaves of the Taxonomy

It should be clear by now that even the most specific 'leaf-level' cue phrases in the taxonomy may still be undefined with regard to certain features. For instance, as noted in Section 6.2.5, *if* is undefined for the ANCHOR feature; but there are no phrases below *if* to distinguish between the alternative values of this feature (CAUSE-DRIVEN and RESULT-DRIVEN.) In fact, it may be that when all the features necessary to describe the taxonomy are found, few if any leaf-level phrases will be defined for every one of them.

It might also be that when the possible combinations of feature values are investigated more extensively, some combinations will be found which are not signalled by *any* cue phrase, even a general one which requires feature values to be inferred. Again, this remains an open question.

In short, the phrases at the leaves of the taxonomy are unevenly distributed over the space of possible feature value combinations. Some phrases, being undefined for various features, can be used for a wide range of possible combinations; there may also be possible combinations which are not signalled by any cue phrase. The mapping between leaf-level phrases and possible feature value combinations is thus many-to-one, and incomplete.

This is in no way a disadvantage of the theory. There is no reason to expect a perfect mapping between leaf-level phrases and possible value combinations, as was stressed in Section 5.2.4. And relations are all abstractions, so the idea of some features being undefined in a given relation is quite understandable.

6.4.2 Relations at Different Levels of Abstraction

A trickier problem is raised by the patterns of substitutability illustrated in Figure 6.17. Consider first the extract in 6.17 (i). In this case, the hypernym is *after*, and the

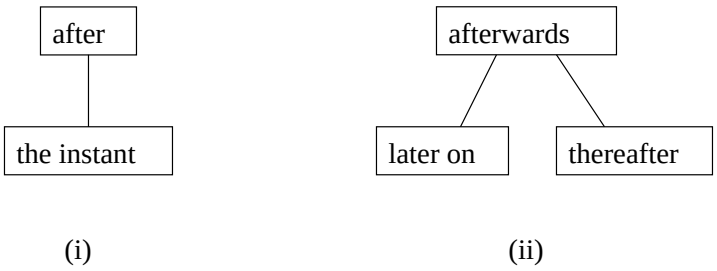

Figure 6.17: Some 'Problem' Extracts from the Taxonomy

hyponym is *the instant*, which in the example we can take to be a leaf node. According to the current hypothesis, *the instant*, being a leaf node, signals all the features of some particular relation: *after* is used to signal this same relation if the extra information carried by *the instant* is easily inferrable from context. The odd thing is, that there should *never* be any need to use *the instant* to signal the relation in question. *After* can only ever mark one possible relation; that is the relation marked by *the instant*. So why does the phrase *the instant* exist at all? If *after* could signal two different relations, then *the instant*'s existence would be understandable: it would be used whenever the information to distinguish between these two relations was not inferrable (in which case *after* would not identify the relation). But there is just one relation in the present case: we know this because *the instant* has no sister phrases, and the whole rationale of the taxonomy is that a cue phrase will exist for *each relation*.

So we have a problem in trying to explain this type of pattern in the taxonomy. In fact, we have exactly the same type of problem with the pattern in 6.17 (ii), where the leaf nodes *later on* and *thereafter* are contingently intersubstitutable, and have a common hypernym *afterwards*. It might be thought that in this case there are two different relations, picked out uniquely by *later on* and *thereafter* respectively. In this case, *afterwards* is able to signal either relation, provided that the information needed to distinguish between them is inferrable from context. However, the phrases *later on* and *thereafter* could not be used to pick out their respective relations uniquely: they are contingently intersubstitutable, which means that in some contexts, they can be substituted for one another. In these contexts, there are no cue phrases for identifying either relation exclusively; again, this goes against the rationale of the whole taxonomy. An alternative suggestion is that *later on* and *thereafter* signal the same relation but identify different subsets of its features, making them suitable for use in different contexts, when different features will be inferrable. On this hypothesis, *afterwards* is to be used to signal this same relation, in contexts where *all* those features unspecified

by either *later on* or *thereafter* are safely inferrable. However, this makes the situation identical to that in 6.17 (i): if there is just a single relation, there is no need for the phrases *later on* and *thereafter* at all.

Patterns such as those illustrated in Figure 6.17 are common in the taxonomy, and some explanation is called for. One explanation is that the relations which cue phrases signal are not all associated with the cue phrases at the leaves of the taxonomy: some are associated with cue phrases at higher levels. In other words, *some relations are more abstract than others*. This idea permits an explanation of patterns like those in Figure 6.17: in each case, the hypernymic phrases signal relations at a higher level of abstraction, subsuming those relations signalled by the hyponymic phrases. Our general assumption that 'for every relation there exists a cue phrase which uniquely identifies it' now actually *requires* such hypernymic phrases to exist.

To give an example: in Figure 6.17 (i) we can posit two relations; one signalled by the leaf node phrase *the instant*, and another (this one more abstract) signalled by the hyponymic phrase *after*. *The instant* uniquely identifies one particular relation. *After*, at another level of abstraction, picks out another relation: there is no other cue phrase which can pick out all the occurrences of this relation.

It might be thought that this scenario permits a different kind of ambiguity: if a writer uses the hypernym *after*, in a situation where the relation associated with *the instant* can be inferred from context, how would the reader know whether the more abstract or the less abstract relation is intended? Both would appear to be possible. However, more careful consideration about what an 'abstract relation' is will show that this kind of 'ambiguity' presents no real problem for the reader.

Levels of Abstraction in Text Processing

In Chapter 3, when the conception of relations as cognitive constructs was being outlined, Rosch's notion of the **basic level** of categorisation—the level of abstraction at which the trade-off between usefulness and generality is optimised—was invoked. Rosch suggested that humans would operate at this level when performing certain tasks with a heavy information-processing load. I argued that tasks like reading and writing are likely to be mediated by constructs at a similar level of abstraction; and proposed to think of coherence relations as modelling such constructs.

The hypothesis now under consideration, that 'relations exist at different levels of abstraction', is quite consistent with this conception of relations. A writer's information-processing load can vary considerably, reflecting factors such as time pressure or the difficulty of the task at hand. It is likely that as the information processing load varies, so too does the optimal level of abstraction at which the writer should operate: the higher the load, the greater the degree of abstraction. To take an example: if the writer has to produce a text in a hurry, or if the elements of the text plan are hard to manipulate, then maybe the writer's planning operators should be less tightly defined. Likewise, if the writer's time constraints are relaxed, then more features of a relation can be taken into account.

Many researchers have postulated relations at different levels of abstraction (eg Maier

and Hovy (1991), Hovy *et al* (1992), Mann and Thompson (1988)). But no-one has yet come up with a clear reason for doing this. However, if we accept the argument on which the present project is founded—that for any relation that people actually use, there will exist a cue phrase which uniquely marks it—then we have a genuine motivation for postulating the existence of relations at different levels of abstraction. And by drawing on the psychological conception of relations developed in Chapter 3, we have a way to understand why different levels of abstraction should ever be used.

Motivating Relations from High-Level Cue Phrases

The question of *when* a high-level cue phrase motivates a 'more abstract' relation is still to some extent an open one. Given patterns like those in Figure 6.17, we have definite evidence for different levels of abstraction. But consider a pattern such as that in Figure 6.18, where two sister leaf nodes are exclusive. Here, two possibilities can

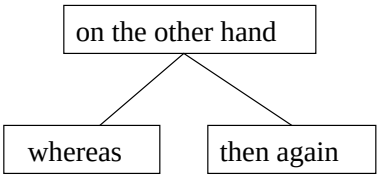

Figure 6.18: Exclusive Sister Leaf Nodes

be discerned. The original story was that *on the other hand* can signal two different relations, which are identified uniquely by *whereas* and *then again*. But, in the light of the current discussion, it might also be that *on the other hand* corresponds directly to a 'more abstract' relation. The answer to this question is a matter for further research. For the moment, I will assume that abstract relations are only motivated where patterns such as those in Figure 6.17 are found. In cases such as that in Figure 6.18, no abstract relation will be postulated.

6.5 Summary

This chapter has presented the beginnings of a feature-theoretical description of the taxonomy of cue phrases. Eight two-valued features have been proposed to account for various dimensions of variation amongst the phrases in the taxonomy. Definitions of the features are summarised in Appendix C, and a preliminary table of relation definitions, expressed in terms of these features, is given in Appendix D.

There remains much work to be done, of course. The taxonomy contains a vast amount of substitutability data, and only a small portion has been looked at. The features so far motivated only cover some portions of the taxonomy. Moreover, feature definitions have been expressed at a relatively informal level: more precise definitions would certainly be preferable, both from a theoretical point of view and as a precursor to implementation. Finally, a great deal more effort is needed to investigate all the possible combinations

of the features presented here. Again, this is an area where much further work is required.

All the same, the preliminary conclusions reported here will be useful as the foundation for a more complete account of the phrases in the taxonomy. At very least, they serve as a convincing demonstration of the utility of the methodology being proposed. In particular, they provide good support for the decision taken in Chapter 5, to look for a completely productive set of features. The directions in which further progress can be made are thus quite clearly indicated by the ideas developed in this chapter.