Chapter 7

An Evaluation of the Substitution Methodology

This chapter contains a discussion of some of the potential problems with the substitution methodology proposed in Chapters 3 to 6.

A first problem relates to the substitutability test itself: there are some cases where it seems unable to provide the kind of clearcut data necessary to motivate a feature-based account of relations. This problem will be discussed in Section 7.1; as a solution, an empirical test is suggested which promises a clearer picture about these cases.

Two other problems concern the set of relations eventually motivated by the substitutability test—there are many respects in which the set of relations seems unable to provide a complete coverage of texts. For one thing, it is not hard to find pairs of sentences or clauses in a coherent text for which no cue phrase at all seems appropriate. Since our set of relations is based on the set of cue phrases, such contexts are beyond the scope of the theory. This problem is considered in Section 7.2; a solution is proposed by appealing to the concept of **focus** as better suited for an explanation of these contexts.

Finally, in Section 7.3, the question of relations between large sections of text is raised. It might be thought that cue phrases are only suitable for signalling relations at a low level of hierarchy in a text. However, it is argued that the presence of anaphoric elements in many cue phrases allows them in practice to signal relations between very large sections of text. A discussion of such 'high-level' cue phrases proves to be of relevance to the issue of the interaction of theories of relations and of focus.

7.1 Limitations of the Substitutability Test: The Case of Presentational Sequences

Some relationships in the taxonomy of cue phrases seem more clearcut than others. No-one, for instance, would claim that the phrases *nevertheless* and *on the grounds that* are anything other than exclusive. However, in other cases, our intuitions seem to be less certain. This is particularly so for decisions about phrases signalling what some researchers have called ‘presentational sequence’—phrases such as *moreover*, *for another thing* and *furthermore*. For instance, consider the following examples:

The Tories have done a terrible job in government over the last decade. Their early policies led to huge unemployment. Their increasing isolationism over Europe wasted a precious opportunity to gain international influence.	$\left\{ \begin{array}{l} \textit{Furthermore,} \\ \textit{? Later,} \end{array} \right\}$	they reduced Britain to a laughing stock with their ‘back to basics’ campaign.	(7.1)

Let us review the sequence of governmental blunders during the past decade. Their early policies led to huge unemployment. Their increasing isolationism over Europe wasted a precious opportunity to gain international influence.	$\left\{ \begin{array}{l} \textit{Later,} \\ \textit{? Furthermore,} \end{array} \right\}$	they reduced Britain to a laughing stock with their ‘back to basics’ campaign.	(7.2)

The taxonomy represents *furthermore* and *later* as exclusive, based on examples such as these. The claim is that in Example 7.1, replacing *furthermore* with *later* changes the text from an argument to a temporal sequence; and that in Example 7.2, replacing *later* with *furthermore* changes the text from a temporal sequence into an argument. Indeed, there does seem to be some kind of difference between the texts. But nevertheless, it is still possible to imagine a writer replacing one phrase by the other. Certainly, it is easier to imagine this than to imagine replacing *nevertheless* with *on the grounds that.*

Another problematic case is given in Example 7.3; here, it is unclear whether or not *whereas* can be replaced by *furthermore*.

It’s crazy to keep Bill and Bob in their present positions: we should swap them. Bill is a better player in attack;	$\left\{ \begin{array}{l} \textit{whereas} \\ \textit{? furthermore,} \end{array} \right\}$	Bob is a better player in defence.	(7.3)

The original phrase seems to imply that there is only one reason why Bill and Bob should be swapped, namely that they are better suited to each other’s positions. *Furthermore*, on the other hand, suggests that there are two independent reasons why Bill and Bob should be swapped. Again, the difference seems to be one of emphasis—ultimately, the same information is derivable from both texts—but there is nonetheless a difference. Should this difference be treated as theoretically significant or not?

The problem is that in order to answer this question, the tester has to stop being a normal reader/writer, and start being a ‘discourse analyst’. Should the alternative versions of the text be given the same analysis, or different ones? Such questions require the kind of ‘post-theoretical’ intuitions which were called into question at the beginning of the thesis in Section 1.3.2—we cannot be sure if they are reliable, and there seems to be no way of resolving any disagreements which might occur. Ideally, therefore, the test for substitutability should not be relied on in such cases.

It is not surprising that disagreements do in fact occur as to how to analyse texts like those above. The disagreements turn on alternative theories of **span structure**—see Section 2.2.3. Two alternative positions can be identified; these are set out in the following two sections.

7.1.1 The Simultaneous Representation Hypothesis

One line of thought is that *later* and *furthermore* do not signal contradictory information; they simply make different unconnected features of the text explicit. On this story, the two phrases would be contingently substitutable. The relation between the two spans in Examples 7.1 and 7.2 contains both temporal information and information about the structure of an argument, and the phrases are used to signal one or other of these two types of information. Likewise, the relation between the two spans in Example 7.3 contains argumentative information as well as information about a contrast between the two premises; so *whereas* and *furthermore* should also be considered as contingently substitutable. The analysis for these texts would look something like that given in Figure 7.1.

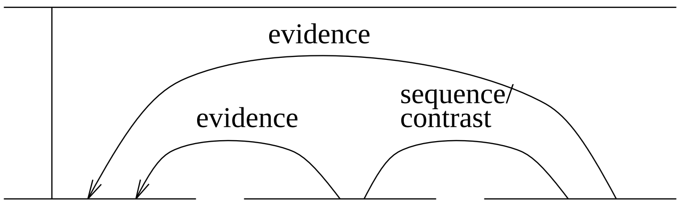

Figure 7.1: The Simultaneous Representation Hypothesis for Texts 7.1–7.3

Such an analysis seems to be espoused by Halliday and Hasan (1976). The following example is referred to:

(7.4) Next, previously to this he had already offered to resign.

The apparent contradiction in this sentence is explained by noting that two different kinds of sequence are being signalled; one INTERNAL (argumentative) and one EXTERNAL (temporal). Martin (1992) also suggests that INTERNAL and EXTERNAL relations can apply simultaneously. Maier (1993) perhaps comes closest to espousing the view outlined above, giving analyses very similar to that in Figure 7.1. For Maier, *furthermore* signals a TEXTUAL relation, and *whereas* and *later* signal *ideational* relations, and these two types of relation can co-occur in a text.

7.1.2 The Dominant Representation Hypothesis

The alternative to the simultaneous representation hypothesis can be referred to as the **dominant representation hypothesis**. According to this hypothesis, there is a significant difference between a text marked with *furthermore* and one marked with *whereas* or *later*, which makes both these latter phrases exclusive with the former.

According to this view, the texts with *whereas* or *later* should be analysed as in the diagram on the left of Figure 7.2, as containing a single complex premise, while the texts with *furthermore* should be analysed as in the diagram on the right, as containing two independent premises.

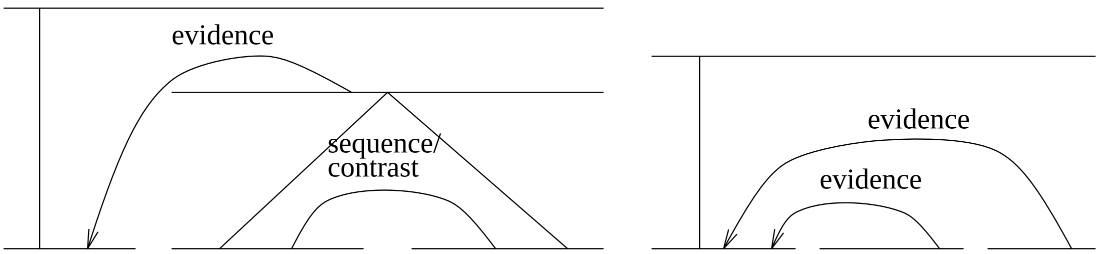

Figure 7.2: The Dominant Representation Hypothesis for Texts 7.1–7.3

This is the view taken by Mann and Thompson (1988): their schemas do not permit the type of diagram shown above. It is also the view taken in this thesis. It will be argued here that although the above texts *can* be analysed using simultaneous relations, thinking of relations as modelling psychological constructs makes this analysis implausible. It seems likely that for a reader or a writer, one of the relations will always be more relevant than the other, and will dominate a representation of the text. This is not to say that when one reading dominates, the alternative reading is *not available at all* to the reader. Clearly, both the argumentative information and the temporal (or contrastive) information can be inferred from the text if it is considered with any serious attention. It is only suggested that this does not typically happen in what we might call 'normal' reading or writing.

This hypothesis can be supported on grounds of cognitive economy. Consider the case of a writer planning the text in Example 7.1. The writer's intention is to persuade the reader that the Tories have done a terrible job over the past decade. Her tactic for achieving this goal is to state a series of facts, each of which will individually push the reader towards this conclusion. The writer's task is then to determine which facts fall into this category: in other words, she needs to find a collection of facts which stand in the relationship of PRAGMATIC ADDITIVITY with each other. Once she has found such facts, there is no need for her to ascertain anything else in order to pursue her strategy. In particular, it is not important to ask about the temporal ordering of the facts: communicating this information will not contribute towards the overall persuasive goal.

Now consider Example 7.2. Here, the overall goal (as set out in the first sentence) is to present to the reader a sequence of events. In order for the writer to achieve this goal, she clearly needs to verify the temporal order of the events. However, it will not be necessary to consider whether or not the events act as premises for some common conclusion: this question is just not relevant to the writer's goals.

Given these facts, it is plausible to suggest that writers concentrate on either the temporal relation or the argumentative relation when constructing text. The reader of a text can be expected to have a similar bias one way or the other, given that he is really trying to make a coherent representation of the text, which ultimately includes a representation of the writer's intentions.

It might be objected to this account that the writer may not have a single overriding intention behind producing a piece of text. It is conceivable that she has two separate intentions; one to persuade the reader of a particular fact, and one to tell the reader about a sequence of events. It thus makes sense to have a strategy for killing the two birds with one stone in a single piece of text: in order to analyse this text, it will be necessary to represent both argumentative and temporal relations between its spans.

Of course, it *is* possible to have two separate intentions simultaneously. To take an unrelated example: I can have an intention to buy an ice cream, and at the same time one to write a letter to my parents. However, there is no reason to expect these two intentions to co-occur with particular regularity; and therefore no reason to expect that any special strategy will have evolved to deal with them both in the same action. Likewise, there is no reason to expect a regular co-occurrence between the intention to convey a temporal sequence of events and the intention to argue for a given conclusion. So, if we are interested in modelling the specialised strategies a writer can make use of for conveying intentions in a text, we are unlikely to need to represent a special mechanism whereby both these intentions can be achieved simultaneously.

7.1.3 An Experimental Design for Testing the two Hypotheses

Whatever the arguments put forward for the alternative hypotheses, it is unlikely that either will be conclusive by itself. However, other empirical means for deciding between them may be more promising. In this section, an experiment is outlined which forces the two hypotheses to make different measurable predictions.

The Form of the Experiment

We will begin by considering the case of Texts 7.1 and 7.2. The dominant representation hypothesis suggests that readers and writers concentrate on the argumentative relation *or* on the temporal relation. The simultaneous representation hypothesis suggests that readers and writers should be able to concentrate on both relations at the same time.

To decide between these alternatives, a experiment is proposed in which subjects read a text containing both temporal and argumentative information, and then answer questions which relate specifically to one type of information or the other. Different reading conditions can be created by varying the cue phrase used in the text: *furthermore* emphasises the argumentative information, while *later* emphasises the temporal information; finally, an unmarked version of the text is neutral between the two.

The text used in the earlier examples can be adapted for this purpose. A neutral introductory sentence is used, so that both *furthermore* and *later* result in coherent texts:

The sequence of governmental blunders during the past decade makes for interesting reading. Their early policies led to huge unemployment. Their increasing isolationism over Europe wasted a precious opportunity to gain international influence. $\left\{ \begin{array}{l} \checkmark \textit{Later,} \\ \checkmark \textit{Furthermore,} \end{array} \right\}$ they reduced Britain to a laughing stock with their 'back to basics' campaign. (7.5)

When the readers have read the text, they are asked to verify a number of statements, including the following pair:

S1 The Tories were isolationists over Europe before 'back to basics' was launched.

S2 The Tories can be criticised on several grounds.

The subjects' reaction times will be monitored, to give an indication of how easy they find it to decide about these statements. We can then make comparisons between the subjects' performance after reading one of the alternative marked texts, and their performance after reading the unmarked text.

Predictions of the Two Hypotheses

It seems reasonable to predict that the text using *later* will facilitate the verification of S1 compared to the unmarked text; given that it makes explicit the temporal relation between the two propositions in question. Likewise, it seems reasonable to predict that the text with *furthermore* will help readers verify S2 more than the unmarked text. The interesting question is how the *later* text influences the decision about S2, and how the *furthermore* text influences the decision about S1. The simultaneous representations model in these cases predicts the null hypothesis; namely that

- verifying S2 will be just as easy after reading the *later* text as after reading the unmarked text;
- verifying S1 will be just as easy after reading the *furthermore* text as after reading the unmarked text.

These predictions follow from the claim that both the temporal and the argumentative relations in the original text should be represented: they are independent aspects of its meaning. By marking just one aspect, we should in no way be inhibiting the other, which should be just as clear as in the unmarked case.

Different predictions are made by the dominant relation hypothesis: according to this, the explicit signalling of the temporal information will *inhibit* the representation of the argumentative information, and vice versa. Thus

- verifying S2 will be *harder* after reading the *later* text than after reading the unmarked text;

- verifying S1 will be *harder* after reading the *furthermore* text than after reading the unmarked text.

If these predictions were borne out, there would be a good case for arguing that temporal and argumentative relations are thought of by the reader as real alternatives to each other, and that they are not expected to co-occur. If the reader finds an explicit temporal marker, it actually prejudices him against looking for an argumentative reading; while if he finds an argumentative marker, it prejudices him against looking for a temporal reading. For this reason, it seems legitimate to represent the two markers as exclusive in the taxonomy: they really provide the reader with contradictory information.

Testing the Hypotheses for *Furthermore* and *Whereas*

A similar experiment can be set up to test the predictions made by the two hypotheses about Example 7.3. This text can be given in three conditions; with *whereas*, with *furthermore*, and with the neutral null cue phrase:

$$\begin{array}{r} \text{It's crazy to keep Bill and Bob in their} \\ \text{present positions: we should swap them.} \\ \text{Bill is a better player in attack;} \end{array} \left\{ \begin{array}{l} ?\ \textit{whereas} \\ ?\ \textit{furthermore,} \\ ? \end{array} \right\} \begin{array}{l} \text{Bob is a} \\ \text{better player} \\ \text{in defence.} \end{array} \qquad (7.6)$$

The statements to be verified in the subsequent decision task are now as follows:

S1 Bill and Bob are different.

S2 There are two reasons why we should swap Bill and Bob.

Again, the dominant representation hypothesis would predict that the text with *whereas* would slow down the subjects' response to S2 as compared with the null marker, while the text with *furthermore* would slow down the response to S1. The simultaneous representation hypothesis would not predict slowed down responses.

A Control Condition: Contingently Substitutable Phrases

In the case of two phrases which are 'genuinely' contingently substitutable, we would not expect to find one phrase inhibiting the reading signalled by the other. We could thus run a control condition, using clearly contingently substitutable phrases like *once* and *as soon as* (see Figure 7.3).

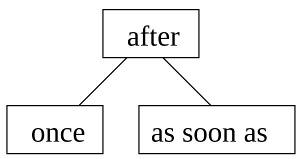

Figure 7.3: Contingently Substitutable Phrases

Both phrases signal a temporal succession between the two related spans, but *once* signals in addition that they are related as cause and effect, while *as soon as* signals in addition that the temporal succession is very rapid. As the following two examples show, *as soon as* can be used in non-causal contexts, and *once* can also be used where the temporal succession is not immediate:

$$\begin{array}{r} \text{We got the fright of our} \\ \text{lives last night coming back} \\ \text{from Sue's party.} \end{array} \left\{ \begin{array}{l} \textit{As soon as} \\ \#\ \textit{Once} \\ \checkmark\ \textit{After} \end{array} \right\} \begin{array}{l} \text{we had stepped out of the} \\ \text{house, we heard a huge clap} \\ \text{of thunder.} \end{array} \tag{7.7}$$

$$\left\{ \begin{array}{l} \textit{Once} \\ \#\ \textit{As soon as} \\ \checkmark\ \textit{After} \end{array} \right\} \text{the war ended, the country's economy gradually improved.} \tag{7.8}$$

Note that the phrase *after*, being a common hyponym, can be used in both contexts: it signals nothing in addition to the temporal succession relationship.

In a context describing an immediate causal succession, a writer could use any one of the three phrases:

$$\begin{array}{r} \text{The thieves didn't have} \\ \text{long to search Jones'} \\ \text{apartment.} \end{array} \left\{ \begin{array}{l} \textit{After} \\ \checkmark\ \textit{As soon as} \\ \checkmark\ \textit{Once} \end{array} \right\} \begin{array}{l} \text{he had fallen asleep, they} \\ \text{set to work as fast as they} \\ \text{could.} \end{array} \tag{7.9}$$

In such a case, there are no niggling suspicions that by substituting *as soon as* for *once* (or vice versa) we are somehow contradicting a presupposition set up in the original text. It seems legitimate to talk about 'causality' and 'immediacy' as independent features of the reader's model of the text, which can both be represented simultaneously. If this is indeed the case, we would not expect that *once* actively inhibits the information about immediacy, or that *as soon as* inhibits the information about causality. Again, we could create decision tasks for subjects to test this prediction; for example about the following alternative statements:

S1 The thieves had to wait for Jones to fall asleep before they could search his apartment.

S2 The thieves searched Jones' apartment immediately after he fell asleep.

Here, our predictions conform to the null hypothesis—that verifying S2 will be no harder after reading the *once* text than after reading the neutral *after* text; and that verifying S1 will be no harder after reading the *as soon as* text than after reading the neutral text.

7.1.4 A Revised Role for The Substitutability Test

This section has presented an empirical means of examining the substitutability relationship between two cue phrases. In some cases, this seems to provide a better means of investigating the relationship than the test for substitutability; it will be interesting to see what results are produced in these cases.

However, the test for substitutability cannot be dispensed with in favour of this new method. There are still many cases where the test for substitutability provides clear information. More importantly, the experiment only works where the alternative cue phrases both lead to coherent texts—most of the time, this is not the case. For instance, it would be impossible to find a context neutral between the phrases *because* and *although*, where the relations marked by both phrases are inferrable from the same context.

7.2 Issues of Descriptive Adequacy: The Problems of ELABORATION and BACKGROUND

Two further problems with the substitution methodology concern the descriptive coverage of the set of relations it eventually produces. One will be discussed in this section, and the other in Section 7.3.

The first problem is simple—many coherent pairs of clauses or sentences can be found for which no cue phrase at all seems appropriate. Given that all the relations in the new set have their basis in the phrases that can signal them, this indicates that the new relations do not by themselves provide a descriptively adequate account of text.

A few examples can be given of contexts where no cue phrase can be used. Consider the following cases:

Dow Associates is one of Britain's largest companies.	{ # *Indeed,* # *Specifically,* # *Furthermore,* # *Incidentally,* }	its head office is in Kensington, where Dow himself presides.	(7.10)
The University is playing host to the 4th VM conference on geochemistry in August.	{ # *Indeed,* # *Specifically,* # *Furthermore,* # *Incidentally,* }	Geochemistry helps in the search for minerals by looking at the origins and natural associations of chemical elements and compounds.	(7.11)

In these texts, it seems impossible to find an appropriate cue phrase to fit between the two sentences. In each case, the best approach is simply to present the sentences one after the other with no explicit cue. No doubt it is texts such as these which led Mann and Thompson to their claim that 'some types of rhetorical relations have *no* corresponding conjunctive signals'.

The lack of a prototypical cue phrase is characteristic of two RST relations in particular: ELABORATION and BACKGROUND. Text 7.10 illustrates the ELABORATION relation. The second sentence provides additional details about Dow Associates, and thus elaborates on the first. Text 7.11 illustrates BACKGROUND. In order properly to understand the first sentence in this case, it is necessary to know what geochemistry is, so this additional information is provided in the second sentence. While it is sometimes hard to distinguish between ELABORATION and BACKGROUND, it seems that neither of them is associated with any cue phrases at all.[1]

[1] Note that RST's ELABORATION can be signalled by relative clauses (Scott and de Souza (1990)):

7.2.1 Re-Assessing the Cue Phrase Hypothesis

The finding that there are some coherent contexts where no cue phrase is appropriate has important consequences for the working hypothesis in this thesis, that a connection can be made between the set of cue phrases and the set of coherence relations. We are forced to choose between the following two possibilities:

- If coherence relations are required to provide a full account of text coherence on their own, then the hypothesis that we can use the set of cue phrases as evidence for the set of relations is false.
- If we wish to maintain the hypothesis, then we have to introduce some other theoretical mechanism alongside relations in order to account for contexts where no cue phrase is appropriate.

The question we need to ask now, of course, is whether there are any principled reasons for proposing different mechanisms for explaining the coherence of texts such as those given above. In the remainder of this section, I will argue that there are such reasons, re-iterating that relations were only intended to model *some aspects* of the phenomenon of coherence, and suggesting that the coherence of texts such as 7.10 and 7.11 is better explained in terms of the metaphor of **focus**. I will argue that the lack of descriptive adequacy shown by the new set of relations may actually be an advantage: it makes for a less redundant account of the interaction of coherence relations and focusing phenomena in text.

7.2.2 The Concept of Focus Revisited

The concept of focus is also of central importance to a theory of coherence. To recap from Section 2.3: the notion of focus (eg Sidner (1983), Grosz, Joshi and Weinstein (1983), Brennan *et al* (1987), McCoy and Cheng (1991)) is used to model those entities in a text's domain of reference which are uppermost in a reader's mind as the text is being read. At any particular moment, certain entities are said to be 'in focus'; focus theories are concerned to chart the constraints on how the focused entities can change from one portion of text to another.

The concept of a **potential focus list** (Sidner (1983)) is commonly invoked in such theories. Every portion of a text is associated with a potential focus list, which specifies all the items to which the focus can legally shift in the next portion. Consider the following three texts:

(7.12) Dow Associates today launched a surprise bid for the beleaguered Harris group. It started buying shares as soon as the market opened, and continued buying all day.

Dow Associates, *whose head office is in Kensington,* is one of Britain's largest companies.

However, relative clauses and cue phrases are two very different kinds of syntactic resource. Spans linked by relative clauses are embedded one within another; those linked by cue phrases are presented consecutively. Furthermore, relative clauses can only operate at low levels of hierarchy within a text; whereas cue phrases such as *this is because* can link quite large text spans.

(7.13) Dow Associates today launched a surprise bid for the beleaguered Harris group. Harris reacted speedily, holding an upbeat press conference.

(7.14) Dow Associates today launched a surprise bid for the beleaguered Harris group. Beans, which contain lots of protein, are good for you.

In Text 7.12, the focused item in the first sentence is *Dow Associates*, and this focus is preserved in the second sentence. Text 7.13 shifts its focus from *Dow Associates* to *Harris*; since *Harris* is in the potential focus list associated with the first sentence, this text is also coherent. However, the potential focus list does not contain the item *beans*—so Text 7.14 is incoherent.

The concept of focus has proved useful primarily for modelling the pattern of anaphora in a text. For instance, it can explain how it is that the focused item in Text 7.12 can be pronominalised, but those in Texts 7.13 and 7.14 cannot be—essentially, an item can only be pronominalised if it is already in focus. Focus theories have been used successfully for resolving anaphora (Sidner, 1983) and for generating it (Dale, 1988, 1988).

7.2.3 Relations and Focus: Two Overlapping Metaphors

In the present context, the important point to note is that an account of focus seems to touch on phenomena which a theory of coherence relations might also be called on to explain. For instance, a theory of focus might equally well be used to explain the coherence of texts such as 7.10 and 7.11, which are presently explained respectively in terms of the ELABORATION and BACKGROUND relations. Conversely, the ELABORATION relation seems appropriate for analysing Texts 7.12 and 7.13, while no relation seems appropriate for the incoherent Text 7.14.

Much of the overlap between theories of relations and of focus can be traced to the relations of ELABORATION and BACKGROUND. The RST definitions for these relations make explicit reference to 'objects' being elaborated, and 'elements' for which further background is required; in this respect, they are unlike any other RST relations. It is thus clear how such relations might be re-described in terms of the focus metaphor.

In fact, the overlap between the focus metaphor and the ELABORATION relation is only partial. The definition for ELABORATION identifies several sub-types of relation, not all of which are elaborations of objects. The sub-relations WHOLE-PART, OBJECT-ATTRIBUTE and SET-MEMBER seem clearly to involve a notion of an entity being elaborated on. However, the sub-relations ABSTRACT-INSTANCE, PROCESS-STEP and GENERALISATION-SPECIFIC elaborate not on entities but on propositions. It should be noted that for these latter species of ELABORATION, cue phrases can readily be found: ABSTRACT-INSTANCE can be signalled by *for instance*; PROCESS-STEP can be signalled by *to do this* or *by*; and GENERALISATION-SPECIFIC can be signalled by *specifically* or *to be precise*.

A clear difference seems to be emerging between those relations which can be marked by cue phrases—which hold between propositions—and those which cannot be marked by cue phrases—which could just as well be thought of as focusing phenomena. The current cue-phrase-based conception of relations thus seems promising as the basis for a principled distinction between relations and focusing phenomena in text.

7.2.4 Why do we need both Relations and Focus?

The previous section has noted the potential for overlap between the metaphors of relations and focus, and consequently for redundancy in any theory in which they both figure. Is it not therefore possible to frame a theory of text solely in terms of one metaphor or the other? This section presents some arguments against such a proposal.

Problems with a Purely Relational Account

Several commentators (e.g. Hovy and McKoy (1989)) have noted that relations by themselves do not provide tight enough constraints on coherence. Consider the text in Figure 7.4. This text can be successfully analysed using relations, but is nonetheless

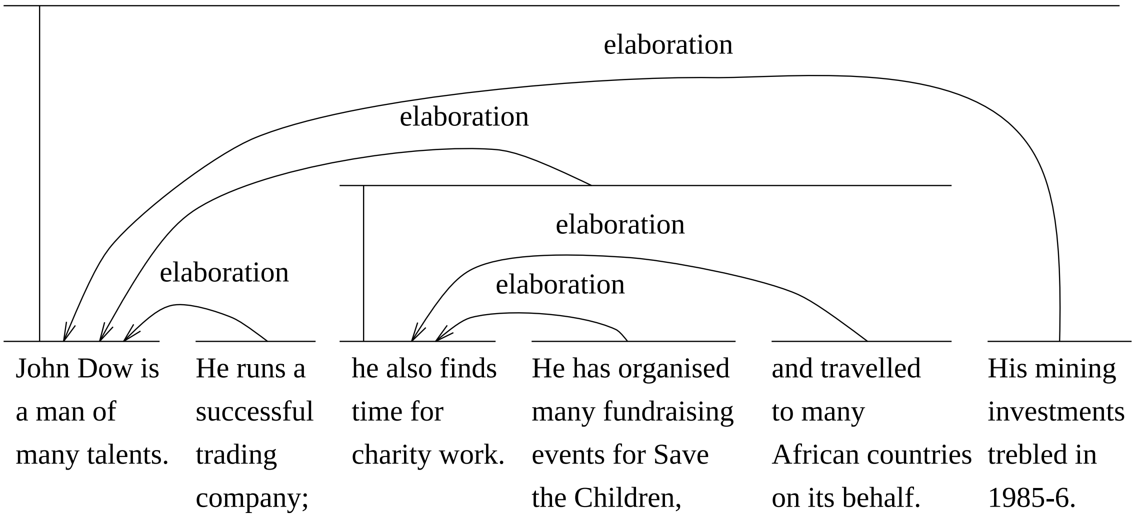

Figure 7.4: A Poorly Structured Text

poorly structured. The problem is with the last sentence, which returns to a topic which had apparently been closed. (It would be more appropriate to include the material from this sentence next to the material about Dow's trading company.) The sentence in its present position can nevertheless be described as an ELABORATION on the first clause—it gives additional information about Dow, which is the only requirement specified by the relation.

It thus seems that the ELABORATION relation is overly permissive. Note, however, that if we are only using relations to analyse texts, then leaving out ELABORATION results in a theory which is overly restrictive. For instance, how could we describe Text 7.15 in relational terms, except by saying that the second sentence elaborates on the first?

(7.15) Dow Associates is one of Britain's largest companies. Its head office is in Kensington.

Given that ELABORATION is the most commonly used relation in text analyses (cf. Mann (1984), p372), we must conclude that there are important aspects of text coherence which do not lend themselves well to description in terms of relations.

Problems with a purely focus-based account.

There are also problems with the focus metaphor when used by itself to model text coherence. Consider Text 7.16:

(7.16) Last year must have been hard for Dow. Oil prices fell by 20%.

This text is clearly coherent, but it seems to contain a sharp focus shift: from *Dow* and *last year* to *oil prices*. In order to account for shifts such as these, focus theorists are forced to abstract away from entities referred to explicitly in a text when creating the potential focus list, and add references to properties or features associated with these items. For instance, Grosz, Joshi and Weinstein allow that entities which are 'elements of the situation' described by an utterance are members of the potential focus list. This would permit them to include *the fall of oil prices* in the list associated with the first sentence of Text 7.16.

However, such abstractions seem to make the focus shifting mechanism itself too permissive: clearly, it will not always be the case that a text's focus can move from *last year* to *oil prices*. It seems more appropriate to explain coherence in a case like Text 7.16 in relational terms, for instance by identifying the second sentence as a JUSTIFICATION for the first. Coherence in this text is due to the rhetorical force of the second sentence, not to any general tendency to shift from one topic to another.

Summary

To sum up thus far: there are some aspects of coherence that seem best described in terms of rhetorical relations, and other aspects that seem best described in terms of focus. However, there is also a significant overlap between the phenomena described by the two metaphors. This leads to what we can term the **redundancy problem**: in a full account of coherence, it is likely that much information is expressed twice.

The redundancy problem is symptomatic of a second, more fundamental problem with relations and focus, which we can call the **grounding problem**. The problem with the two metaphors is that they are too expressive: theorists are free to use them to model as much or as little as they want. What is needed, therefore, are *empirically-grounded* conceptions of relations and focus which allow us to specify *in advance* what is to count as a relation or a focused entity.

7.2.5 Recent Attempts to Link Relations and Focus

Many discourse theorists have appreciated the need to integrate relational and focus-based approaches for a full account of coherence. Two strategies which have been suggested are discussed in this section.

Grosz and Sidner's (1986) theory presents one influential hypothesis—that the relational structure of a text has a role in *determining* shifts in focus. The DOMINANCE relations in a text, which model the hierarchy of the writer's goals, are used during its interpretation to determine pushes and pops of a **focus stack**. The concept of a 'potential focus list' does not feature in this model; its work is done by the intentional structure of the text. Abrupt changes of focus such as that illustrated in Text 7.16 therefore present no problems. However, as has already been noted in Section 2.4.3, the level of detail of this theory does present problems, in particular when it is used as the basis for automatic text processing applications.

Of the recent text planning systems to have incorporated both relational and focusing devices, the most notable are those of McKeown (1985), Hovy and McCoy (1989), and Hovy *et al* (1992). These are all characterised by their use of relations and focus as **multiple simultaneous constraints** on coherence. As a text is built, each new element is fitted into an appropriate rhetorical structure, and is also checked with adjacent elements to ensure that a legal focusing move is made. While this strategy certainly guarantees text conforming to both relational and focusing constraints, it is not guaranteed to be efficient: the large overlap between relations and focus means that, in all likelihood, many constraints are being checked twice.

7.2.6 A New Proposal about the Interaction of Relations and Focus

In order to overcome the redundancy and grounding problems, we need to develop independent conceptions of relations and focus, which model clearly separable aspects of discourse coherence, and which can be empirically motivated.

The new set of relations is promising as the basis for such a model. We have already seen how ELABORATION and BACKGROUND are responsible for much of the overlap between relations and focus, and how ELABORATION is overly permissive in its own right. The new set of relations, in which the exclusion of ELABORATION and BACKGROUND are motivated on independent grounds, seems a good starting point for a joint model of relations and focus.

As was mentioned in Section 2.3, a theory of relations goes hand in hand with a theory of **span structure**, which specifies whereabouts in a coherent text relations are expected to be found. The theory of span structure determines those points in a text where coherence is to be attributed to relations, and consequently, those parts of a text where coherence is to be accounted for by other devices. As was argued in Section 7.2.4, there are good reasons for adopting a theory of span structure which does not force relations to do all the explanatory work.

Based on the present conception of relations, a new suggestion about the interaction of relations and focus can be made. The principal claim is that text coherence where no cue phrase can be used is better explained with the focus metaphor than with the relational metaphor. If one clause elaborates on another, the focus is likely to remain the same, or to shift to some other item that has been explicitly mentioned in the text. A specific hypothesis can be advanced:

H1 Where there is no coherence relation between two text units—i.e. where no cue phrase appears and none can be inserted—coherence is ensured by **focusing constraints**.

It should now be possible to give the term **focusing constraints** a much more concrete interpretation; one that can be grounded in surface textual phenomena. (A simplistic hypothesis could be, for instance, that when no cue phrase is appropriate between two spans, the entity in focus in the second span must appear as a lexical item in the first span. Of course this hypothesis is likely to need refinement, for example to account for phenomena such as bridging references.) We can also propose a hypothesis about how coherence is preserved between text units which are linked by relations:

H2 Focusing constraints are *not* needed to ensure coherence between text units linked by relations.

The main new idea in these hypotheses is that relations and focusing constraints account not just for different aspects of a text, but *different portions* of it. To ensure a text is coherent, it is not necessary that there be a relation between every pair of adjacent text units; but where there is no relation, separate constraints on focus must be met. Conversely, where there is a relation, focusing constraints are not needed to ensure coherence. This idea differs from existing models of coherence, in which relational and focusing constraints apply simultaneously at every point in a text. By using relations and focus to model different portions of text, the new model presents an appealing solution to the redundancy problem. Moreover, since it permits a more concrete conception of focusing constraints, it addresses the grounding problem as well.

There thus seem to be good reasons for proposing an account where relations and focus interact in the way proposed above. However, it should be borne in mind that these are only suggestions, and that much further research is required to explore them in detail. Meanwhile, we should reiterate that the hypothesis proposed in this thesis, of a connection between cue phrases and coherence relations, is contingent on these suggestions proving fruitful.

7.3 Relations at Different Levels of Hierarchy

A final potential problem with the substitution methodology concerns the issue of relations between large sections of text. The idea that relations can apply between text units of any size was noted as one of their attractive features. In many theories of relations, structures as large as entire paragraphs are linked together by relations. The hierarchical, recursive analyses which result from this conception of relations make it attractive from a computational point of view, and also from the standpoint of theoretical parsimony.

It may seem that the decision to associate relations with cue phrases threatens to tie them to text spans of a particular size. For instance, the cue phrase *because* is best suited for joining two clauses within a single sentence:

(7.17) Mary was in a good mood, because she had passed her exam.

While it is possible to use *because* in more complex sentences, it is often inappropriate to use the phrase to link units larger than clauses. Consider Text 7.18:

(7.18) Mary was in a good mood. Because she had passed her exam.

While some such texts may border on acceptability, *because* is not commonly used to link whole sentences, and would probably lead to incoherence if used to link units of several sentences.

Other cue phrases such as *nevertheless* are better able to link whole sentences. However, the problem with such phrases is essentially the same: there is always a limit to the size of spans which they are suitable for joining. For instance, we would not expect to find the word *nevertheless* linking whole chapters of a book, or even sections of an article. How can the current theory of relations deal with relations between large units of this kind?

An answer is suggested here which draws on the fact that cue phrases can contain anaphoric elements.

7.3.1 Cue Phrases and Propositional Anaphora

The test for cue phrases as outlined in Section 4.2 permits phrases which contain propositional anaphora, such as *this is because*, or *following this*. This decision is justified in two ways. Firstly, though the test for cue phrases calls for the replacement of all anaphoric elements by their antecedents, in the case of these phrases the replacement renders the outcome of the test a foregone conclusion—when the antecedent is used, the two clauses originally related are effectively re-expressed in a single clause, which can therefore stand on its own. For instance, consider what happens when the test is applied to the following isolated clause:

(7.19) Because of *this*, Mary was in a good mood.

Replacing *this* by a non-anaphoric NP would result in something of the following form:

(7.20) Because of *the fact that she had passed her exam*, Mary was in a good mood.

Text 7.20, unlike Text 7.19, does not need additional context to be interpreted.

The second reason for allowing anaphora in cue phrases is that many simple connectives have evolved from more complex expressions containing anaphora. As has been noted by Halliday and Hasan (1976), the etymology of words like *therefore* and *thereby* shows that they derive from phrases containing anaphora.

If simple propositional anaphora such as *this* and *that* are used, the resulting cue phrases are still inappropriate for signalling very large relations; readers tend not to look for antecedents beyond the previous sentence. However, if more specific anaphoric expressions are used, antecedents can be further away. Consider the following text:

(7.21) Although the developers had scheduled construction to begin in January, an important archaeological find was made while the foundations were being dug. A dig has quickly been organised to make the most of the find; heading the investigation is Professor W Percus, an authority on Roman Britain.

Because of this discovery, the building has been delayed several months. It is now unlikely to be finished on time...

In this example, the cue phrase *because of this discovery* can be seen as signalling a relation between two whole paragraphs. It should be clear how such expressions could be used to link even larger sections of text.

7.3.2 Different Relations at Different Levels?

The substitutability test deliberately ignores the issue of span size when comparing two phrases. It is possible, for instance, that two phrases are classed as synonymous even though they are suitable for linking spans of different sizes. So it remains an open question whether the set of relations used at low levels of hierarchy is the same as that used for higher levels. For the most part, the examples in this thesis have been of relations between single sentences or single clauses; whether the same set of relations emerges with larger texts is a matter for further investigation.

Nevertheless, some observations can already be made. As was noted in Section 4.3.2, many cue phrases can be systematically modified to include anaphora—for instance, *in contrast* becomes *in contrast to this*; *instead* becomes *instead of this*; *as a result* becomes *as a result of this*.

An interesting link can also be noted between the devices used to signal relations at high levels and those which perform the same function within single clauses. Consider the following texts:

(7.22) Because of this problem, the experiment failed.

(7.23) Because of a small leak in the hydraulic system, the experiment failed.

In Text 7.22, *this problem* refers back to some previously described proposition, and a large relation appears to be signalled. However, in Text 7.23, the problem is being described for the first time—in effect, a causal relation is being signalled inside a single clause. Note that the only difference between the two clauses is whether or not the noun phrase is interpreted as anaphoric.

7.3.3 Relations, Focus and Nominalisation

The use of nominalisation to refer to previous propositions places texts such as 7.22 within the scope of a theory of focus. Are we to interpret this clause as one span of a high-level CAUSE relation, or should we consider it simply as an additional statement about a topic introduced earlier? The fact that an anaphoric expression can be produced, and correctly interpreted, might be taken as support for the focus-based account. However, propositional anaphora are rather different from nominal or verbal anaphora. They introduce a meta-level of description: the propositions in which they appear are not about objects in the domain of reference, but about propositions about these objects.

Nonetheless, it is implausible that all predications about propositions should be interpreted as 'indicators of text structure'. Predications do not have to provide the sort of information given by relations. Consider the following text:

(7.24) Due to a leak in the hydraulic system, the experiment failed. The problem was spotted by a junior lab assistant.

We would not want to suggest that this is an instance of a relation called SPOTTED-BY.

In many cases, therefore, it remains unclear how a joint theory of relations and focus would handle texts containing nominalisations. Perhaps it would be best to leave a decision about such cases until a detailed theory of the interaction of relations and focus has been established for more concrete texts. However, if this theory did not lead to a clear distinction between relational and focusing phenomena in these more abstract cases, this would certainly tell against it.

7.4 Summary

This chapter has discussed in some detail three potential problems with the substitution methodology. Firstly, a limitation with the substitutability test was noted—in some cases it appears to draw on post-theoretical intuitions. For these cases, an alternative method for investigating substitutability relationships was suggested, based on an empirical study of subjects' response times. Secondly, contexts were presented where no cue phrase is appropriate, and hence in the current model no relation is present. It was suggested that these contexts are better explained in terms of the focus metaphor, and the outlines of a joint theory of relations and focus were sketched. Finally, the question of relations between large sections of text was raised. It was partially answered by noting the existence of cue phrases containing propositional anaphora, which are able to link quite large sections of text. The possibility of focusing mechanisms applying at these high levels was also discussed.

Clearly, all of these issues call for further investigation. However, none of the objections raised should yet be seen as fatal to the theory being proposed in this thesis. On the contrary, they all suggest interesting new avenues of research.